Stonesoft 5.5

Administrator's Guide

Security Engines

Management Center

STONESOFT

Legal Information

End-User License Agreement

The use of the products described in these materials is subject to the then current end-user license agreement, which can be found at the Stonesoft website:
www.stonesoft.com/en/support/eula.html

Third Party Licenses

The Stonesoft software includes several open source or third-party software packages. The appropriate software licensing information for those products can be found at the Stonesoft website:
www.stonesoft.com/en/customer_care/support/third_party_licenses.html

U.S. Government Acquisitions

If Licensee is acquiring the Software, including accompanying documentation on behalf of the U.S. Government, the following provisions apply. If the Software is supplied to the Department of Defense (“DoD”), the Software is subject to “Restricted Rights”, as that term is defined in the DOD Supplement to the Federal Acquisition Regulations (“DFAR”) in paragraph 252.227-7013(c) (1). If the Software is supplied to any unit or agency of the United States Government other than DOD, the Government’s rights in the Software will be as defined in paragraph 52.227-19(c) (2) of the Federal Acquisition Regulations (“FAR”). Use, duplication, reproduction or disclosure by the Government is subject to such restrictions or successor provisions.

Product Export Restrictions

The products described in this document are subject to export control under the laws of Finland and the European Council Regulation (EC) N:o 1334/2000 of 22 June 2000 setting up a Community regime for the control of exports of dual-use items and technology (as amended). Thus, the export of this Stonesoft software in any manner is restricted and requires a license by the relevant authorities.

General Terms and Conditions of Support and Maintenance Services

The support and maintenance services for the products described in these materials are provided pursuant to the general terms for support and maintenance services and the related service description, which can be found at the Stonesoft website:
www.stonesoft.com/en/customer_care/support/

Replacement Service

The instructions for replacement service can be found at the Stonesoft website:
www.stonesoft.com/en/customer_care/support/rma/

Hardware Warranty

The appliances described in these materials have a limited hardware warranty. The terms of the hardware warranty can be found at the Stonesoft website:
www.stonesoft.com/en/customer_care/support/warranty_service/

Trademarks and Patents

The products described in these materials are protected by one or more of the following European and US patents: European Patent Nos. 1065844, 1189410, 1231538, 1231754, 1259028, 1271283, 1289183, 1289202, 1304830, 1304849, 1313290, 1326393, 1361724, 1379037, and 1379046 and US Patent Nos. 6,650,621; 6,856,621; 6,912,200; 6,996,573; 7,099,284; 7,127,739; 7,130,266; 7,130,305; 7,146,421; 7,162,737; 7,234,166; 7,260,843; 7,280,540; 7,325,248; 7,360,242; 7,386,525; 7,406,534; 7,461,401; 7,573,823; 7,721,084; and 7,739,727 and may be protected by other EU, US, or other patents, or pending applications. Stonesoft, the Stonesoft logo and StoneGate, are all trademarks or registered trademarks of Stonesoft Corporation. All other trademarks or registered trademarks are property of their respective owners.

Disclaimer

Although every precaution has been taken to prepare these materials, THESE MATERIALS ARE PROVIDED "AS-IS" and Stonesoft makes no warranty to the correctness of information and assumes no responsibility for errors, omissions, or resulting damages from the use of the information contained herein. All IP addresses in these materials were chosen at random and are used for illustrative purposes only.

Revision: SGAG_20130327

TABLE OF CONTENTS

GETTING STARTED

CHAPTER 6

MONITORING

CHAPTER 7

CONTROLLING ENGINES

Chapter 16

Management Center Configuration

Chapter 17

Chapter 18

Chapter 19

Chapter 20

Chapter 21

ENGINE ELEMENT CONFIGURATION

CHAPTER 28

CHAPTER 29

CHAPTER 30

CHAPTER 31

CHAPTER 32

CHAPTER 33

CHAPTER 34

ROUTING

TRAFFIC INSPECTION POLICIES

CHAPTER 41

CHAPTER 42

CHAPTER 43

CHAPTER 44

USERS AND AUTHENTICATION

CHAPTER 52

VIRTUAL PRIVATE NETWORKS

CHAPTER 53

CHAPTER 54

Configuring IPsec VPNs

CHAPTER 55

Managing VPN Certificates

CHAPTER 56

Reconfiguring Existing VPNs

MAINTENANCE AND UPGRADES

TROUBLESHOOTING

REFERENCE

GETTING STARTED

In this section:

CHAPTER 1

USING STONESOFT DOCUMENTATION

This chapter describes how to use this guide and related documentation. It also provides directions for obtaining technical support and giving feedback on the documentation.

The following sections are included:

Using This Documentation

This documentation is intended for Stonesoft Management Center (SMC) administrators. It includes step-by-step instructions for the configuration, operation, and maintenance of the SMC and all of the various components that the SMC controls. Initial system installation is not covered here. For other documentation, see Additional Documentation (page 27).

Typographical Conventions

The following conventions are used throughout the documentation:

Table 1.1 Typographical Conventions

Formatting	Informative Uses
User Interface text	Text you see in the User Interface (buttons, menus, etc.) and any other interaction with the user interface are in **bold-face**.
References, terms	Cross-references and first use of acronyms and terms are in *italics*.
`Command line`	File names, directories, and text displayed on the screen are `monospaced`.
`User input`	User input on screen is in **`monospaced bold-face`**.
`Command parameters`	Command parameter names are in *`monospaced italics`*.

We use the following ways to indicate important or additional information:

Note – Notes prevent commonly-made mistakes by pointing out important points.

Caution – Cautions prevent breaches of security, information loss, or system downtime. Cautions always contain critical information that you must observe.

Tip – Tips provide additional helpful information, such as alternative ways to complete steps.

Example Examples present a concrete scenario that clarifies the points made in the adjacent text.

Prerequisites: Prerequisites point out tasks you must perform before the procedure you are reading. Obvious prerequisites (such as installing a firewall to be able to configure a firewall feature) are not included.

What's Next?

▶ The *What's Next* lists at the ends of tasks guide you to closely related tasks that you must perform in order to configure features. If several of the procedures listed apply, pick the first one; a new *What's Next* list is available at the bottom of the first task.

Additional Documentation

Stonesoft technical documentation is divided into two main categories: product documentation and support documentation.

Product Documentation

The table below lists the available product documentation.

Table 1.2 Product Documentation

Guide	Description
Reference Guide	Explains the operation and features of the Stonesoft system comprehensively. Demonstrates the general workflow and provides example scenarios for each feature area. Available as separate guides for Stonesoft Management Center and Stonesoft Firewall/VPN, and as a combined guide for Stonesoft IPS and Stonesoft Layer 2 Firewall.
Installation Guide	Instructions for planning, installing, and upgrading a Stonesoft system. Available as separate guides for Stonesoft Management Center and Stonesoft Firewall/VPN, and as a combined guide for Stonesoft IPS and Stonesoft Layer 2 Firewall.
Online Help	Describes how to configure and manage the system step-by-step. Accessible through the Help menu and by using the Help button or the F1 key in any window or dialog. Available in the Stonesoft Management Client and the Stonesoft Web Portal. An HTML-based system is available in the Stonesoft SSL VPN Administrator through help links and icons.
Administrator’s Guide	Describes how to configure and manage the system step-by-step. Available as a combined guide for Stonesoft Firewall/VPN, Stonesoft IPS, and Stonesoft Layer 2 Firewall, and as separate guides for Stonesoft SSL VPN and Stonesoft IPsec VPN Client.
User’s Guide	Instructions for end-users. Available for the Stonesoft IPsec VPN Client and the Stonesoft Web Portal.
Appliance Installation Guide	Instructions for physically installing and maintaining Stonesoft appliances (rack mounting, cabling, etc.). Available for all Stonesoft hardware appliances.

PDF guides are available at http://www.stonesoft.com/en/customer_care/documentation/current/. The *Stonesoft Administrator’s Guide*, and the *Reference Guides* and *Installation Guides* for Stonesoft Management Center, Stonesoft Firewall/VPN, Stonesoft IPS, and Stonesoft Layer 2 Firewall are also available as PDFs on the Management Center DVD.

Using Online Help Locally

By default, the Management Client's Online Help is accessed through the Internet at the Stonesoft help server. Alternatively, you can configure the Management Client to use a copy of the Online Help from a local machine (for example, an intranet server or your own computer). This way the Online Help is available even when there is no Internet connectivity. If you want to use the Online Help locally, it is highly recommended to save the local copy on an intranet server. The Online Help is context-sensitive only if it is used from a server.

Note – If you use a local copy of the Online Help, you must manually update the Online Help when a new version of the Online Help becomes available.

▼ To use the Online Help locally

1. Go to the product documentation pages of the Stonesoft website at http://www.stonesoft.com/en/customer_care/documentation/current/.
2. Download the Management Client's Online Help as a `.zip` file to a suitable location on the local machine.
3. Extract the Online Help `.zip` file to a suitable location.
4. Browse to the `$USER_HOME/.stonegate/data` folder.
5. Open the `SGClientConfiguration.txt` file.
6. Add a new parameter *HELP_SERVER_URL* and enter a path to the main folder under which the Online Help files are stored as the value for the parameter:
   - Enter **`http://`** and the path if you have saved the Online Help on an intranet server.
   - Enter **`file:///`** and the path if you have saved the Online Help in the local file system.

Example **If you have extracted the Online Help to a folder called** `SMC_Online_Help/5.3.2` **on an intranet server, enter** **`HELP_SERVER_URL=http://`*`<intranet.server>`*`/SMC_Online_Help/5.3.2`.**

7. Save the changes to the `SGClientConfiguration.txt` file and wait a few seconds. The Management Client automatically starts using the Online Help from the specified folder after a short while.

Support Documentation

The support documentation provides additional and late-breaking technical information. These technical documents support the guide books, for example, by giving further examples on specific configuration scenarios.

The latest technical documentation is available on the Stonesoft website at http://www.stonesoft.com/support/.

System Requirements

The certified platforms for running Stonesoft engine software can be found at the product pages at http://www.stonesoft.com/en/products/appliances/.

The hardware and software requirements for the Management Center and version-specific details for all software products can be found in the *Release Notes* available on the software download page at the Stonesoft website. The Management Center Release Notes are also included on the Management Center DVD.

Supported Features

Not all features are supported on all platforms. See the Appliance Software Support Table at the Stonesoft Support Documentation pages for more information.

Contact Information

For street addresses, phone numbers, and general information about Stonesoft products and Stonesoft Corporation, visit our website at http://www.stonesoft.com/.

Licensing Issues

You can view your current licenses at the License Center section of the Stonesoft website at https://my.stonesoft.com/managelicense.do.

For license-related queries, e-mail order@stonesoft.com.

Technical Support

Stonesoft offers global technical support services for Stonesoft's product families. For more information on technical support, visit the Support section at the Stonesoft website at http://www.stonesoft.com/support/.

Your Comments

We want to make our products fulfill your needs as well as possible. We are always pleased to receive any suggestions you may have for improvements.

- To comment on software and hardware products, e-mail feedback@stonesoft.com.
- To comment on the documentation, e-mail documentation@stonesoft.com.

Security Related Questions and Comments

You can send any questions or comments relating to the Security Engine inspection features and network security to security-alert@stonesoft.com.

Other Queries

For queries regarding other matters, e-mail info@stonesoft.com.

CHAPTER 2

NEW IN THIS RELEASE

This section lists major changes since the previous release. Most new or reworked features in the software are listed here. Changes that do not significantly affect the way the Stonesoft components are configured are not listed. For a full list of changes in the software and detailed version-specific information, consult the Release Notes.

The following sections are included:

Other Changes in SMC 5.5

Changes to Domains

A Domain is active on the Management Server that controls the Domain. Each Domain is active on only one Management Server at a time. In an environment with more than one Management Server, it is possible to activate different Domains on different Management Servers.

- For more details, see Getting Started with Domains (page 284).

Changes to Management Server High Availability

There is no longer the concept of active and standby Management Servers. Some commands and configuration changes must be done on the Management Server that controls the Shared Domain. However, all Management Servers are always active, and changes to the database are replicated incrementally between the Management Servers.

- For more details, see About Additional SMC Servers (page 330).

Changes to the Tasks Branches

The Running Tasks and Executed Tasks branches have been merged into a single History branch. There are now only two branches for Tasks: Definition and History. The Definition branch shows custom Task Definitions and predefined System Task Definitions. By default, the History branch shows Tasks that are currently running. Optionally, you can also view Executed Tasks in the History branch. System Tasks are not shown in the History branch unless the Show Executed Tasks option is enabled.

- For more details, see Getting Started with Tasks (page 1032).

Increased Security Strength of Management Connections

It is now possible to use 256-bit encryption for the connection between Security Engines and the Management Server. This requires both the engines and the Management Server to be version 5.5 or higher. You must also use an Internal ECDSA Certificate Authority to sign certificates for system communication.

- For more details, see Enabling 256-bit Security Strength for Engines (page 234).

Internal ECDSA Certificate Authority

The Elliptic Curve Digital Signature Algorithm (ECDSA) is a digital signature algorithm that uses elliptic curve cryptography. It is now possible to use an Internal ECDSA Certificate Authority to sign the certificates that components use to identify each other in system communications.

Caution – When you start using an ECDSA Certificate Authority, you must recertify all SMC servers and you may also need to make initial contact between the engines and the Management Server. Engines with versions lower than 5.5 and SSL VPN gateways cannot communicate with the SMC when an ECDSA Certificate Authority is used.

- For more details see Creating a New Internal ECDSA Certificate Authority (page 348).

Obsolete Elements

In SMC versions lower than 5.5, deleted elements were sent to the Trash branch in environments that had multiple Management Servers. In version 5.5, elements to be deleted are first marked as obsolete and they are shown in the Obsolete Elements branch from which they can be deleted permanently. The Obsolete Elements branch is displayed regardless of the number of Management Servers.

- For more details, see Marking Elements as Obsolete (page 88).

Changes in Firewall/VPN 5.5

Loopback Interface Configuration

A Loopback IP address allows the firewall to communicate with itself. It is now possible to configure any IP address IP address that is not already used as a Cluster Virtual IP Addresses (CVI) or Node Dedicated IP Addresses (NDI) on another interface as a loopback IP address. You can add several loopback IP addresses to each Firewall. Loopback IP addresses can also be used as End-Point IP addresses in policy-based VPNs and in the Route-Based VPN.

- For more details, see Configuring Loopback IP Addresses for Firewalls (page 450).

Master Engines and Virtual Security Engines

Virtual Security Engines are logically-separate engines that run as virtual machines on a physical engine device. It is now possible to use a physical Security Engine device as a Master Engine to provide resources for Virtual Security Engines. Virtual Security Engines do not require individual licenses. Instead, the Security Engine license for the Master Engine defines how many Virtual Resources can be created. The number of Virtual Resources limits the number of Virtual Security Engines. Virtual Security Engines can currently be used only in the Firewall role.

- For more details, see Getting Started with Virtual Engines (page 404).

QoS Enhancements

Multiple enhancements have been made to the current bandwidth management and traffic prioritization features.

The new QoS Mode option in the properties of a Physical, VLAN, ADSL, Tunnel, or SSID Interface allows you define in more detail how QoS is applied to the interface.

- It is now possible to read and/or write DSCP markers for traffic without configuring Access rules to apply a QoS class to the traffic. The matching is done based only on the QoS Policy.
- QoS Class-based statistics items are now available even when QoS is not used for bandwidth management and traffic prioritization. The QoS class for the packet comes from the QoS Classes that are applied in the Access rules.

For VPN traffic, the DSCP mark for the ESP packet is normally inherited from the plaintext packet. It is now also possible to mark the ESP packet after encryption. Simple limits can now be applied to VPN traffic. Because the total throughput is undefined, Guarantees and Priorities cannot be used.

New Active Queue Management (AQM) features reduce the volume of dropped or retransmitted packets when there is network congestion. AQM monitors the average queue size and uses a scheduling algorithm to determine the statistical probability for dropping incoming packets.

It is now possible to assign a weight to QoS Classes so that different QoS Classes with the same priority can be assigned to the queue according to their weight when the QoS Class Guarantee is reached and traffic must be queued. This allows more granular control of traffic prioritization, but does not act as a guarantee.

- For more details, see Getting Started with QoS (page 782).

Changes in IPS and Layer 2 Firewalls 5.5

QoS Enhancements

Multiple enhancements have been made to the current bandwidth management and traffic prioritization features.

The new QoS Mode option in the properties of Physical interfaces that are part of an Inline interface pair on IPS engines and Layer 2 Firewalls allows you define in more detail how QoS is applied to the interface.

- It is now possible to read and/or write DSCP markers for traffic without configuring Access rules to apply a QoS class to the traffic. The matching is done based only on the QoS Policy.
- QoS Class-based statistics items are now available even when QoS is not used for bandwidth management and traffic prioritization. The QoS class for the packet comes from the QoS Classes that are applied in the Access rules.

New Active Queue Management (AQM) features reduce the volume of dropped or retransmitted packets when there is network congestion. AQM monitors the average queue size and uses a scheduling algorithm to determine the statistical probability for dropping incoming packets.

It is now possible to assign a weight to QoS Classes so that different QoS Classes with the same priority can be assigned to the queue according to their weight when the QoS Class Guarantee is reached and traffic must be queued. This allows more granular control of traffic prioritization, but does not act as a guarantee.

- For more details, see Getting Started with QoS (page 782).

Chapter 3

Using the Management Client

The sections listed below provide an introduction to the Stonesoft Management Client.

The following sections are included:

Overview to the Management Client

Prerequisites: None

This section presents a general overview to the Management Client. For detailed information on using these views, see the Related Tasks at the end of this section.

The Management Client offers several task-specific views. There are alternative ways to switch between the different views:

- The main menu and the toolbar shortcuts are always available.
- Additional links are provided, for example, in the right-click menus of elements and in the logs. You can also bookmark your most frequently visited views.

You have several options for opening a new view:

- Clicking replaces the current view with the new one.
- Ctrl-clicking opens the new view as a new tab.
- Shift-clicking opens the new view as a new window.

You can open a new empty tab in any view, by using the keyboard shortcut Ctrl+T. You can also do it by clicking the plus icon on the right of the previous tab. From the list of views that opens, select the view to be shown in the new tab.

The System Status View

Illustration 3.1 The System Status View

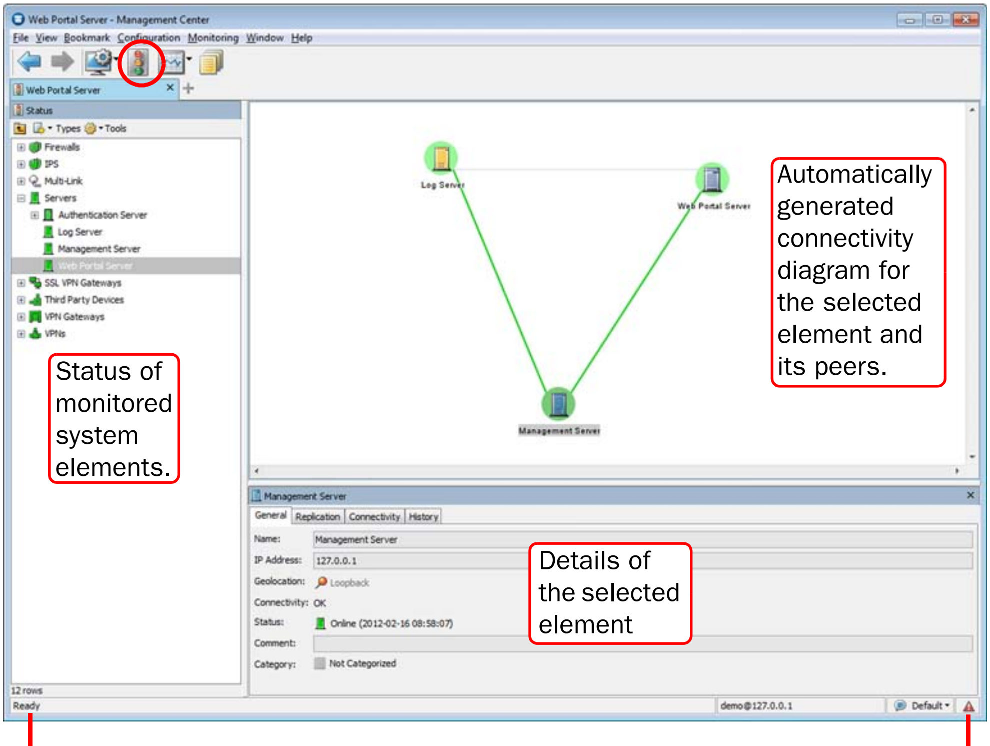

Name: Management Server

IP Address: 127.0.0.1

Geolocation: Loopback

Connectivity: OK

Status: Online (2012-02-16 08:58:07)

Comment:

Category: Not Categorized

Automatically generated connectivity diagram for the selected element and its peers.

Status of monitored system elements.

Details of the selected element

Information messages of the Management Client itself.

Active alert notification.

By default, when you launch the Management Client, you see the System Status view. This view provides the operating and connectivity status of system components and third party components that are set up to be monitored through the Stonesoft Management Center. The right-click menus of the elements offer shortcuts to many configuration and maintenance tasks.

The Configuration Views

Illustration 3.2 The Security Engine Configuration View

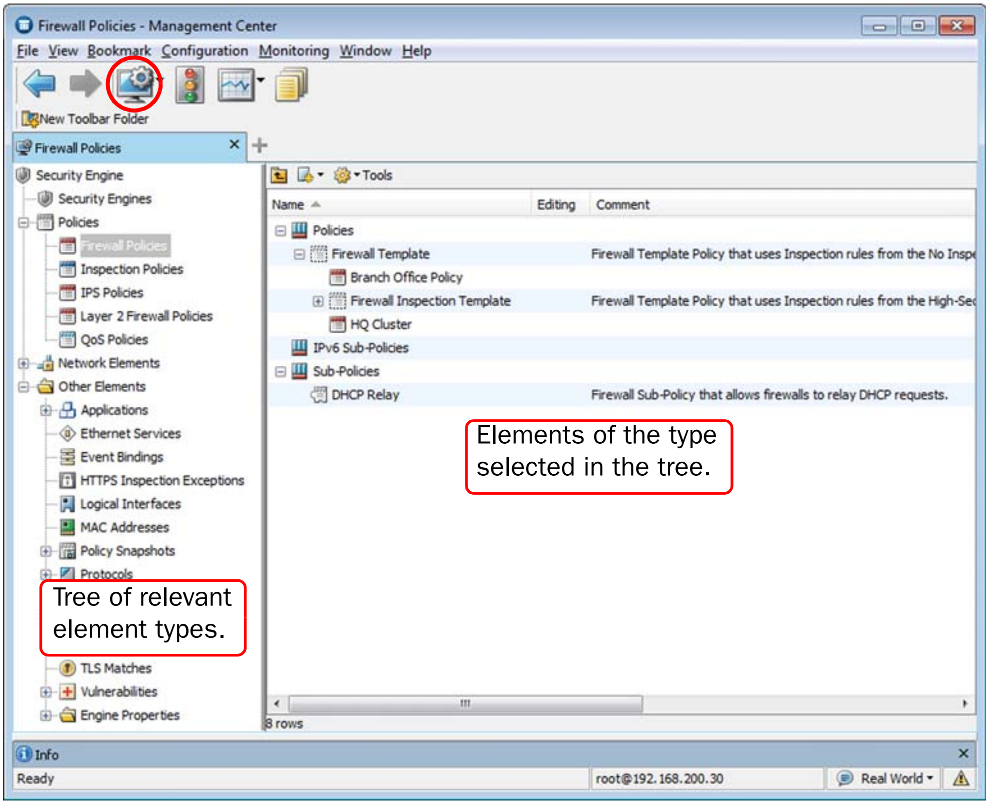

Name	Editing	Comment
Policies		
Firewall Template		Firewall Template Policy that uses Inspection rules from the No Inspe
Branch Office Policy		
Firewall Inspection Template		Firewall Template Policy that uses Inspection rules from the High-Sec
HQ Cluster		
IPv6 Sub-Policies		
Sub-Policies		
DHCP Relay		Firewall Sub-Policy that allows firewalls to relay DHCP requests.

Elements of the type selected in the tree.

Tree of relevant element types.

There are separate configuration views for Security Engine, User Authentication, VPN, Administration, and Monitoring configuration tasks. These views can be opened both through the main menu and the shortcut buttons in the toolbar.

The Configuration views allow you to set up and modify all features of the system. The configurations are stored as elements, which are shown in a tree structure. Elements are created and modified through the right-click menus that open when you right-click a tree branch or an element.

Overviews

Illustration 3.3 An Example Overview

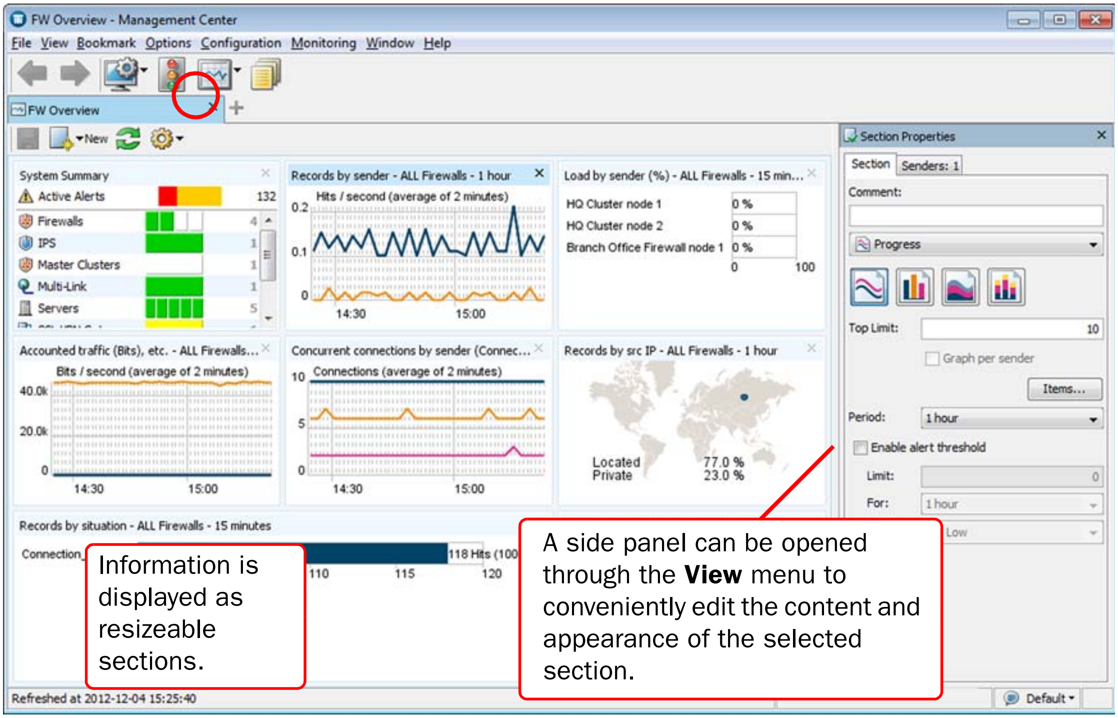

FW Overview - Management Center
File View Bookmark Options Configuration Monitoring Window Help
FW Overview
New
System Summary
Active Alerts 132
Firewalls 4
IPS 1
Master Clusters 1
Multi-Link 1
Servers 5
Records by sender - ALL Firewalls - 1 hour
Hits / second (average of 2 minutes)
0.2
0.1
0
14:30
15:00
Load by sender (%) - ALL Firewalls - 15 min...
HQ Cluster node 1 0 %
HQ Cluster node 2 0 %
Branch Office Firewall node 1 0 %
0
100
Accounted traffic (Bits), etc. - ALL Firewalls...
Bits / second (average of 2 minutes)
40.0k
20.0k
0
14:30
15:00
Concurrent connections by sender (Connec...
Connections (average of 2 minutes)
10
5
0
14:30
15:00
Records by src IP - ALL Firewalls - 1 hour
Located 77.0 %
Private 23.0 %
Records by situation - ALL Firewalls - 15 minutes
Connection
118 Hits (100
110
115
120
Section Properties
Section
Senders: 1
Comment:
Progress
Top Limit: 10
Graph per sender
Items...
Period: 1 hour
Enable alert threshold
Limit: 0
For: 1 hour
Low
Refreshed at 2012-12-04 15:25:40
Default

Information is displayed as resizeable sections.

A side panel can be opened through the **View** menu to conveniently edit the content and appearance of the selected section.

Overviews can contain information on the system's status, bookmarks to views you use often (such as logs filtered with specific criteria), and statistical charts on the system's operation (such as engine load) and the traffic flow. You can create several different overviews for different purposes. Several default overviews are provided as templates.

Logs View

Illustration 3.4 The Logs View in the Records Arrangement

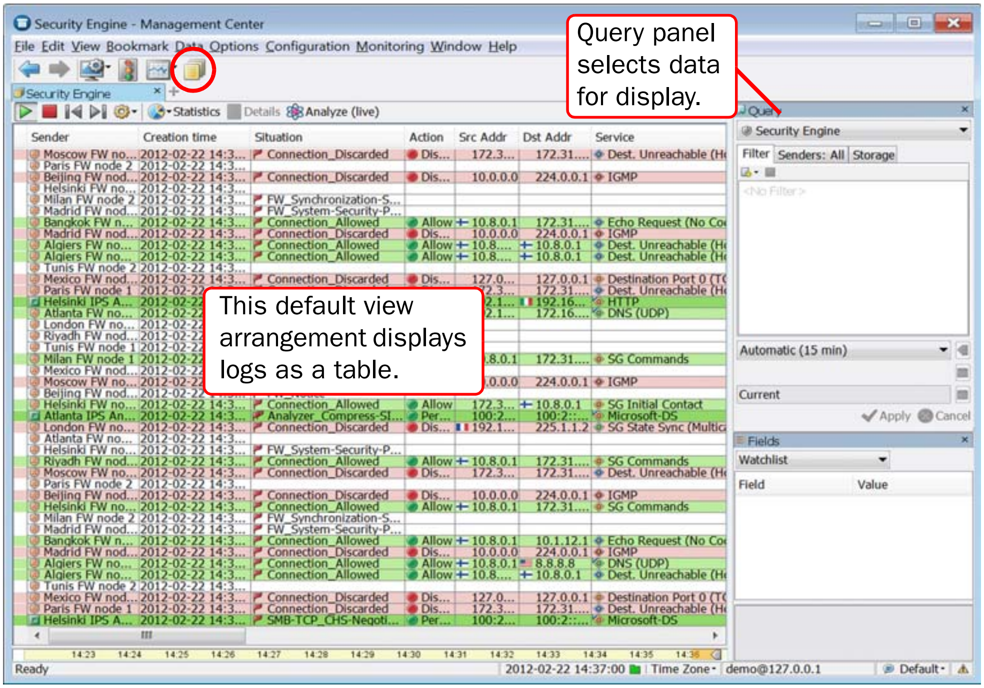

Sender	Creation time	Situation	Action	Src Addr	Dst Addr	Service
Moscow FW no...	2012-02-22 14:3...	Connection_Discarded	Dis...	172.3...	172.31....	Dest. Unreachable (H
Paris FW node 2	2012-02-22 14:3...					
Beijing FW nod...	2012-02-22 14:3...	Connection_Discarded	Dis...	10.0.0.0	224.0.0.1	IGMP
Helsinki FW no...	2012-02-22 14:3...					
Milan FW node 2	2012-02-22 14:3...	FW_Synchronization-S...				
Madrid FW nod...	2012-02-22 14:3...	FW_System-Security-P...				
Bangkok FW n...	2012-02-22 14:3...	Connection_Allowed	Allow	10.8.0.1	172.31....	Echo Request (No Co
Madrid FW nod...	2012-02-22 14:3...	Connection_Discarded	Dis...	10.0.0.0	224.0.0.1	IGMP
Algiers FW no...	2012-02-22 14:3...	Connection_Allowed	Allow	10.8....	10.8.0.1	Dest. Unreachable (H
Algiers FW no...	2012-02-22 14:3...	Connection_Allowed	Allow	10.8....	10.8.0.1	Dest. Unreachable (H
Tunis FW node 2	2012-02-22 14:3...					
Mexico FW nod...	2012-02-22 14:3...	Connection_Discarded	Dis...	127.0...	127.0.0.1	Destination Port 0 (T
Paris FW node 1	2012-02-2			2.3...	172.31....	Dest. Unreachable (H
Helsinki IPS A...	2012-02-2			2.1...	192.16...	HTTP
Atlanta FW no...	2012-02-2			2.1...	172.16....	DNS (UDP)
London FW no...	2012-02-2					
Riyadh FW nod...	2012-02-2					
Tunis FW node 1	2012-02-2					
Milan FW node 1	2012-02-2			.8.0.1	172.31....	SG Commands
Mexico FW nod...	2012-02-2					
Moscow FW no...	2012-02-2			.0.0.0	224.0.0.1	IGMP
Beijing FW nod...	2012-02-2					
Helsinki FW no...	2012-02-22 14:3...	Connection_Allowed	Allow	172.3...	10.8.0.1	SG Initial Contact
Atlanta IPS An...	2012-02-22 14:3...	Analyzer_Compress-SI...	Per...	100:2...	100:2::...	Microsoft-DS
London FW no...	2012-02-22 14:3...	Connection_Discarded	Dis...	192.1...	225.1.1.2	SG State Sync (Multic
Atlanta FW no...	2012-02-22 14:3...					
Helsinki FW no...	2012-02-22 14:3...	FW_System-Security-P...				
Riyadh FW nod...	2012-02-22 14:3...	Connection_Allowed	Allow	10.8.0.1	172.31....	SG Commands
Moscow FW no...	2012-02-22 14:3...	Connection_Discarded	Dis...	172.3...	172.31....	Dest. Unreachable (H
Paris FW node 2	2012-02-22 14:3...					
Beijing FW nod...	2012-02-22 14:3...	Connection_Discarded	Dis...	10.0.0.0	224.0.0.1	IGMP
Helsinki FW no...	2012-02-22 14:3...	Connection_Allowed	Allow	10.8.0.1	172.31....	SG Commands
Milan FW node 2	2012-02-22 14:3...	FW_Synchronization-S...				
Madrid FW nod...	2012-02-22 14:3...	FW_System-Security-P...				
Bangkok FW n...	2012-02-22 14:3...	Connection_Allowed	Allow	10.8.0.1	10.1.12.1	Echo Request (No Co
Madrid FW nod...	2012-02-22 14:3...	Connection_Discarded	Dis...	10.0.0.0	224.0.0.1	IGMP
Algiers FW no...	2012-02-22 14:3...	Connection_Allowed	Allow	10.8.0.1	8.8.8.8	DNS (UDP)
Algiers FW no...	2012-02-22 14:3...	Connection_Allowed	Allow	10.8....	10.8.0.1	Dest. Unreachable (H
Tunis FW node 2	2012-02-22 14:3...					
Mexico FW nod...	2012-02-22 14:3...	Connection_Discarded	Dis...	127.0...	127.0.0.1	Destination Port 0 (T
Paris FW node 1	2012-02-22 14:3...	Connection_Discarded	Dis...	172.3...	172.31....	Dest. Unreachable (H
Helsinki IPS A...	2012-02-22 14:3...	SMB-TCP_CHS-Negoti...	Per...	100:2...	100:2::...	Microsoft-DS

Query panel selects data for display.

This default view arrangement displays logs as a table.

Related Tasks

- Bookmarking Views (page 41)
- Changing the Startup View (page 44)
- Managing Elements (page 75)
- Getting Started with System Monitoring (page 94)
- Creating Overviews (page 104)
- Getting Started with the Logs View (page 148)

Rearranging the General Layout

Prerequisites: None

You can select different panels and view options through the **View** menu. Some layout options are specific to the currently active view, and some options are global.

You can drag and drop the panels to several preconfigured places within each specific view. The layout is saved as your preference for further use, but you can reset it through the **View→Layout→Reset Layout** option in the main menu.

- To *resize* a panel, drag by the outer edge of the panel as usual when resizing.
- To *move* a panel, drag by the title bar at the top like you would move a window (see illustration below).

Illustration 3.5 Positioning a Panel

You can move the panels in a number of positions that are highlighted as you drag the panel around. Drop the panel where you prefer to have it.

If the highlighted area completely covers some other panel, the second panel adds a new tab.

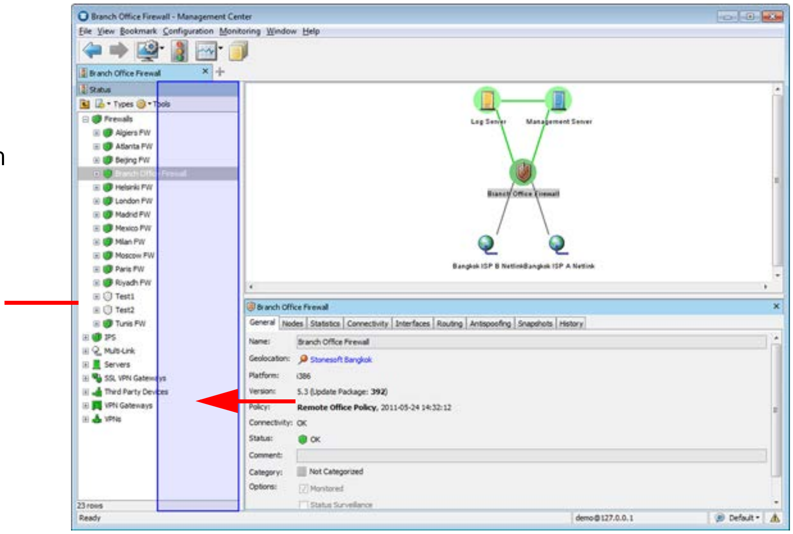

Name: Branch Office Firewall
Geolocation: Stonesoft Bangkok
Platform: i386
Version: 5.3 (Update Package: **392**)
Policy: **Remote Office Policy**, 2011-05-24 14:32:12
Connectivity: OK
Status: OK
Comment:
Category: Not Categorized
Options: ☑ Monitored
☐ Status Surveillance

You can also bookmark alternative layouts to quickly return to a specific view and layout at any later time. See Bookmarking Views (page 41).

Bookmarking Views

Prerequisites: None

Bookmark-related actions can be found through the **Bookmark** menu. Bookmarks can also be added to the toolbar. Bookmarks can store many of the settings you have selected in views. For example, you can bookmark a Logs view with particular filtering settings. Several windows and tabs can be stored behind a single click when you combine bookmarks into *Bookmark Folders*.

Bookmarks in the default **Shared Bookmarks** folder are shown to all administrators that log in to the same Management Server. Other bookmarks are private to the Management Clients of individual administrators.

Managing Bookmarks

Bookmarks can be managed in the Administration Configuration view under the **Bookmarks** branch or by selecting **Bookmark→Manage Bookmarks**.

You can, for example, copy Bookmarks and Bookmark Folders from one group to another. All actions are available through the right-click menu for the Bookmarks and Bookmark folders.

Illustration 3.6 Bookmarks

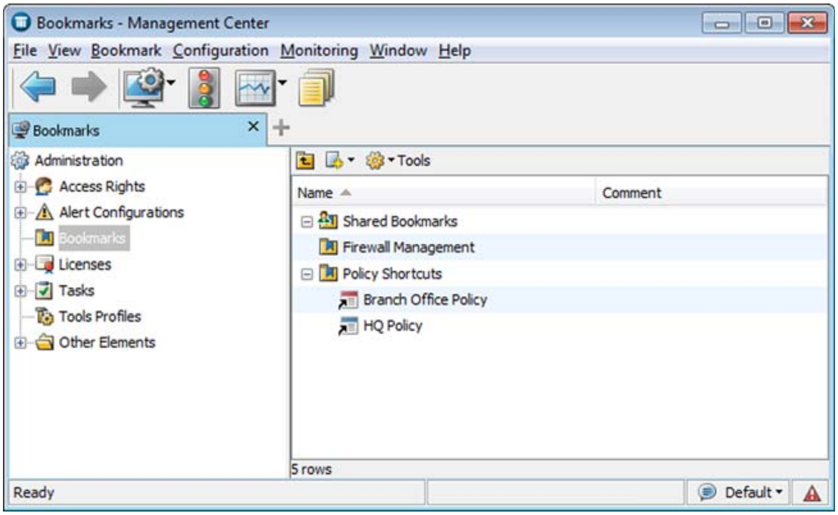

Name	Comment
Shared Bookmarks	
Firewall Management	
Policy Shortcuts	
Branch Office Policy	
HQ Policy	

Creating New Bookmarks

You can create a bookmark for the currently active tab and other window-level elements in the configuration you select. Some view-specific options are also stored in bookmarks, such as the currently active filter in the Logs view, or the type of elements that are listed in a Configuration view at the time the bookmark is created.

Bookmarking is a main window action, so the properties dialogs for the various elements are never included in the bookmark.

▼ To bookmark the active view

1. Arrange the view as you would like to see it when the bookmark is opened.
2. Select **Bookmark→Add Bookmark**. The Bookmark Properties dialog opens.

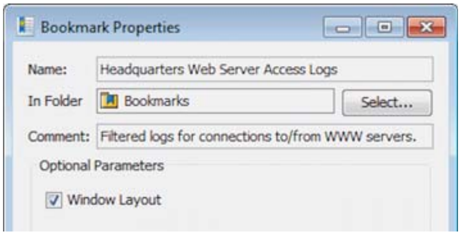

Name: Headquarters Web Server Access Logs

In Folder: Bookmarks

Comment: Filtered logs for connections to/from WWW servers.

Optional Parameters

☑ Window Layout

3. (*Optional*) Change the **Name** and add a **Comment** for your reference.
   - The default name is taken from the bookmarked view's title.
4. (*Optional*) Next to the **In Folder** field, click the **Select** button and select the folder where the Bookmark is placed.
   - The default **Bookmarks** creates the folder at the top level of the bookmarks tree.
   - Select the **Shared Bookmarks** folder if you want other administrators to see this Bookmark. All other folders are private to your Management Client.
   - Select the **Toolbar** folder or one of its sub-folders to add the bookmark folder to the toolbar. If the **Toolbar** folder is not available, activate the feature as explained in Adding Bookmarks to the Toolbar (page 43).
5. (*Optional*) Deselect **Window Layout** if you prefer the bookmark to not change the layout of the window when you open it.
6. Click **OK**.

▼ To bookmark all tabs in the window

1. Open the tabs you want to bookmark in the same window (and close any tabs you do not want to bookmark).
2. Select **Bookmark→Bookmark All Tabs**. The Bookmark Folder Properties dialog opens.
3. Fill in the Bookmark Folder properties (see Creating New Bookmark Folders) and click **OK**. A new Bookmark Folder is created, containing a bookmark for each tab you had open.

Related Tasks

▶ To further edit and organize the bookmarks, see Managing Bookmarks (page 41).

Creating New Bookmark Folders

Bookmark folders organize bookmarks and make it easier to open several bookmarks at once. The folders you create are also added as items in the **Bookmark** menu.

As an alternative to the steps explained below, you can bookmark all tabs that are open in the current view in a bookmark folder that is automatically created to contain the new bookmarks. See To bookmark all tabs in the window (page 42) for detailed instructions.

▼ To create a new bookmark folder

1. Select **Bookmark→Manage Bookmarks**. The Bookmarks tree is displayed.
2. Right-click **Bookmarks** in the tree and select **New Bookmark Folder**. The Bookmark Folder Properties dialog opens.
3. Enter a **Name** and optionally also a **Comment** for your reference.
4. (*Optional*) For **In Folder**, click the **Select** button and select the folder where the Bookmark is placed.
   - The default **Bookmarks** creates the folder at the top level of the bookmarks tree.
   - Select the **Shared Bookmarks** folder if you want other administrators to see this Bookmark. All other folders are private to your Management Client.
   - Select the **Toolbar** folder or one of its sub-folders to add the bookmark folder to the toolbar. If the **Toolbar** folder is not available, activate the feature as explained in Adding Bookmarks to the Toolbar.
5. Click **OK**.

Adding Bookmarks to the Toolbar

You can add your bookmarks to the toolbar below the shortcut icons. After you activate the bookmarks toolbar as explained below, you can add any new bookmark to the toolbar by storing the bookmark in the **Toolbar** folder or one of its sub-folders. You can also move any existing bookmark to the toolbar by dragging and dropping the bookmark or bookmark folder to the **Toolbar** folder in the **Bookmark→Manage Bookmarks** tree.

▼ To activate the bookmarks toolbar

1. Select **View→Layout→Bookmark Toolbar**. The bookmark toolbar is shown below the toolbar icons.

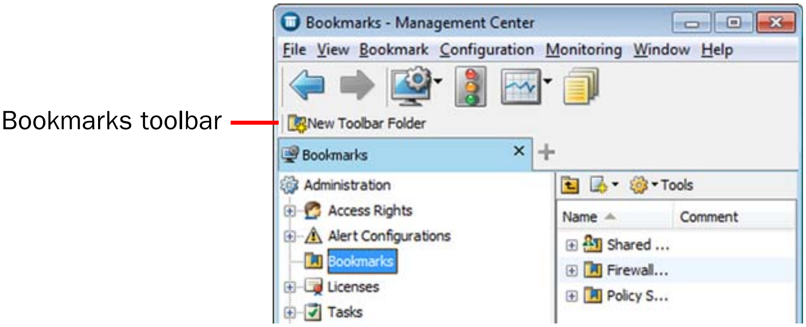

2. Click the default **New Toolbar Folder** item. The Create Toolbar Folder dialog opens.
3. Enter the name for the first folder to add to the toolbar and click **OK**. The first folder appears in the toolbar and the **Toolbar** folder is added to the bookmark hierarchy, allowing you to add, remove, and edit the bookmarks in the toolbar.

Changing the Startup View

Prerequisites: None

You can freely choose which view opens whenever you log in to the Management Client, replacing the default System Status view.

▼ To change the startup view

1. Arrange the screen with the windows, tabs, and views you want to have open at each login.
2. Select **Bookmark→Save as Startup Session**.

Using the Search Features

Prerequisites: None

Using Basic Element Search

You can use the basic search to quickly find elements by a simple element property, such as name, comment, or IP address.

▼ To search for elements

1. Select **View→Search→Search**. The basic element search opens.

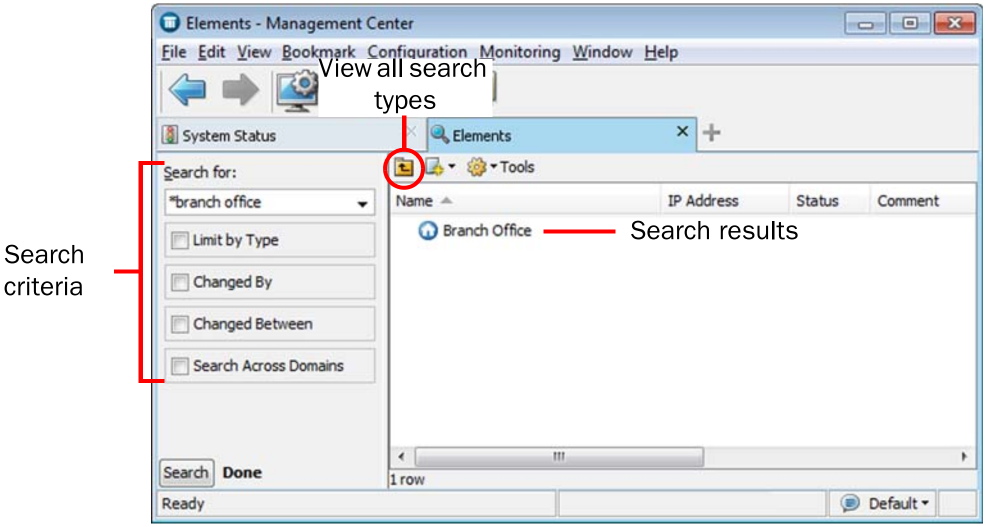

View all search types

Search criteria

Search results

2. Enter the search criteria (name, IP address, comment, vulnerability reference, etc.) in the **Search for** field.
   - You can use wildcard symbols in your query (for example, win*, or 192.168*100).
   - You can use the + symbol to require search criteria.
   - You can use the - symbol to exclude search criteria.

3. (*Optional*) Select options to limit the search:

Option	Explanation
Limit by Type	Restrict the search to one type of element.
Changed by	Restrict the search to elements modified by a particular administrator. Select the administrator from the list of recently used elements or click **Other** to see all Administrator elements.
Changed Between	Search elements that were last changed between certain dates. You can type in the dates or select them from a calendar.
Search Across Domains	Search within all configured administrative Domains. This option is only visible if Domain elements have been configured.

Illustration 3.7 Search Panel

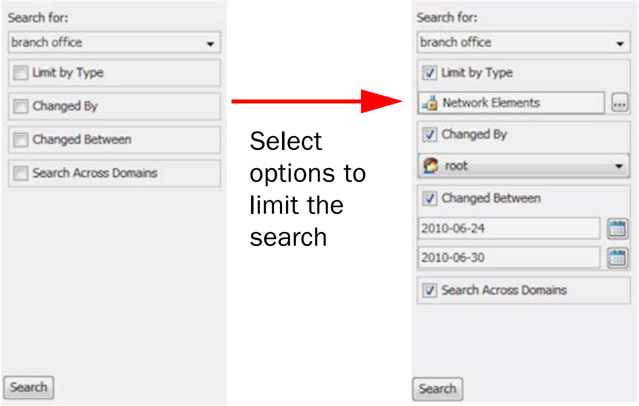

Search for: branch office

□ Limit by Type

□ Changed By

□ Changed Between

□ Search Across Domains

Search

Select options to limit the search

Search for: branch office

☑ Limit by Type

Network Elements

☑ Changed By

root

☑ Changed Between

2010-06-24

2010-06-30

☑ Search Across Domains

Search

4. Click the **Search** button to launch the search. The search results are displayed.

Tip – If the element you searched for does not exist, you can right-click the empty search results and create a new Host or Network that has the Name and/or IP address configured according to your search terms.

Related Tasks

- Searching for Element References (page 46)
- Using the DNS Search (page 46)
- Searching for Duplicate IP Addresses (page 47)
- Searching for Unused Elements (page 47)
- Searching for Users (page 48)
- Searching for Obsolete Elements (page 49)
- Using Type-Ahead Search (page 50)

Searching for Element References

If you need to check where an element is used, you can quickly search for references to it. This is useful, for example, when you want to delete an element (as referenced elements cannot be deleted until the references are removed).

▼ To search for element references

1. Right-click the element and select **Tools→References**. A References view opens and shows the elements that reference the element you right-clicked.
2. (*Optional*) Select **Show References** in the Search panel if you also want to view further references for the elements that are found and click the **Search** button to relaunch the search. The search results are now displayed as a tree that allows you to follow the trail of references.

Tip – Click the Up icon in the toolbar to view a list of all available searches.

Using the DNS Search

You can use the DNS search to look for hosts by their DNS name. The Management Server must have Zone Transfer rights to be able to make queries on the DNS server.

▼ To use the DNS search

1. Select **View→Search→Search DNS**. The DNS view opens.
2. Enter the **DNS Server IP Address** and the **DNS Domain**.
3. Click the **Search** button to launch the search. The search results display a list of available host names and their IP addresses.
   - If a Host element that has a matching name is already in the system, the Host element is shown in the Network Element column.

What's Next?

▶ If the search did not find an element, you can add it based on the search as explained in Creating Host Elements Based on DNS Queries.

Creating Host Elements Based on DNS Queries

If the Network Element column in the DNS search results does not show an existing Host element for a host name, you can create a new Host element based on the host name and IP address.

▼ To create a Host element based on DNS query results

1. Right-click one or more IP addresses in the DNS view and select **Create Network Element**. The Host Creation dialog opens.
   - Ctrl-click or Shift-click to select multiple IP addresses if you want to create Host elements for all of them.
2. (*Optional*) If you want to add the new Host element(s) to a Group, select **Group** and enter the **Group Name**.
   - If such a Group element does not yet exist in the system, a new Group is automatically created.

3. Define a name for the new Host element(s):
   - **Full Name**: the Host name is the full name retrieved from the DNS.
   - **Name without Domain**: the Host name is created by removing the domain part from the name retrieved from the DNS.
   - **Specify Name**: Enter a name for the Host(s).
   - **Prefix + IP Address**: Enter a Prefix that is added to IP address to form the Host name(s).
4. Click **OK**. The created Host(s) are displayed in the Network Element column.

Searching for Duplicate IP Addresses

▼ To search for duplicate IP addresses

1. Select **View→Search→Search Duplicate IPs**. The Duplicate IPs view opens.
2. Select the search criteria as described in the table below:

Option	Explanation
Limit by Type	Restrict the search to one type of element.
Changed by	Restrict the search to elements modified by a particular administrator. Select the administrator from the list of recently used elements or click Other to see all Administrator elements.
Changed Between	Search elements that were last changed between certain dates. You can type in the dates or select them from a calendar.
Search Across Domains	Search within all configured administrative Domains. This options is only visible if Domain elements have been configured.

3. Click the **Search** button to launch the search. The search results are displayed.

Searching for Unused Elements

The unused element search helps you find elements that may be unnecessary. The search looks for element references and lists elements that are not referenced by any other elements. You must use your own judgement to determine which of the listed elements can safely be removed.

▼ To search for unused elements

1. Select **View→Search→Search Unused Elements**. The Unused Elements view opens.
2. Select the search criteria as described in the table below:

Table 3.1 Search Options

Option	Explanation
Limit by Type	Restrict the search to one type of element.
Changed by	Restrict the search to elements modified by a particular administrator. Select the administrator from the list of recently used elements or click Other to see all Administrator elements.

Table 3.1 Search Options (Continued)

Option	Explanation
Changed Between	Restrict the search to elements changed between certain dates. Enter the date information or click the button next to the date fields to select the dates from a calendar.
Search Across Domains	Search within all configured administrative Domains. This option is only visible if Domain elements have been configured.
Ignore references to Categories	Select this to include unused elements which have been included in categories. If you do not select this, unused elements which have been included in a category are left out of the search results.

3. Click the **Search** button to launch the search. A list of elements that are not used in any configuration or within any elements is displayed and the number of the elements found is shown at the bottom of the Search panel.

Searching for Users

▼ To search for users

1. Select **View→Search→Search Users**. The Users view opens.
2. Select the LDAP Domain in which you want to search for the user(s).
3. Select the user attributes for the search as described in the table below:

Option	Explanation
UserID	User’s login name.
Name	The name of the user in the User properties.
Authentication Service	The authentication method configured for the user.
Activation Date	The first date (yyyy-mm-dd) on which the user is allowed to log in through the firewall.
Expiration Date	The last date (yyyy-mm-dd) on which the user is allowed to log in through the firewall.
Expiration Delay	The number of days the user is allowed to log in through the firewall (from the Activation Date to the Expiration Date).

4. Specify the Query type as described in the table below:

Option	Explanation
Contains	The user information contains the value (for example, a name) that is searched for.
Greater Than	The Activation Date or Expiration Date of the user is the same or later than the date entered for the search (yyyy-mm-dd).
Less Than	The Activation Date or Expiration Date of the user is the same or earlier than the date entered for the search (yyyy-mm-dd).
Is Defined	Returns all user accounts that have some value for the selected attribute.

5. (*Depending on attribute selected*) Enter the required value (for example, a name or date) that you want to search for among the user information.
6. Click the **Search** button. A list of users that match the search criteria is displayed.

Searching for Obsolete Elements

You cannot delete elements directly. You must first mark the elements as obsolete, after which you can delete the elements if there are no references to them. For more information on Obsolete Elements, see Marking Elements as Obsolete (page 88).

▼ To search Obsolete Elements

1. Select **View→Search→Search Obsolete Elements**. The Obsolete Elements view opens.
2. Select the search criteria as described in the table below.

Option	Explanation
Limit by Type	Restrict the search to one type of element.
Changed by	Restrict the search to elements modified by a particular administrator. Select the administrator from the list of recently used elements or click Other to see all Administrator elements.
Changed Between	Search elements that were last changed between certain dates. You can type in the dates or select them from a calendar.
Search Across Domains	Search within all configured administrative Domains. This options is only visible if Domain elements have been configured.

3. Click the **Search** button to launch the search. A list of obsolete elements that correspond to the search criteria is shown in the Search panel on the right.

Using Type-Ahead Search

Prerequisites: None

Type-ahead search allows you to quickly filter elements. You can use type-ahead search in all views in which element lists, tables, or trees are displayed. You can also use type-ahead search in rule cells in Policies.

When you start typing, the system immediately filters the display to the elements that include what you typed in the element's basic details, for example:

- Element name.
- IP address.
- Comment.
- TCP/UDP Port.

Illustration 3.8 Type-Ahead Search

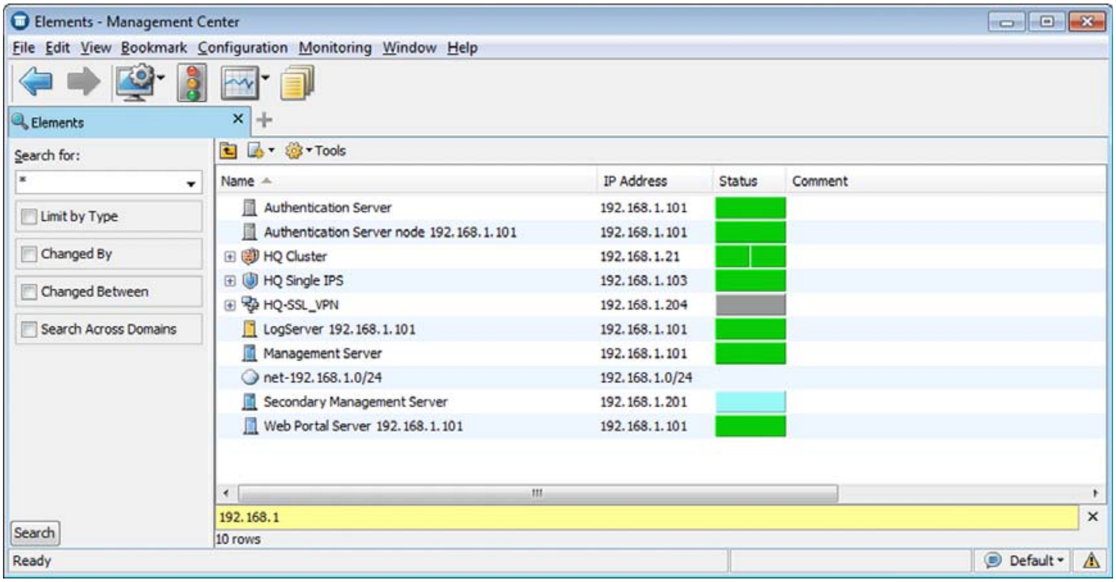
Elements - Management Center
File Edit View Bookmark Configuration Monitoring Window Help
Elements
Tools
Limit by Type
Changed By
Changed Between
Search Across Domains
Search
Default

Search for: *

Name	IP Address	Status	Comment
Authentication Server	192.168.1.101	🟩	
Authentication Server node 192.168.1.101	192.168.1.101	🟩	
HQ Cluster	192.168.1.21	🟩🟩	
HQ Single IPS	192.168.1.103	🟩	
HQ-SSL_VPN	192.168.1.204	⬛	
LogServer 192.168.1.101	192.168.1.101	🟩	
Management Server	192.168.1.101	🟩	
net-192.168.1.0/24	192.168.1.0/24		
Secondary Management Server	192.168.1.201	🟦	
Web Portal Server 192.168.1.101	192.168.1.101	🟩	

192.168.1

10 rows

Ready

Saving Elements, Log Data, Reports, and Statistics

Prerequisites: None

You can save lists of elements, logged data, reports, and statistics in PDF format or as HTML by selecting **File→Print**. Saving as PDF allows you to use a background template. The specifics of PDF printing are described below.

PDF Output Settings

Logged data, reports, and statistics can be saved as a PDF file. The associated settings are explained in the table below.

Table 3.2 Print to PDF Settings

Tab	Section	Settings	Description
General	Printing Options	Print to PDF Reader	Opens the output file in the default PDF viewing application on your computer.
		Print to File	Allows you to save the output as a PDF file.
	Page Setup	Paper Size	Defines the background paper size. The available sizes are A3, A4, A5, and Letter.
		Orientation	Select either a vertical (portrait) or horizontal (landscape) page.
	Records (*Logs view only*)	Style Template	Defines the background for the generated PDF. See Adding Style Templates for PDF Output (page 52).
		Selected Records	Prints only records that are currently selected in the Logs view.
		Filtered Records from...	Prints all records that match the current Query.
Layout (*Logs view only*)		Table	Prints a table with one record per row. Enter the number of columns that are included.
		Records	Prints a table in which each record occupies several rows in a column layout. Enter the number of columns you want to include (counting from left to right).
		Details	Prints all relevant details according to your selection using one page per record.

Adding Style Templates for PDF Output

Background style templates can be used when saving as PDF. The templates can be used in the Management Client and the Web Portal as permitted by account permissions and Domain boundaries.

The style template is a standard PDF file you create with some or all of the following elements:

- One or more cover pages that are attached to the printed PDF before the content pages.
- A content page background with a header and/or footer. The height of the header and footer can be adjusted. The same single page is used as the background for all content pages.
- One or more trailing pages that are attached to the printed PDF after the content pages.
- Your PDF template file can contain additional pages that you do not want to be used; these are ignored.
- A one-page PDF file is used as a content page. Your PDF template file must contain at least two pages if you want to add cover and trailing pages.

You can create the template, for example, by creating a suitable document in a word processor and saving it as a PDF. For best results, design separate templates for the different page sizes (A3, A4, A5, or Letter) and orientations (portrait or landscape) you anticipate using.

▼ To add a new Style Template

1. Create a PDF file that contains a template page for the content and optionally one or more cover and trailing pages.
2. Open the Print to PDF dialog, for example, by right-clicking a log entry and selecting **Print**. The Print to dialog opens.
3. Under Page Setup, select **New** from the **Style Template** list. The Style Template Properties dialog opens.
4. Enter a unique **Name** for the new Style Template.
5. Click **Browse** and select the PDF file you want to use as a template.
6. Select how the pages are used:
   - (*Optional*) The **Front Pages from** are inserted before the content pages without modifications. Fill in just the first field for a single-sheet front page.
   - The **Content Page** is used as the background for all pages that have system-generated content.
   - The **Header Height** and **Footer Height** define how much space is left (in millimeters) between the top and bottom of the page and the first/last line of system-generated content. This prevents the generated content from overlapping text/graphics on your content page.
   - (*Optional*) The **Back Pages from** are inserted after the content pages without modifications. Fill in just the first field for a single-sheet trailing page.

Tip – You can use the same pages for different roles. For example, you can select the same page as a content page and a back page to add an empty page at the end of the generated PDFs. The PDF template file must have at least 2 pages to do this, even if you only use one of the pages.

7. Click **OK**.

Managing PDF Style Templates

▼ To manage Style Templates

1. Open the Print to dialog, for example, by right-clicking a log entry and selecting **Print**.
2. Under Page Setup, select **Other** from the **Style Template** list. The Select dialog opens.
3. Right-click a Style Template and select an action from the right-click menu. You can **Delete** the selected Style Template or select **Properties** to adjust the template settings (see Adding Style Templates for PDF Output (page 52) for information on the settings).

Sending Messages to Other Administrators

Prerequisites: None

The administrator messaging feature allows you to communicate with other administrators who are currently logged in to the Management Client. For example, you can inform administrators about service breaks before upgrading the Management Center. Administrator messaging is enabled by default.

What's Next?

- If you want to enable/disable administrator messaging, proceed to Enabling/Disabling Administrator Messaging.
- To send messages to administrators, proceed to Sending Messages to Other Administrators (page 54).

Enabling/Disabling Administrator Messaging

Only an administrator with unrestricted permissions can enable/disable administrator messaging. If Domain elements have been configured, the setting is applied in all the Domains.

▼ To enable/disable administrator messaging

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Expand the **Access Rights** branch in the Administration tree.
3. Right-click **Administrators** and select/deselect **Administrator Messaging Enabled**.

Sending Messages to Other Administrators

Administrators who are logged in can send messages to all other administrators. Only administrators who have the Manage Administrators permission can send messages to individual administrators. Each administrator must be logged in to a unique administrator account for individual messages to be sent.

▼ To send a new message to other administrators

1. Click the Send Message button at the bottom right of the Management Client window to open the Conversation Properties dialog.

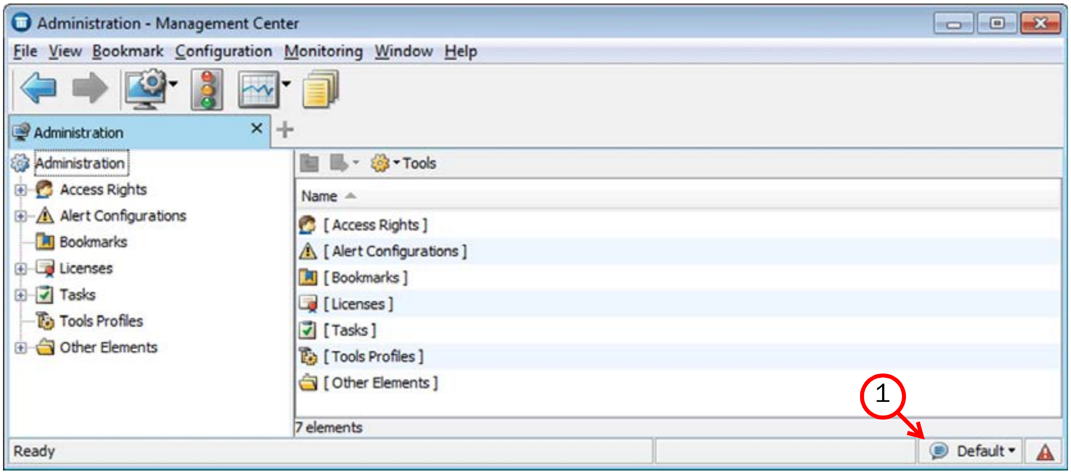

2. Select the **Administrators** who receive the message:
   - (*Administrators with the Manage Administrators permission*) Click the **Select** icon (the icon with the black arrow) to select individual Administrator(s).
   - Click the **Set All Administrators** icon (the “list” icon) to send the message to all administrators.
3. Type in your **Message** and click **Send**. The message is sent to all selected administrators. The Administrator Messaging dialog automatically opens for each administrator so that they can reply to your message.

Adding Custom Commands to Element Menus

Prerequisites: None

You can add commands (for example, *tracert, SSH,* or *ping*) to an element's right-click menu with Tools Profiles. The commands are added in the **Tools** submenu.

What's Next?

- To define new commands, proceed to Creating a Tools Profile.
- To reuse an existing Tools Profile without modifications, proceed to Attaching a Tools Profile to an Element (page 56).

Creating a Tools Profile

Tools Profiles add commands to the right-click menus of elements. You can include information dynamically from the element definition in the command. Only one Tools Profile can be selected for each element, but each Tools Profile can include several commands.

The commands are launched on the workstation that is running the Management Client. Commands are operating-system-specific, so you must add a separate command for each operating system in use. Administrators see commands according to their operating system (for example, a Linux command is not shown if the Management Client is running in Windows).

▼ To create a Tools Profile

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Right-click **Tools Profile** and select **New Tools Profile**. The Tools Profile Properties dialog opens.

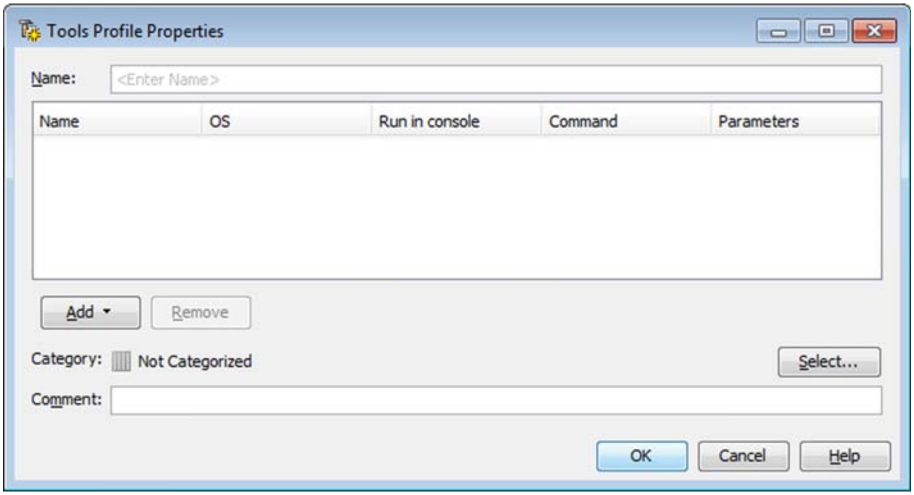

Tools Profile Properties

Name: <Enter Name>

Name	OS	Run in console	Command	Parameters

3. Enter the **Name** for the new Tools Profile. This is added as an item under the Tools submenu of elements that the Tools Profile is attached to.
4. Click **Add** and select the operating system. A row is added to the table.
5. Double-click the **Name** cell, and enter the item to add.

6. *(Optional)* Select **Run in Console** if you want the command to be run in a console application, such as the command prompt in Windows or terminal in Linux.
7. Double-click the **Command** cell and enter the command or the full path to the application.
8. *(Optional)* Double-click the **Parameters** cell and enter the parameters for the command. In addition to static parameters, the following two variables can be used:
   - **`${IP}`**: the primary IP address of the element that is right-clicked.
   - **`${NAME}`**: the name of the element that is right-clicked.
9. Click **OK**.

What's Next?

▶ Proceed to Attaching a Tools Profile to an Element.

Attaching a Tools Profile to an Element

You can attach a Tools Profile to Network Elements that have a single primary IP address.

▼ To attach a Tools Profile to an element

1. Right-click an element and select **Properties**. The element's Properties dialog opens.
2. Select the Tools Profile in one of the following ways:
   - Select a Tools Profile from the list.
   - Select **Other** and browse to the Tools Profile.
   - Select **New** and create a new Tools Profile as explained in Creating a Tools Profile (page 55).
3. Click **OK**. The commands defined in the Tools Profile are added to the right-click menu of the element in the **Tools** submenu.

CHAPTER 4

SETTING UP THE SYSTEM

The sections listed below provide an introduction to the SMC and shortcuts to feature-specific instructions.

The following sections are included:

Getting Started with the Management Center

Prerequisites: Management Center installation (covered in a separate *Installation Guide*)

This section is meant to help you get started after you have completed the Management Center installation.

To familiarize yourself with the Management Client, see Overview to the Management Client (page 36).

The basic administration tasks you must complete after installation include the following:

- Scheduling automatic backup tasks to keep safe the essential configuration information stored on the Management Server as instructed in Backing Up and Restoring System Configurations (page 1007).
- Setting up automated tasks to manage the gathered log data and prevent the Log Server storage space from filling up with logs as instructed in Managing Log Data (page 1015).

Additionally, we highly recommend that you set up the following features:

- Defining additional administrator accounts and delegating administrative tasks as instructed in Administrator Accounts (page 247).
- Reviewing settings for automatic updates and making sure the feature works to ensure that your system stays current. See Getting Started with Automatic Updates and Engine Upgrades (page 242).
- Defining custom alerts and alert escalation policies as instructed in Alert Escalation (page 265).

To efficiently manage the system, you must also familiarize yourself with the following basic tasks:

- Monitoring the system operation as instructed in Monitoring the System (page 93) and Browsing Logged Data (page 147).

If you have installed the optional Authentication Server, you must complete the configuration to be able to use the Authentication Server's authentication services. See Integrating Authentication Server Services (page 861).

Getting Started with the Firewall

Prerequisites: Management Center and Firewall installation (covered in separate *Installation Guides*)

This section is meant to help you get started after you have completed the installation, installed a basic policy, and turned the Firewalls online as instructed in the *Installation Guide*. If you have also installed a new Management Center, first see Getting Started with the Management Center (page 58).

Note – The configuration information is stored on the Management Server, and most changes are transferred to the engines only when you install or refresh the Firewall Policy after making the changes.

The basic administration tasks you must learn or complete next include the following:

- Reading and controlling the operating state of Firewall engines as explained in Controlling Engine Operation (page 217).
- Adjusting the automatic tester that monitors the operation of the Firewall and the surrounding network as explained in Configuring the Engine Tester (page 519).
- Developing your Firewall Policies further as instructed in Getting Started with Policies (page 646).

The most typical customization steps include:

- Configuring multiple network connections for load-balanced, highly available networking as explained in Getting Started with Outbound Traffic Management (page 610).
- Configuring traffic management for incoming connections to groups of servers as explained in Getting Started with Inbound Traffic Management (page 618).
- Setting up bandwidth management and traffic prioritization policies as explained in Getting Started with QoS (page 782).
- Configuring the firewall to utilize external content inspection servers as explained in Getting Started with External Content Inspection (page 810).
- Configuring secure connectivity between different locations and for travelling users as explained in Getting Started With IPsec VPNs (page 918).

Also consult the product-specific *Reference Guides*, which contain background information that helps you better understand the system and its features.

Getting Started with the IPS

Prerequisites: Management Center and IPS installation (covered in separate *Installation Guides*)

This section is meant to help you get started after you have completed the installation, installed a basic policy, and turned the IPS system online as instructed in the *Installation Guide*. If you have also installed a new Management Center, first see Getting Started with the Management Center (page 58).

Note – The configuration information is stored on the Management Server, and most changes are transferred to the engines only when you install or refresh the IPS policy after making the changes.

The basic administration tasks you must learn or complete next include the following:

- How to read and control the operating state of IPS engines as explained in Controlling Engine Operation (page 217).
- Adjusting the automatic tester that monitors the operation of the IPS engines and the surrounding network as explained in Configuring the Engine Tester (page 519).

After you have installed your first IPS policy, your next task is gathering information about the events detected in your networks during a “tuning period” (see Browsing Logged Data (page 147)). Once you have enough information on what kind of traffic—malicious and harmless—can be seen in your network, you can modify your policies to improve the detection accuracy and to get rid of false alarms. The most typical customization steps include:

- Creating your own policy or policy template as explained in Creating and Managing Policy Elements (page 645).
- Modifying the Ethernet rules, Access rules, and Inspection rules as explained in Editing Ethernet Rules (page 668), Editing Access Rules (page 672), and Editing Inspection Policies (page 699).
- Creating your own custom Situations as explained in Defining Situations (page 755).

Also consult the product-specific *Reference Guides*, which contain background information that helps you better understand the system and its features.

Getting Started with the Layer 2 Firewall

Prerequisites: Management Center and Layer 2 Firewall installation (covered in separate *Installation Guides*)

This section is meant to help you get started after you have completed the installation, installed a basic policy, and turned the Layer 2 Firewall system online as instructed in the *Installation Guide*. If you have also installed a new Management Center, first see Getting Started with the Management Center (page 58).

Note – The configuration information is stored on the Management Server, and most changes are transferred to the engines only when you install or refresh the Layer 2 Firewall Policy after making the changes.

The basic administration tasks you must learn or complete next include the following:

- How to read and control the operating state of the Layer 2 Firewalls as explained in Controlling Engine Operation (page 217).
- Adjusting the automatic tester that monitors the operation of the Layer 2 Firewalls and the surrounding network as explained in Configuring the Engine Tester (page 519).

After you have installed your first Layer 2 Firewall Policy, your next task is gathering information about the events detected in your networks during a "tuning period" (see Browsing Logged Data (page 147)). Once you have enough information on what kind of traffic—malicious and harmless—can be seen in your network, you can modify your policies to improve the detection accuracy and to get rid of false alarms. The most typical customization steps include:

- Creating your own policy or policy template as explained in Creating and Managing Policy Elements (page 645).
- Modifying the Ethernet rules, Access rules, and Inspection rules as explained in Editing Ethernet Rules (page 668), Editing Access Rules (page 672), and Editing Inspection Policies (page 699).
- Creating your own custom Situations as explained in Defining Situations (page 755).

Also consult the product-specific *Reference Guides*, which contain background information that helps you better understand the system and its features.

CHAPTER 5

CONFIGURING SYSTEM COMMUNICATIONS

This section provides an overview to communications between system components.

The following sections are included:

Getting Started with System Communications

Prerequisites: None

System communications involve traffic between Stonesoft components and traffic between Stonesoft components and external components that are a part of the system configuration.

Firewalls and Layer 2 Firewalls do not automatically allow any system communications. The predefined Firewall Template Policy and Layer 2 Firewall Template Policy contain rules that allow most types of system communications between the Firewall or Layer 2 Firewall and the components it interacts with. You must create additional rules to allow any other communications through Firewalls or Layer 2 Firewalls, such as the communications of some other system component that pass through the Firewall or Layer 2 Firewall. A list of ports used in system communications is presented in Default Communication Ports (page 1167).

System Communications Through a NAT Device

If NAT is applied between two system components, you must define the translated IP address as a *contact address*. A single component can have several contact addresses. Only IPv4 addresses are used in system communications.

Location elements define when a contact address is used and which of the defined contact addresses is used. When NAT is applied between two communicating system components, you must separate them into different Locations. Components that are in the same Location use the primary IP address when communicating with each other and ignore all contact addresses. When components contact a component that belongs to a different Location, they use the defined contact address.

There is one system element related to contact addresses, the **Default** Location. If you do not select a Location for an element that has the Location option, the element’s Location is set to Default.

You can define a Default contact address for contacting a component (defined in the Properties dialog of the element). The element’s Default contact address is used in communications when components that belong to another Location contact the element and the element has no contact address defined for its Location.

Example of Using Contact Addresses and Locations

Illustration 5.1 Example Network With Locations

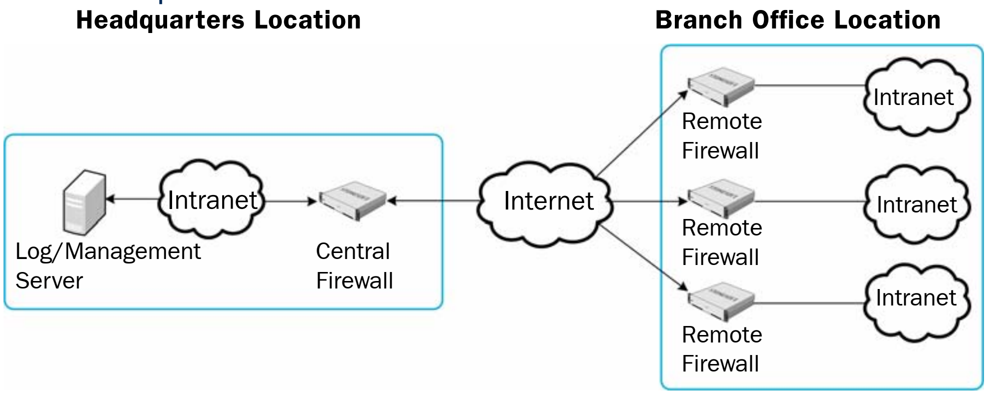

In the example scenario above, a Management Server and a Log Server manage system components both at a company's headquarters and at three branch offices.

- The SMC servers and the Central Firewall are at the "Headquarters" Location.
- The Remote Firewalls are all at the "Branch Office" Location.

In this scenario, contact addresses are typically needed as follows:

- The Firewall at the headquarters or an external router may provide the SMC servers external IP addresses on the Internet. The external addresses of the SMC servers must be defined as contact addresses for the "Branch Office" location, so that the components at the branch offices can contact the servers across the Internet.
- The Branch Office Firewall or an external router may provide external addresses for the system components at the branch office. The external IP addresses of the engines must be defined as contact addresses for the "Headquarters" Location so that the Management Server can contact the components.
- Alternatively, the external address of each component can be defined as a Default contact address without adding a specific entry for "Headquarters" or "Branch Office". The Default contact address is used when a component does not have a specific contact address definition for the contacting component's Location. Note that the components must still be divided into separate Locations for the contact address to be used.

If there are Management Clients used at any of the branch offices, each administrator must also select "Branch Office" as their Location in the Management Client to be able to view logs from the remote Log Server that is behind a NAT device.

What's Next?

- ▶ Proceed to Defining Locations (page 66) to start defining contact addresses.
- ▶ Proceed to Defining Contact IP Addresses (page 66) to add a new contact address to an existing Location.

Related Tasks

- ▶ Defining IP Addresses (page 719)
- ▶ Network Interface Configuration (page 419)
- ▶ Selecting the Management Client Location (page 73)

Defining Locations

Prerequisites: None

If NAT (network address translation) is applied between communicating system components, the components must be assigned to different Locations in the configuration. You create the Locations and add elements to them based on how your network is set up.

Example **If a system has several Locations, but each component still has the same external IP address no matter where the contact is made from, each element only needs a single contact address: the Default contact address. When new system elements are added, they have to be assigned a specific Location, but they only need a Default contact address.**

▼ To create a new Location

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Expand the **Other Elements** branch in the Administration tree.
3. Right-click **Locations** and select **New Location**. The Location Properties dialog opens.
4. Enter a **Name** and an optional **Comment**.
5. Browse to the type of elements you want to assign to the Location in the Resources panel.
6. Select the element(s) and click **Add**. The selected elements are added to the Content panel on the right.
7. Click **OK**.

What's Next?

- ▶ Continue by Defining Contact IP Addresses.

Defining Contact IP Addresses

Prerequisites: Defining Locations

You can define contact addresses for Firewalls, IPS engines, and Layer 2 Firewalls, and most types of server elements. The contact addresses are defined directly in the element properties. The contact addresses are based on Location elements. You can also define a Default contact address that is used whenever no contact address is defined for a certain Location.

What's Next?

- ▶ If you need to define a contact address for an engine, start by Defining Engine Location.
- ▶ If you need to define contact addresses for a server, proceed to Defining Server Contact Addresses (page 70).
- ▶ If you need to define a contact address for an external VPN gateway, see Defining a Contact Address for External Security Gateway End-Point (page 72).
- ▶ If you need to define a contact address for an internal VPN gateway, see Defining End-Points for Internal Security Gateways (page 924).

Defining Engine Location

You must define a Location and a contact address for an engine if NAT is applied to the communications between the engine and some other component that needs to contact the engine. If you use Multi-Link, remember to add contact addresses for each NetLink.

▼ To define a Location for an engine element

1. Right-click the engine element and select **Properties** from the menu. The Properties dialog for the element opens.

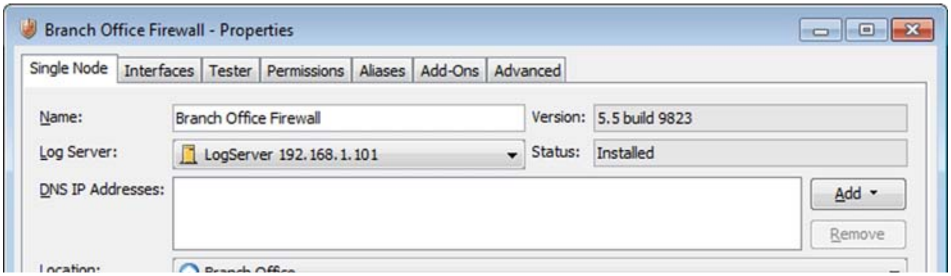

Name: Branch Office Firewall

Version: 5.5 build 9823

Log Server: LogServer 192.168.1.101

Status: Installed

DNS IP Addresses:

Location: Branch Office

2. Select the **Location** for this element.
3. Switch to the **Interfaces** tab in the engine element's properties.

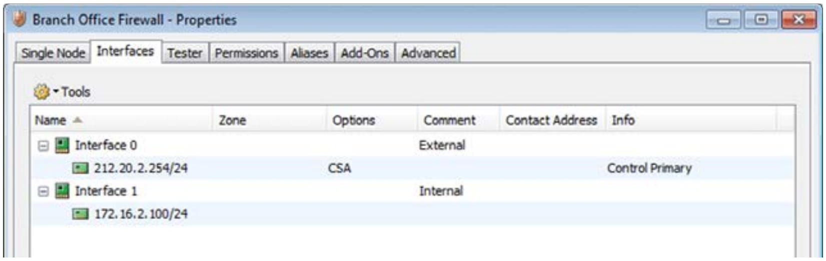

Name	Zone	Options	Comment	Contact Address	Info
Interface 0			External		
212.20.2.254/24		CSA			Control Primary
Interface 1			Internal		
172.16.2.100/24					

4. In the tree view, expand the tree and double-click the Cluster Virtual IP Address (CVI), Node Dedicated IP Address (NDI), or the IP address for which you want to define a contact address. The IP Address Properties dialog opens.
   - On Firewall Clusters, the CVI contact address is used for VPNs and NDI contact addresses are used for other system communications.

Proceed to one of the following depending on interface or engine type:

- Defining Contact Addresses for a Single Engine or a Cluster Virtual IP Address (page 68)
- Defining Contact Addresses for Node Dedicated IP Addresses (page 69)
- Defining Contact Addresses for an IPS Cluster or a Layer 2 Firewall Cluster (page 70)

Defining Contact Addresses for a Single Engine or a Cluster Virtual IP Address

▼ To define Contact Addresses for a single engine or a Cluster Virtual IP Address

1. In the IP Address Properties dialog, define the Default contact address. The Default contact address is used by default whenever a component that belongs to another Location connects to this interface.
   - If the interface has a static Default contact address, enter the **Default** contact address in the **Default** field. If the interface has a dynamic IP address, disable **Dynamic** (next to the **IPv4 Address** field) before entering the Default contact address.
   - If the interface has a dynamic Default contact address, select **Dynamic** (next to the **Default** field).

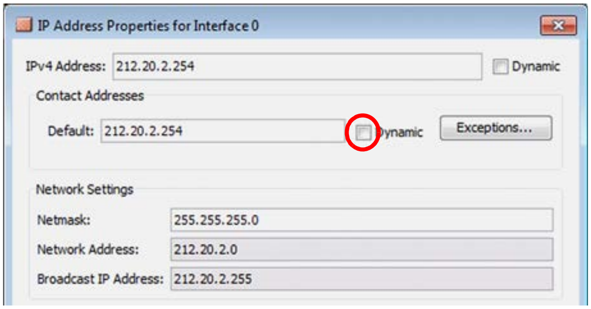

IPv4 Address: 212.20.2.254 □ Dynamic

Default: 212.20.2.254 □ Dynamic

Netmask: 255.255.255.0

Network Address: 212.20.2.0

Broadcast IP Address: 212.20.2.255

2. If components from some Locations cannot use the Default contact address to connect to the interface, click **Exceptions** to define Location-specific contact addresses. The Exceptions dialog opens.

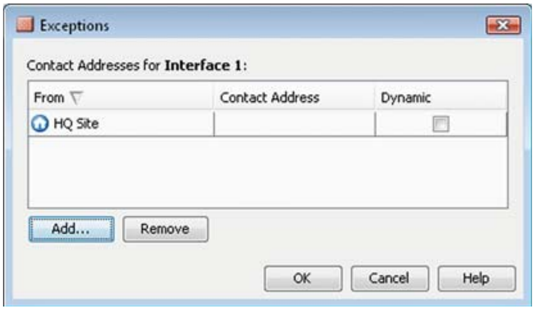

Contact Addresses for **Interface 1**:

From	Contact Address	Dynamic
HQ Site		□

3. Click **Add** and select the Location. A new row is added to the table.
4. Click the **Contact Address** column and enter the IP address that the components belonging to this Location must use when they contact the interface or select **Dynamic** if the interface has a dynamic contact address.

Note – Elements that belong to the same Location element always use the primary IP address (defined in the Properties dialog of the element) when contacting each other. Elements that do not belong to a specific Location are considered to belong to the Default Location.

5. Click **OK** to close the Exceptions dialog.

What's Next?

- If you want to define contact addresses for Firewall Cluster nodes, continue by Defining Contact Addresses for Node Dedicated IP Addresses.
- Otherwise, click **OK** to close the IP Address Properties dialog.

Defining Contact Addresses for Node Dedicated IP Addresses

▼ To define Contact Addresses for Node Dedicated IP Addresses

1. In the IP Address Properties dialog, define the IP contact address for each node in the Node Dedicated IP Address section by double-clicking the Contact Address cell for each node. The Exceptions dialog opens.

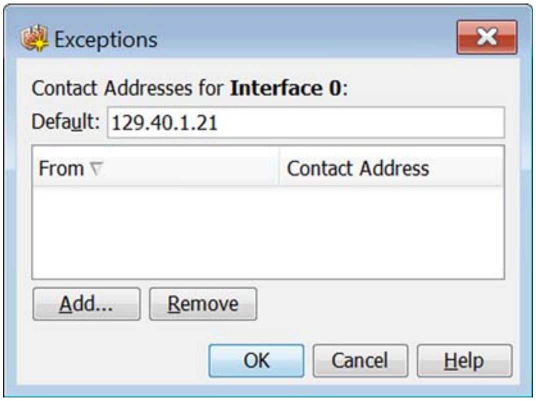

Default: 129.40.1.21

2. Enter the **Default** contact address at the top of the dialog. The Default contact address is used by default whenever a component that belongs to another Location connects to this interface.
3. If components from some Locations cannot use the Default contact address, click **Add** to define Location-specific contact addresses. A new row is added to the table.
4. Click the **Contact Address** column and enter the IP address that the components assigned to this Location must use when they contact the node.

Note – Elements that belong to the same Location element always use the primary IP address (defined in the Properties dialog of the element) when contacting each other. Elements that do not belong to a specific Location are considered to belong to the Default Location.

5. Click **OK** to close the Exceptions dialog.
6. One you have defined the contact address(es) for each node, click **OK** to close the IP Address Properties dialog.

Defining Contact Addresses for an IPS Cluster or a Layer 2 Firewall Cluster

▼ To add Contact Addresses for an IPS Cluster or a Layer 2 Firewall Cluster

1. In the IP Address Properties dialog, double-click the Contact Address cell. The Exceptions dialog opens.

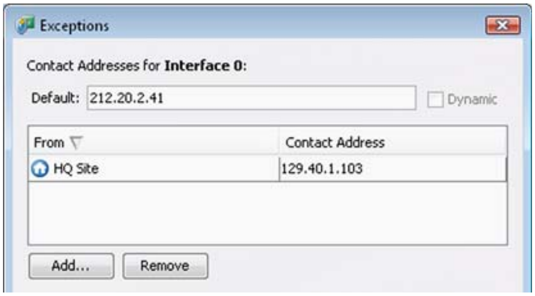

Contact Addresses for **Interface 0:**

Default: 212.20.2.41 □ Dynamic

From	Contact Address
HQ Site	129.40.1.103

2. Enter the **Default** contact address at the top of the dialog. The Default contact address is used by default whenever a component that belongs to another Location connects to this interface.
3. If components from some Locations cannot use the Default contact address, click **Add** to define Location-specific contact addresses. A new row is added to the table.
4. Click the **Contact Address** column and enter the IP address that the components belonging to this Location must use when they contact the interface.

Note – Elements that belong to the same Location element always use the primary IP address (defined in the Properties dialog of the element) when contacting each other. Elements that do not belong to a specific Location are considered to belong to the Default Location.

5. Click **OK** to close the Exceptions dialog.
6. Click **OK** to close the IP Address Properties dialog.

Defining Server Contact Addresses

In addition to situations where NAT is used between system components, server contact addresses are also needed for external servers, such as LDAP and Authentication Servers, if they provide services for system components.

You can configure multiple contact addresses for each Management Server and Log Server, and for an optional Authentication Server. Multiple contact addresses allow remote components to connect to a Management Server or a Log Server even if one of the Management Servers or Log Servers fails. If you use Multi-Link, it is recommended that the Management Server and the Log Server have a separate contact address for each NetLink so that if a NetLink goes down, the engines can still be managed (in case of reverse monitoring between the Management Server and the engines) and the engines can still send status and log data to the Log Server. Each Authentication Server node can have a single contact address for each Location.

▼ To define Server contact addresses

1. In the Security Engine Configuration view, browse to **Network Elements→Servers**.
2. Right-click the server element for which you want to define a contact address and select **Properties** from the menu. The Properties dialog for that Server opens.

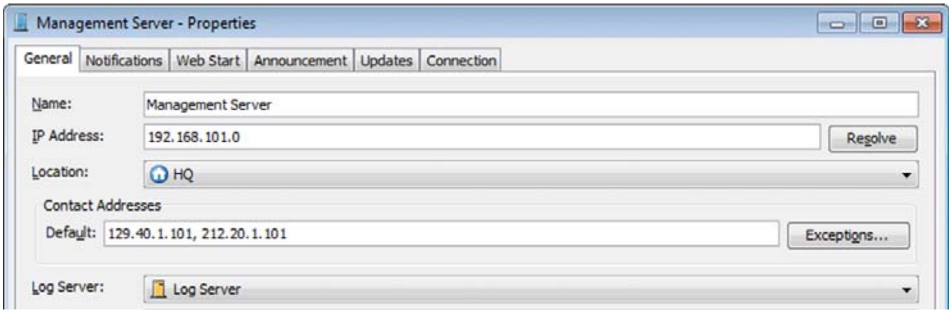

Name: Management Server

IP Address: 192.168.101.0

Location: HQ

Default: 129.40.1.101, 212.20.1.101

Log Server: Log Server

3. Select the **Location** of this server.
4. Enter the **Default** contact address. If the server has multiple Default contact addresses, separate the addresses with a comma.
5. (*Optional*) Click **Exceptions** to define further contact addresses for contacting the server from specific Locations. The Exceptions dialog opens.

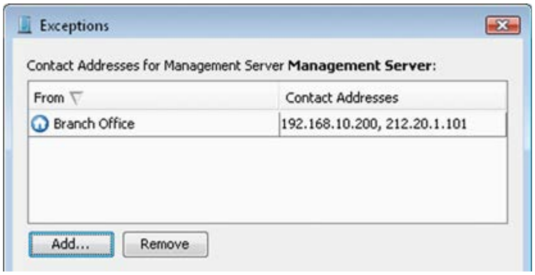

Contact Addresses for Management Server **Management Server**:

From	Contact Addresses
Branch Office	192.168.10.200, 212.20.1.101

6. Click **Add** and select a Location. A new row is added to the table.
7. Click the **Contact Address** column and enter the IP address(es) that the components belonging to this Location must use when they contact the Server.
   - You can enter several contact addresses per Location for Management Servers and Log Servers. Separate the contact addresses with a comma.

Note – Elements that belong to the same Location element always use the primary IP address (defined in the Properties dialog of the element) when contacting each other. Elements that do not belong to a specific Location are considered to belong to the Default Location.

8. Click **OK** to close the Exceptions dialog.
9. Click **OK** to close the Server properties dialog.

Defining a Contact Address for External Security Gateway End-Point

You must define a contact address for the end-point of an External Security Gateway if the IP address for contacting the Gateway is different from the IP address the Gateway actually has on its interface (for example, because of NAT).

▼ To define end-point contact address

1. In the External Gateway properties, switch to the **End-Points** tab.
2. Right-click an End-Point and select **Properties**. The End-Point Properties dialog opens.

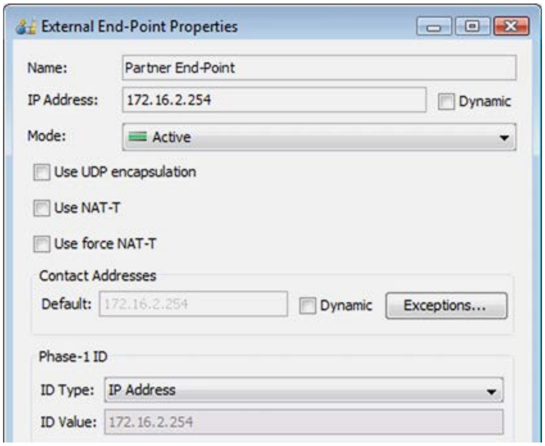

Name: Partner End-Point

IP Address: 172.16.2.254 □ Dynamic

Mode: Active

□ Use UDP encapsulation

□ Use NAT-T

□ Use force NAT-T

Contact Addresses

Default: 172.16.2.254 □ Dynamic

Phase-1 ID

ID Type: IP Address

ID Value: 172.16.2.254

3. Enter the **Default** contact address or select **Dynamic** if the Default contact address is dynamic.
   - The Default contact address is used by default whenever a component that belongs to another Location connects to this end-point.
4. (*Optional*) If components belonging to some Location cannot use the Default contact address, click **Exceptions** to define contact addresses that components belonging to specific Locations use to connect to this end-point. The Exceptions dialog opens.

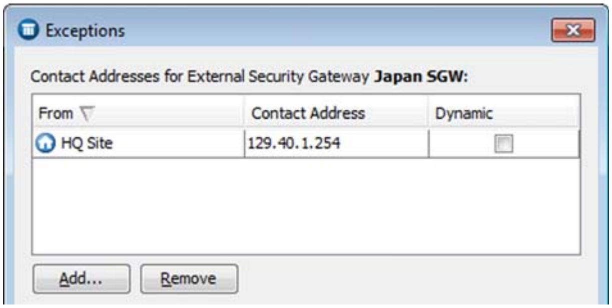

Contact Addresses for External Security Gateway **Japan SGW:**

From	Contact Address	Dynamic
HQ Site	129.40.1.254	□

5. Click **Add** and select a Location. A new row is added to the table.
6. Click the **Contact Address** column and enter the IP address that the components belonging to this Location must use when they contact the end-point or select **Dynamic** if the end-point has a dynamic contact address.
7. Click **OK** to close the Exceptions dialog.
8. Click **OK** to close the End-Point Properties dialog.

Selecting the Management Client Location

Prerequisites: Defining Locations, Defining Server Contact Addresses

If NAT is performed between the Management Client and a Log Server, you may need to change the Location of the Management Client to be able to view the logs. Which Location you need to select depends on the system configuration. The “Default” selection is appropriate if the Log Server is assigned a specific Location and the Log Server’s Default contact address is correct for your current network connection.

▼ To select the Management Client Location

➥ Click the location name in the status bar at the bottom right corner of the Management Client window and select the Location.

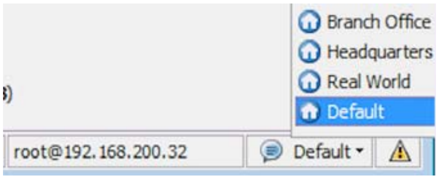

If there are no suitable configurations available, you may need to add a new Location and define a contact address for this specific Location in the Log Server’s properties.

Related Tasks

- ▶ Defining Locations (page 66)
- ▶ Defining Contact IP Addresses (page 66)

Configuring Multi-Link System Communications

Prerequisites: Defining a Multi-Link Route

If a remotely managed Firewall has Multi-Link, we recommend that you add a primary and a secondary management interface for different ISP connections to ensure connectivity if one of the Firewall’s local ISP connections fails. Make sure you configure these addresses consistently for the actual interface address on the Firewall, for the external contact addresses (if applicable), and in the NAT rules of the Firewall that protects the SMC servers (as necessary).

If a Management Server or Log Server is located behind Multi-Link in relation to components that contact them, define a contact address for each network connection and make sure your NAT rules translate from each external address to the correct internal address of the SMC server.

Related Tasks

- ▶ Setting Interface Options for Firewalls (page 453)

Chapter 6

Managing Elements

There are certain tasks that are common to most elements. Some of these tasks are not mandatory for defining an element, but are still helpful as you get your system up and running.

The following sections are included:

Using Categories

Prerequisites: None

Categories allow you to filter only a subset of elements for viewing in the Management Center. When you activate a Category filter, elements that do not belong to one of the selected Categories are filtered out of your view.

Categories help you manage large networks by filtering the elements displayed to you. You can, for example, create separate Categories for elements that belong to a Firewall, IPS, or Layer 2 Firewall configuration and select the correct category when you want to configure just one of the products. You can freely select how to assign the Categories and quickly and flexibly change which combinations of Categories are shown according to your tasks.

Configuration Overview

The general workflow for using Categories is as follows:

1. Create a new Category. See Creating New Categories.
2. Associate elements with the Category. See Selecting Categories for Elements (page 77).
3. Select the defined Category as the active Category Filter. See Activating Categories (page 77).
4. Change the Category selection as needed in the toolbar list. See Activating Categories (page 77).

Creating New Categories

You can create as many Categories as you need. You can base the categorization on any criteria. One element can belong to several Categories.

By default, the Category filter also includes a Not Categorized item which you can select on its own or in combination with a Category defined in your system.

▼ To define a new Category

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Expand the **Other Elements** branch of the Administration tree.
3. Right-click **Categories** and select **New Category**. The Category Properties dialog opens.
4. Give the new Category a unique name.
5. (*Optional*) Enter a **Comment** for your own reference.
6. Click **OK**. The Category is created and appears in the **Other Elements→Categories** branch of the tree. Categories are stored as elements in the system and they are displayed to other administrators as well.

What's Next?

- ▶ Selecting Categories for Elements

Related Tasks

- ▶ Activating Categories (page 77)

Selecting Categories for Elements

You can select one or more Categories for all types of elements. If you have activated a Category filter, all elements you create are by default added to the currently selected Categories, but you can override this selection.

If you want to view only existing elements that are not categorized yet, select the Not Categorized item in the Category Filter Toolbar. See Activating Categories (page 77) for more information.

▼ To select Categories for an element

1. Click **Select** next to the **Category** field in the properties of a new or existing element. The Category Selection dialog for the element opens.

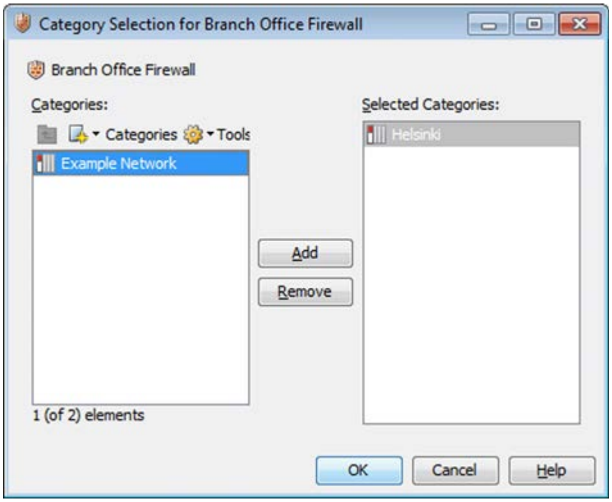

2. Add or remove Categories as needed.
3. Click **OK**.

Tip – Alternatively, right-click an element and select **Add Category** or **Remove Category**→**Remove** to open a dialog for adding or removing Categories.

Activating Categories

In most views, you can select Categories to restrict which elements are displayed. You can also display the elements that do not belong to any Category by selecting the Not Categorized filter. You can also select more than one Category at a time. See Filtering With Several Categories (page 78).

Note – The selected Category is applied *in all views* until you select a different Category or deselect the Use Category Filter button.

▼ To activate a Category

1. If the Category selection is not visible in the toolbar, select **View→Layout→Category Filter Toolbar**. The Category selection appears in the toolbar.

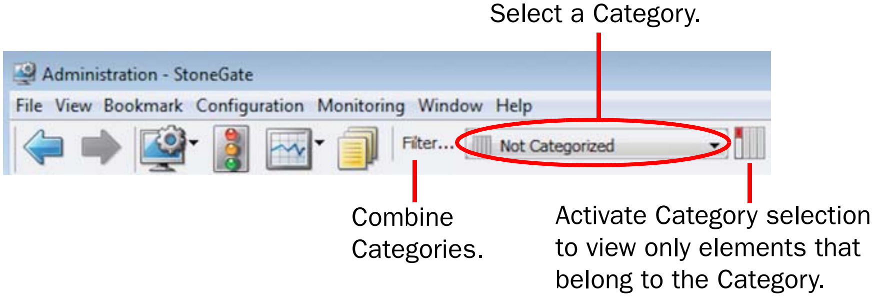
Select a Category.
Administration - StoneGate
File View Bookmark Configuration Monitoring Window Help
Filter...
Not Categorized
Combine Categories.
Activate Category selection to view only elements that belong to the Category.

2. Select one of the existing Categories.
3. Click the **Use Category Filter** button to the right of the Category selection. Only the elements that belong to the selected Category are displayed.

To display all the elements again, deselect the **Use Category Filter** button.

Filtering With Several Categories

You can select more than one Category at a time to display elements that are in any of the selected Categories.

▼ To create combined filters

1. Click the **Filter** button in the Category Filter toolbar. The Category Filter dialog opens.

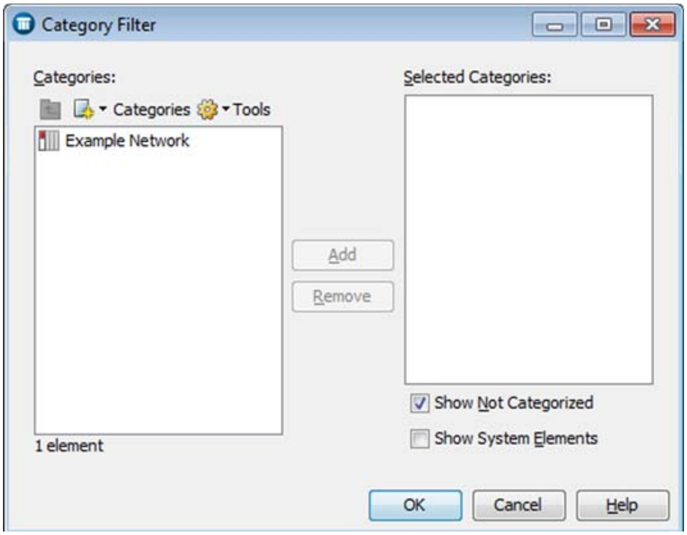

2. Select the Categories you want to add (Shift-click or Ctrl-click to select multiple Categories) and click **Add.** The selected Categories are added to the list of Selected Categories.
   - (*Optional*) If you want to view elements that do not have a Category (they belong to the Not Categorized Category), select **Show Not Categorized**.
   - (*Optional*) If you want to view elements that are predefined elements in the system (they belong to the System Category) select **Show System Elements**.

3. Click **OK**. The combined list of Categories is displayed as the current Category Filter in the Category selection.
4. Click the **Use Category Filter** button in the toolbar. Only the elements assigned to the selected categories are displayed.

Related Tasks

- Creating New Categories (page 76)
- Activating Categories (page 77)

Exporting, Importing, and Restoring Elements

Prerequisites: None

You can import and export most kinds of elements. This allows you to use the same elements in a different system without having to create the elements again.

You can also import old versions of elements or deleted elements by restoring them from a Policy Snapshot or from an Element Snapshot. If you have configured additional Management Servers, you can import, export, and restore elements that have been marked as obsolete.

What's Next?

- Exporting Elements (page 79)
- Importing Elements (page 81)
- Restoring Elements From Policy Snapshots (page 83)
- Restoring Elements From Element Snapshots (page 86)
- Restoring Obsolete Elements (page 89)

Related Tasks

- Importing Users from an LDIF File (page 852)
- Exporting Users to an LDIF File (page 853)

Exporting Elements

This section explains how to export elements using the Management Client. For command line exporting, see the command `SgExport` in Command Line Tools (page 1145). For PDF or HTML output, see Saving Elements, Log Data, Reports, and Statistics (page 51).

Note – The exported file is meant for importing elements into the database of a Management Server. It is not meant to be viewed or edited in external applications.

What's Next?

- Exporting Selected Elements (page 80)
- Exporting All Elements (page 80)

Exporting Selected Elements

▼ To export selected elements

1. Select **File→Export→Export Elements**. The Export dialog opens.

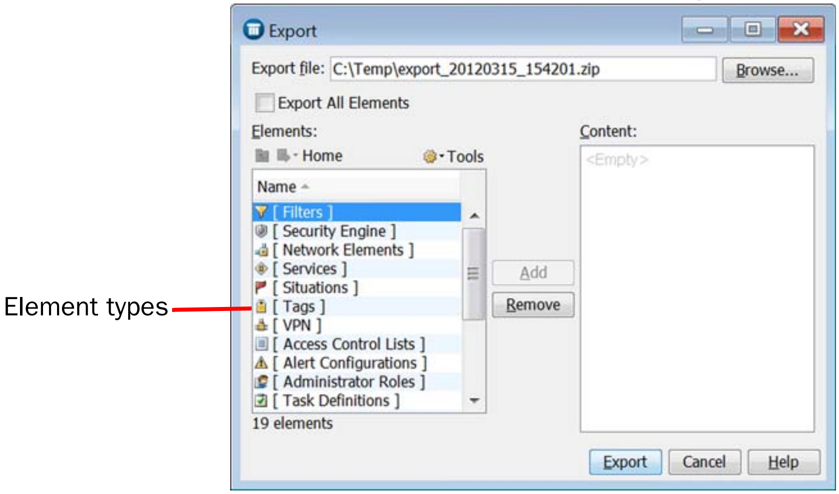

Element types

Export file: C:\Temp\export_20120315_154201.zip

□ Export All Elements

2. Enter a file name for the export file, or click **Browse** to select the location where you want to create the file.

3. (*Optional*) Select **Tools→Show Obsolete** if you want to view and export elements that have been marked as obsolete. See Marking Elements as Obsolete (page 88).

4. Select the element(s) that you want to export and click **Add**. The elements are added to the Content list.
   - You can export most but not all kinds of elements. The elements that you cannot export are elements that contain particularly sensitive information (for example, administrator accounts and certificates). To export an element that references an element that cannot be exported, you must manually create before the export a corresponding element that has the same name as the original referenced element. Otherwise, the export fails.

5. When you have finished adding elements to the Content list, click **Export**. A new tab opens to show the progress of the export.

Exporting All Elements

▼ To export all elements

1. Select **File→Export→Export All Elements**. The Export dialog opens with the **Export All Elements** option selected by default.

2. Enter a file name for the export file, or click **Browse** to select the location where you want to create the file.
   - (*Optional*) Select **Tools→Show Obsolete** if you want to view and export elements that have been marked as obsolete. See Marking Elements as Obsolete (page 88).

3. Click **Export**. A new tab opens to show the progress of the export.

Importing Elements

This section explains how to import elements using the Management Client. You can import elements from a CSV (comma-separated value) file, a TSV (tab-separated value) file, or a ZIP file of an element exported from the Management Client. See Exporting Elements (page 79). You can also import elements from the command line. See the command `sgImport` in Command Line Tools (page 1145).

Creating a CSV File or a TSV File

You can import elements from a CSV (comma-separate value) file or a TSV (tab-separated value) file. In a CSV file, all the values are separated by a comma. In a TSV file, all the values are separated by a tabulator.

▼ To create a CSV or TSV file

1. Create a new CSV or TSV file (for example, in a spreadsheet application).
2. Enter the header (the first row) into the file:
   - In a CSV file: enter **`ip,name,comment`**
   - In a TSV file: enter **`ip`**<tab>**`name`**<tab>**`comment`**
   - Only the **`ip`** parameter is mandatory in the header row. Note that all data entered in the file must follow the format of the header row.
3. Enter the IP address of the element and optionally a name and a comment on the row below the header. Use the same format as in the header.
   - Only the IP address field is mandatory. However, each row in the file must have the same format as the header row. If you have other parameters in addition to **`ip`** in the header row, you must enter a separator (a comma or a tabulator) even if you do not enter a name and/or a comment on a row.

Example **Enter `10.1.1.1,,` if the header row is `ip,name,comment` but you want to omit the name and the comment in a CSV file.**

Example **Enter `10.1.1.1`<tab><tab> if the header row is `ip`<tab>`name`<tab>`comment` but you want to omit the name and the comment in a TSV file.**

   - If you do not enter a name, a name is automatically generated for the element based on its IP address. The system automatically detects the element type based on the IP address syntax: 10.10.10.10 - a Host, 10.10.10.0/24 - a Network, and 10.10.10.10-20.20.20.20 - an Address Range.

4. (*Optional*) Repeat step 3 to add more rows.
5. Save the file.

What's Next?

▶ Importing Elements From a File (page 82)

Importing Elements From a File

▼ To import elements

1. Select **File→Import→Import Elements**. The Import File dialog opens.
2. Select the file(s) you want to import and click **Import**. The Management Server automatically checks if any elements in the file to be imported have the same name and XML content as any elements that already exist in the system. A new tab opens.

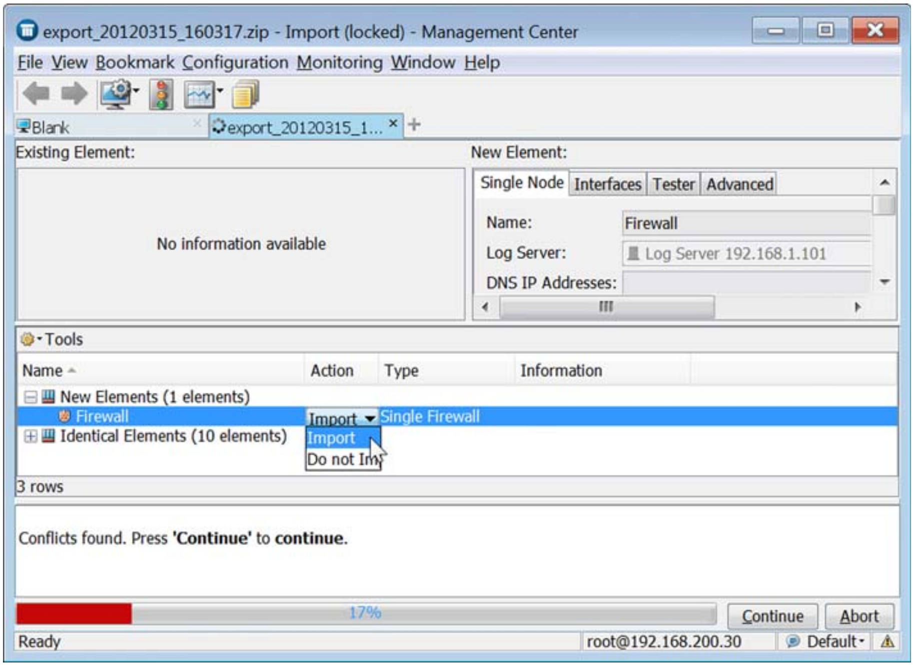

Existing Element:

No information available

New Element:

Name: Firewall

Log Server: Log Server 192.168.1.101

DNS IP Addresses:

Name	Action	Type	Information
New Elements (1 elements)			
Firewall	Import	Single Firewall	
Identical Elements (10 elements)			

3 rows

Conflicts found. Press **'Continue'** to **continue.**

3. If any conflicts are found between elements in the import file and elements that currently exist in the system, select the **Action** for each conflict according to the conflict type:

Table 6.1 Conflict Types and Actions

Type	Action	Explanation
New Elements	Import	The element that already exists in the system is overwritten with the element in the import file.
	Do not Import	The element is not imported.
Identical Elements	Do not Import	The element is not imported. Because *Identical Elements* do not produce any changes when imported, the Do not Import action is automatically selected for Identical Elements.
	Undelete Existing	An obsolete element in the system and the element to be imported are identical. The obsolete element is restored.
Conflicts	Import	The element that already exists in the system is overwritten with the element in the import file.
	Do not Import	The element is not imported.

Table 6.1 Conflict Types and Actions

Type	Action	Explanation
Conflicts (*cont.*)	Rename	The element in the import file is imported as a new element in the system and renamed.
	Undelete Existing and Import	An obsolete element in the system and the element to be imported are identical. The Undelete Existing and Import action restores the obsolete element and it is overwritten by the element to be imported.

- To view the elements in XML format, select **Tools→Show XML**. If there is a conflict between elements in the system and elements in the import file, the differences in the XML format are shown in color.

4. If there are no conflicts in the system or when you have selected the **Action** for conflicting elements, click **Continue** to start the import.
5. When the importing is done, click **Close**.

Restoring Elements From Policy Snapshots

This section explains how to restore elements from Policy Snapshots using the Management Client. You can either restore all elements of a snapshot, or select the elements to be restored.

What's Next?

▶ Restoring All Elements From a Policy Snapshot

▶ Restoring Selected Elements From a Policy Snapshot

Restoring All Elements From a Policy Snapshot

▼ To restore all elements from a Policy Snapshot

1. Select **Configuration→Configuration→Security Engine**. The Security Engine Configuration view opens.
2. Browse to **Other Elements→Policy Snapshots**. The list of Policy Snapshot types opens.
3. Expand the Policy Snapshot type to open the list of Policy Snapshots.
4. Right-click the Policy Snapshot from which you want to restore elements and select **Tools→Restore**. The restoring of the elements starts.
5. If there are *Conflicts* between the elements in the Policy Snapshot and the elements that currently exist in the system, resolve them by selecting the **Action**:

Table 6.2 Conflict Types and Actions

Type	Action	Explanation
New Elements	Import	The element that already exists in the system is overwritten with the element in the import file.
	Do not Import	The element is not imported.

Table 6.2 Conflict Types and Actions

Type	Action	Explanation
Identical Elements	Do not Import	The element is not imported. Because *Identical Elements* do not produce any changes when imported, the Do not Import action is automatically selected for Identical Elements.
	Undelete Existing	An obsolete element in the system and the element to be imported are identical. The obsolete element is restored.
Conflicts	Import	The element that already exists in the system is overwritten with the element in the import file.
	Do not Import	The element is not imported.
	Rename	The element in the import file is imported as a new element in the system and renamed.
	Undelete Existing and Import	An obsolete element in the system and the element to be imported are identical. The Undelete Existing and Import action restores the obsolete element and it is overwritten by the element to be imported.

- To view the elements in XML format, select **Tools→Show XML**. If there is a conflict between elements in the system and elements in the Policy Snapshot, the differences in the XML format are shown in color.

6. Click **Continue**.
7. When the importing is done, click **Close**.

Restoring Selected Elements From a Policy Snapshot

▼ To restore selected elements from a Policy Snapshot

1. Select **Configuration→Configuration→Security Engine**. The Security Engine Configuration view opens.
2. Browse to **Other Elements→Policy Snapshots**. The list of Policy Snapshots Types opens.
3. Expand the branch for the Policy Snapshot type. The list of Policy Snapshots opens.
4. Right-click the Policy Snapshot from which you want to restore elements, and select **View Policy Snapshot**. A preview of the Policy Snapshot opens.
5. Select one or several elements to restore, right-click and select **Tools→Restore**. The restoring of the element(s) starts.

6. If there are *Conflicts* between the elements in the Policy Snapshot and the elements that currently exist in the system, resolve them by selecting the **Action**:

Table 6.3 Conflict Types and Actions

Type	Action	Explanation
New Elements	Import	The element that already exists in the system is overwritten with the element in the import file.
	Do not Import	The element is not imported.
Identical Elements	Do not Import	The element is not imported. Because *Identical Elements* do not produce any changes when imported, the Do not Import action is automatically selected for Identical Elements.
	Undelete Existing	An obsolete element in the system and the element to be imported are identical. The obsolete element is restored.
Conflicts	Import	The element that already exists in the system is overwritten with the element in the import file.
	Do not Import	The element is not imported.
	Rename	The element in the import file is imported as a new element in the system and renamed.
	Undelete Existing and Import	An obsolete element in the system and the element to be imported are identical. The Undelete Existing and Import action restores the obsolete element and it is overwritten by the element to be imported.

- To view the elements in XML format, select **Tools→Show XML**. If there is a conflict between elements in the system and elements in the Policy Snapshot, the differences in the XML format are shown in color.

7. Click **Continue**.

8. When the importing is done, click **Close**.

Restoring Elements From Element Snapshots

This section explains how to restore elements from Element Snapshots stored in Audit logs using the Management Client.

▼ To restore an element from an Element Snapshot

1. Display Element Snapshots in the Logs view. See To view and compare Element Snapshots (page 118).
2. Right-click the audit entry of an element and select **Compare to Current Element**. The Compare Elements dialog opens.
   - If the Element Snapshot properties differ from the properties of the element that currently exists in the system, a red border is displayed around the Audit Log Version (snapshot) and the Current Version of the element.

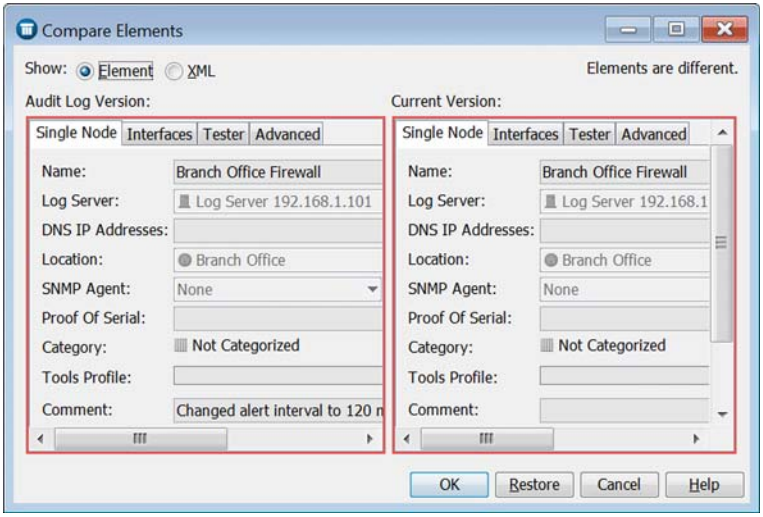

Show: ● Element ○ XML

Elements are different.

	Audit Log Version:	Current Version:
	Single Node / Interfaces / Tester / Advanced	Single Node / Interfaces / Tester / Advanced
Name:	Branch Office Firewall	Branch Office Firewall
Log Server:	Log Server 192.168.1.101	Log Server 192.168.1
DNS IP Addresses:		
Location:	Branch Office	Branch Office
SNMP Agent:	None	None
Proof Of Serial:		
Category:	Not Categorized	Not Categorized
Tools Profile:		
Comment:	Changed alert interval to 120 n	

   - Select **Show: XML** to view in XML format the differences between the snapshot and the element that currently exists in the system. Differences are indicated in red.

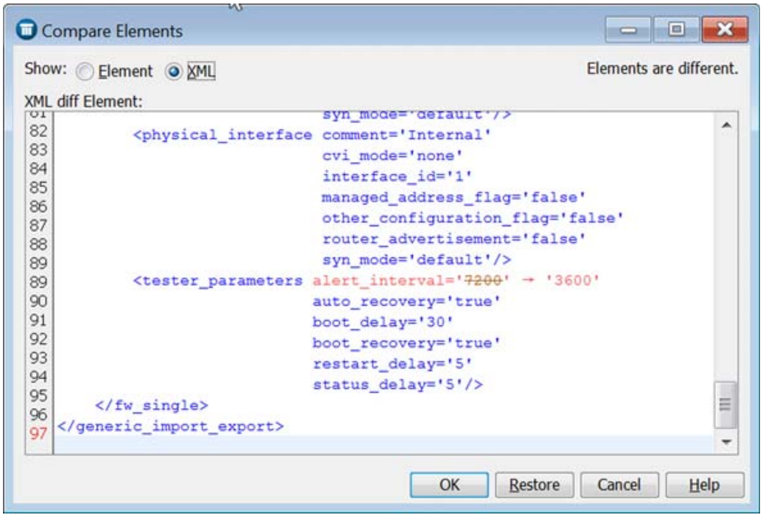

Show: ○ Element ● XML

Elements are different.

XML diff Element:

```
                              syn_mode='default'/>
        <physical_interface comment='Internal'
                            cvi_mode='none'
                            interface_id='1'
                            managed_address_flag='false'
                            other_configuration_flag='false'
                            router_advertisement='false'
                            syn_mode='default'/>
        <tester_parameters alert_interval='~~7200~~' → '3600'
                           auto_recovery='true'
                           boot_delay='30'
                           boot_recovery='true'
                           restart_delay='5'
                           status_delay='5'/>
    </fw_single>
</generic_import_export>
```

3. Click **Restore** to restore the properties of the Element Snapshot to the current element. The restoring of the element properties starts.

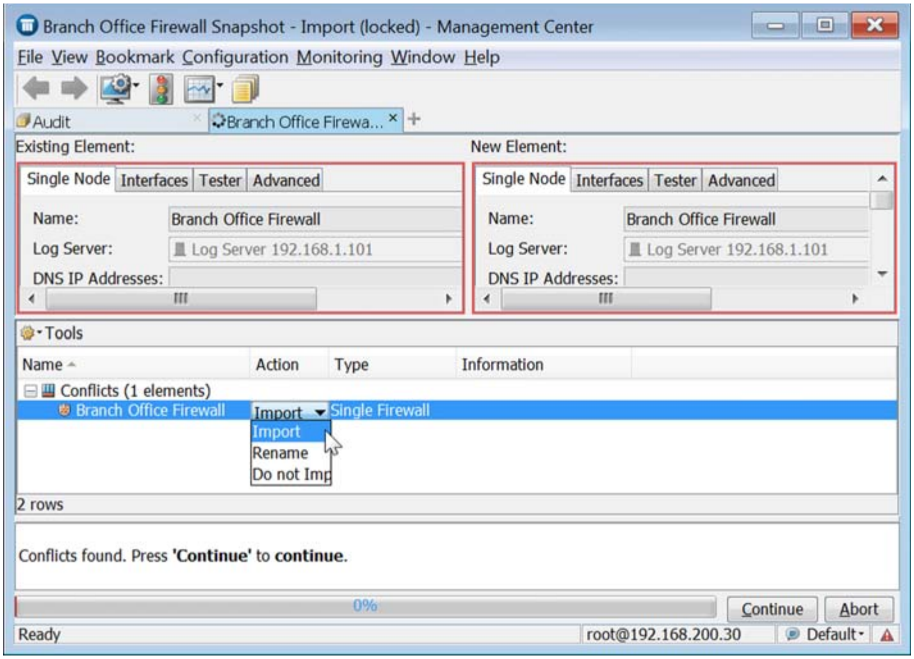

Existing Element:

Name: Branch Office Firewall
Log Server: Log Server 192.168.1.101
DNS IP Addresses:

New Element:

Name: Branch Office Firewall
Log Server: Log Server 192.168.1.101
DNS IP Addresses:

Name	Action	Type	Information
Conflicts (1 elements)			
Branch Office Firewall	Import	Single Firewall	

2 rows

Conflicts found. Press **'Continue'** to **continue.**

0%

4. If there are *Conflicts* between the Element Snapshot and the element that currently exists in the system, resolve them by selecting the **Action**:
   - **Import**: The element that already exists in the system is overwritten with the element in the Element Snapshot.
   - **Rename**: The element in the Element snapshot is renamed and imported as a new element.
   - **Do not Import**: The element is not imported.
5. Click **Continue**. Importing the elements of the Element Snapshot starts.

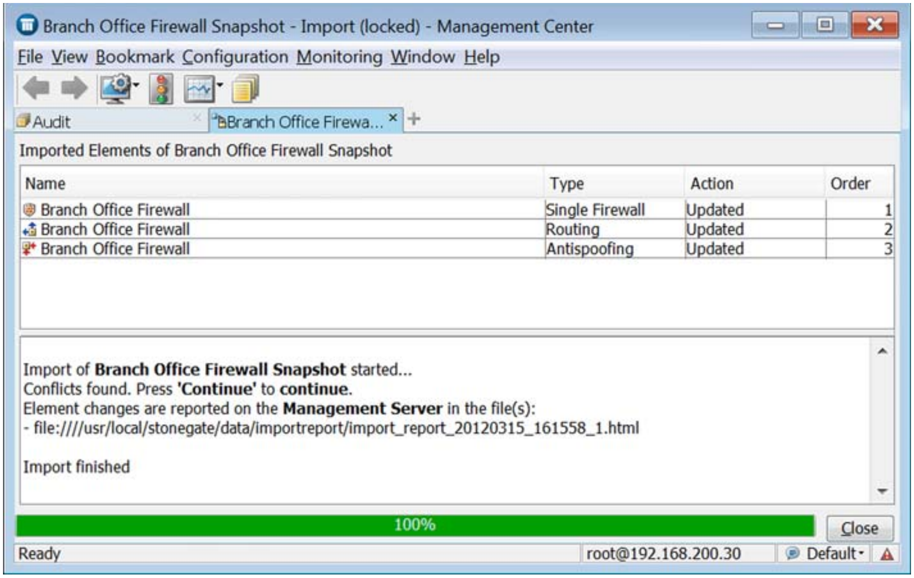
Branch Office Firewall Snapshot - Import (locked) - Management Center
File View Bookmark Configuration Monitoring Window Help
Audit
Branch Office Firewa...
Close
Ready
root@192.168.200.30
Default

Imported Elements of Branch Office Firewall Snapshot

Name	Type	Action	Order
Branch Office Firewall	Single Firewall	Updated	1
Branch Office Firewall	Routing	Updated	2
Branch Office Firewall	Antispoofing	Updated	3

```
Import of Branch Office Firewall Snapshot started...
Conflicts found. Press 'Continue' to continue.
Element changes are reported on the Management Server in the file(s):
- file:///usr/local/stonegate/data/importreport/import_report_20120315_161558_1.html

Import finished
```

100%

6. When the importing is done, click **Close**.

Related Tasks

- ▶ Viewing and Comparing Element Snapshots (page 118)
- ▶ Restoring Elements From Policy Snapshots (page 83)

Locking and Unlocking Elements

Prerequisites: None

An administrator who is allowed to edit an element can lock the element and add a comment to explain the reason for locking it. You must be logged in to the Domain in which the element is stored to be able to lock or unlock the element. Locked elements are displayed with a lock symbol. You must unlock a locked element before modifying or deleting it. Any administrator who is allowed to edit the locked element can remove the lock.

You cannot lock predefined system elements or elements that are marked as obsolete.

▼ To lock an element

1. Right-click the element you want to lock and select **Tools→Lock**. The Lock Properties dialog opens.
2. Enter a **Comment** explaining the reason for locking the element.
3. Click **OK**. The element is now locked and a lock symbol is displayed on its icon.

▼ To unlock an element

1. Right-click the element you want to unlock and select **Tools→Unlock**.
2. Click **Yes** in the dialog that opens to confirm the deletion of the lock.

Marking Elements as Obsolete

Prerequisites: None

You cannot delete elements directly. You must first mark the elements as obsolete. You can view the elements that have been marked as obsolete by selecting **View→Obsolete Elements**. To search for elements that have been marked as obsolete, select **View→Search→Search Obsolete Elements**. An administrator with unrestricted permissions can search for obsolete elements in all administrative domains.

An obsolete element is still valid in any previous configuration where the element was used before you marked it as obsolete. However, you cannot add an obsolete element to any new configuration.

When you export and import elements that have been marked as obsolete, all references to the elements remain valid. An element that has been exported from Obsolete Elements remains obsolete when imported to an environment with several Management Servers.

You can also restore obsolete elements. An obsolete element is permanently deleted only when you delete the elements from Obsolete Elements. See Deleting Obsolete Elements (page 89).

▼ To mark an element as obsolete

1. Right-click the element and select **Mark as Obsolete**. A confirmation dialog opens.
   - If the element you are marking as obsolete is currently used in any configuration, the confirmation dialog also shows a list of all references to the element. Click **Open References** to view the references. To remove the references, right-click each element and select **Edit**.
2. Click **Yes**.

Related Tasks

- ▶ Restoring Obsolete Elements
- ▶ Deleting Obsolete Elements
- ▶ Searching for Element References (page 46)
- ▶ Disabling Administrator Accounts (page 264)
- ▶ Deleting Domains (page 292)

Restoring Obsolete Elements

▼ To restore an Obsolete Element

1. Select **View→Obsolete Elements**. The elements that have been marked as obsolete are displayed.
2. Right-click the obsolete element that you want to restore and select **Undelete**. A confirmation dialog opens.
   - To view the references to the element you are restoring, click **Open References** in the confirmation dialog.
3. Click **Yes**. The element is restored, and you can, for example, add it to new configurations.

Deleting Obsolete Elements

You can permanently delete an element that you have marked as obsolete. If you have selected **Tools→Show Obsolete** to view the elements that have been marked as obsolete, you can delete the obsolete element in your current view. Otherwise, you can either delete a single element from the Obsolete Elements branch or delete all obsolete elements at the same time.

An element can only be deleted by administrators who have sufficient privileges, and only if the element is not used in any configuration, for example, in a policy.

A history of the changes that an administrator makes to elements is saved in the Management Center. If you delete the administrator account, all the history information on the changes the administrator has made is lost. If you do not want to lose the history information, disable the administrator account instead of deleting the account: right-click the administrator account and select **Disable Administrator**.

Caution – Deletion is permanent. There is no undo. To recover a deleted element, you must either recreate it, or restore it from a previously created backup or XML export file that contains the element.

▼ To delete an element from Obsolete Elements

1. Select **View→Obsolete Elements**. The elements that have been marked as obsolete are displayed.
2. Right-click the element that you are deleting and select **Delete**. A confirmation dialog opens.
   - If the element you are deleting is currently used in any configuration, click **Open References** in the confirmation dialog to view the reference(s). Right-click each element and select **Edit** to remove the reference(s).
3. Click **Yes**. The element is permanently deleted.

▼ To delete all Obsolete Elements

1. Select **View→Obsolete Elements**. The elements that have been marked as obsolete are displayed.
2. Right-click the **Obsolete Elements** branch and select **Delete Obsolete Elements**.
3. Click **Yes**. All elements in **Obsolete Elements** are permanently deleted.

MONITORING

In this section:

CHAPTER 7

MONITORING THE SYSTEM

All system components can be monitored through the Management Client.

The following sections are included:

Getting Started with System Monitoring

Prerequisites: None

There are several ways to monitor the system in the Management Client. You can:

- Monitor the status of individual components and view a summary of the system status. See Monitoring the System Status.
- Create customizable overviews of the system. See Creating Overviews (page 104).
- Monitor enforced blacklists, open connections, active VPN SAs, active users, and routing. See Monitoring Connections, Blacklists, VPN SAs, Users, and Routing (page 111).
- View connection information on a map. See Monitoring Connections on a Map (page 120).
- Check which configurations and policies are currently applied in the system. See Monitoring Configurations and Policies (page 124).
- Check which actions administrators take. See Monitoring Administrator Actions (page 125).
- Check the status of Tasks that schedule commands to run automatically. See Monitoring Task Execution (page 125).
- Monitor the status of the maintenance contract. Checking Maintenance Contract Information (page 129).
- Monitor the status of internal certificates and internal certificate authorities. See Checking When Internal Certificates or Internal CAs Expire (page 130).

Related Tasks

- ▶ Getting Started with the Logs View (page 148)
- ▶ Getting Started with Reports (page 168)
- ▶ Getting Started with Third-Party Device Monitoring (page 132)

Monitoring the System Status

Prerequisites: None

The Management Server automatically keeps track of the operation of the system components. You can check the status of the system and individual components in the System Status view.

You can also monitor the status of third party devices in the Management Client. See Getting Started with Third-Party Device Monitoring (page 132) for more information.

There are several ways to open the System Status view. For example:

- Select **Monitoring→System Status** from the menu or click the corresponding toolbar icon.
- Right-click an element that is monitored, and select **Monitoring→System Status**.

What's Next?

- ▶ For an overview of the System Status View, start in Default Arrangement of System Status View (page 95).

Default Arrangement of System Status View

The System Status view displays the monitored elements in the Status tree with icons that indicate their status. The System Summary provides more detailed status information. The statuses are constantly updated, so you can quickly check which components in your system may need your attention at any given time.

Illustration 7.1 System Status View - Default Arrangement

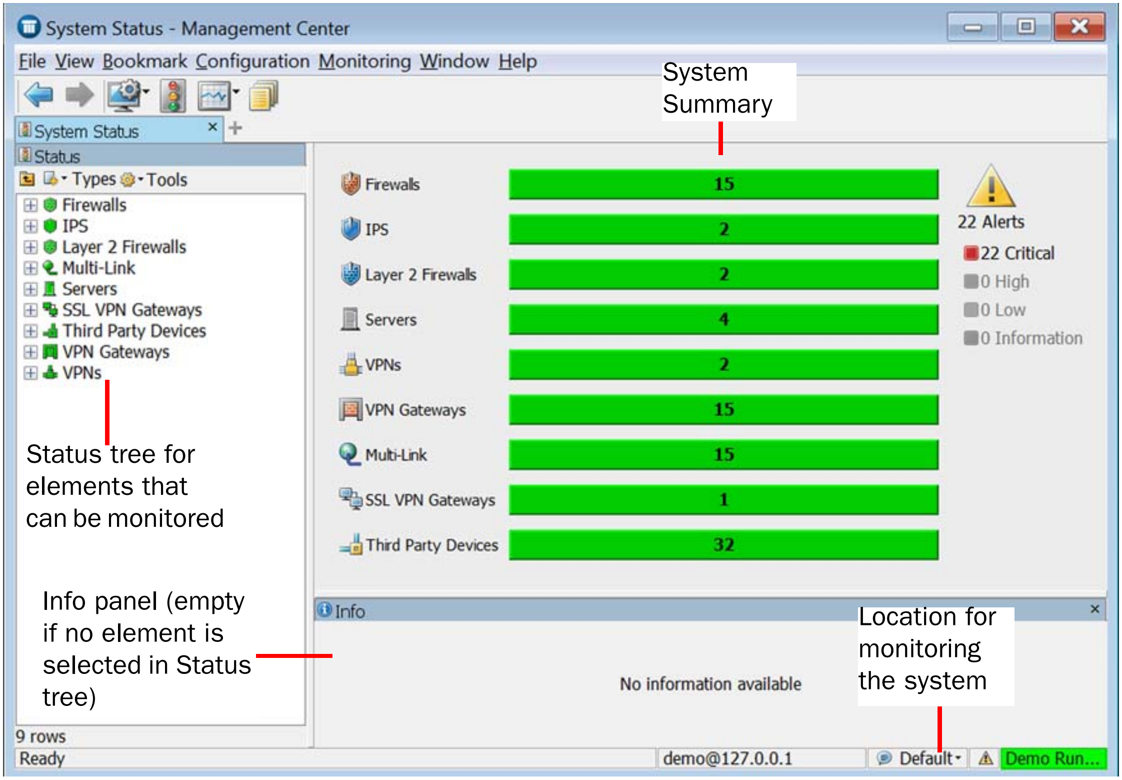

Firewalls	15
IPS	2
Layer 2 Firewalls	2
Servers	4
VPNs	2
VPN Gateways	15
Multi-Link	15
SSL VPN Gateways	1
Third Party Devices	32

22 Alerts

- 🟥 22 Critical
- ■ 0 High
- ■ 0 Low
- ■ 0 Information

No information available

System Summary

Status tree for elements that can be monitored

Info panel (empty if no element is selected in Status tree)

Location for monitoring the system

To check the status of a component in the Status tree, expand the tree and place the mouse pointer over any element to see its IP address and status as text in a tooltip. For information on the different statuses, see Reading Component Statuses (page 100).

System Summary

The System Summary provides an overview of the status of monitored components and displays the number of active alerts in the system. You can view the System Summary in the System Status view or insert it as a section in your own custom monitoring views (Overviews, Reports). Double-click the status information for more details.

Illustration 7.2 System Summary in the System Status View

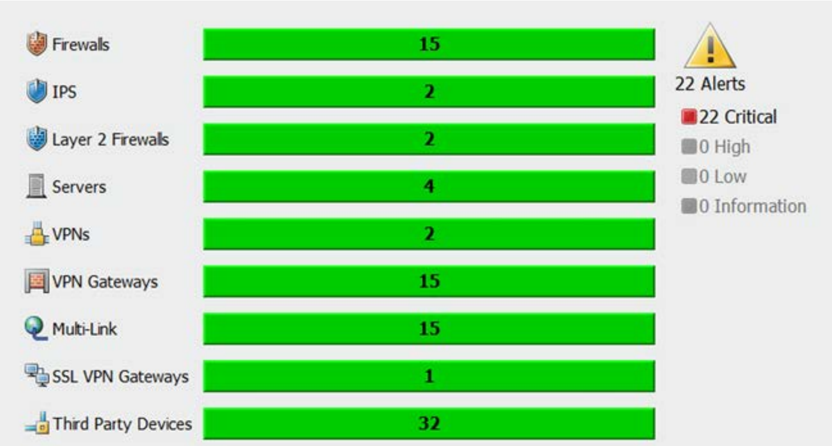

Firewalls	15
IPS	2
Layer 2 Firewalls	2
Servers	4
VPNs	2
VPN Gateways	15
Multi-Link	15
SSL VPN Gateways	1
Third Party Devices	32

22 Alerts

🟥 22 Critical
■ 0 High
■ 0 Low
■ 0 Information

Viewing System Status for a Selected Element

When you select an element in the Status tree, the System Status view automatically shows the element status as a connectivity diagram (select **View→Draw Diagram on Selection**). The Info panel displays detailed information on the selected element.

When you select **View→Filter Diagram Content**, you see the selected elements and their related connections in the automatically generated diagram. Otherwise, all system elements and connections are shown.

Illustration 7.3 System Status Information for Selected Element

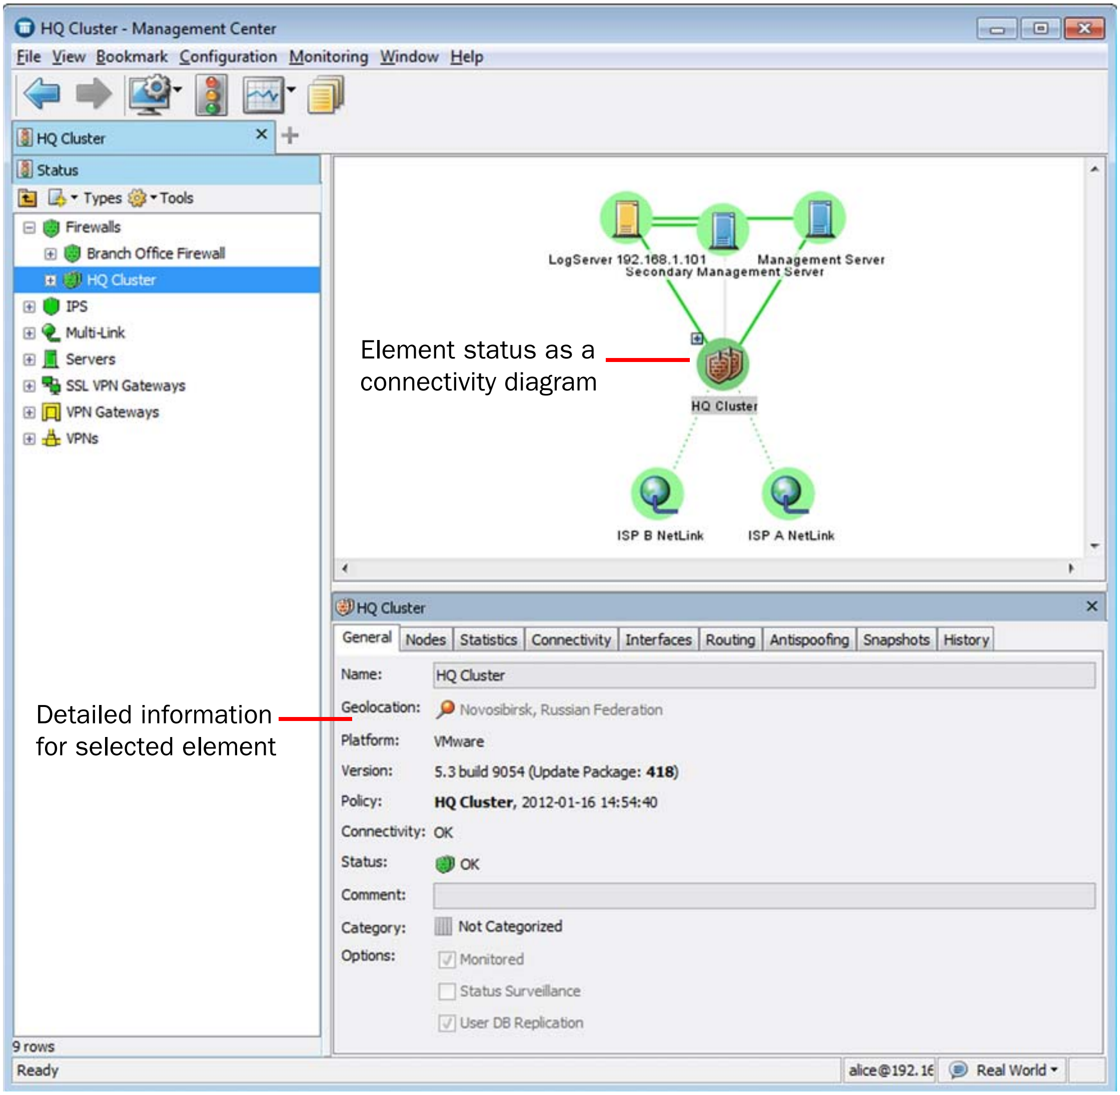

Name: HQ Cluster

Geolocation: Novosibirsk, Russian Federation

Platform: VMware

Version: 5.3 build 9054 (Update Package: **418**)

Policy: **HQ Cluster**, 2012-01-16 14:54:40

Connectivity: OK

Status: OK

Comment:

Category: Not Categorized

Options:
☑ Monitored
☐ Status Surveillance
☑ User DB Replication

Element status as a connectivity diagram

Detailed information for selected element

Viewing Appliance Configuration Status

When you configure a Stonesoft appliance using the plug-and-play configuration method, you can view the status of the configuration process by selecting **File→System Tools→View Appliance Configuration Status**.

Illustration 7.4 Appliance Configuration Status Dialog

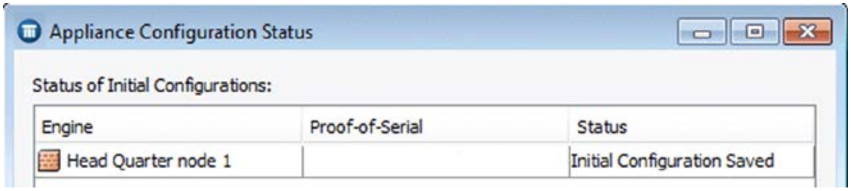

Status of Initial Configurations:

Engine	Proof-of-Serial	Status
Head Quarter node 1		Initial Configuration Saved

Info Panel

The Info panel shows detailed information on the component that you select. The level of the selected component in the System Status tree determines which tabs are shown. For example, the Interface Status tab shows information on the network ports of the selected engine (such as speed/duplex). The Appliance Status tab shows the hardware status of the selected device. Select **View→Info** to see the Info panel.

Illustration 7.5 Node Status in the Info panel

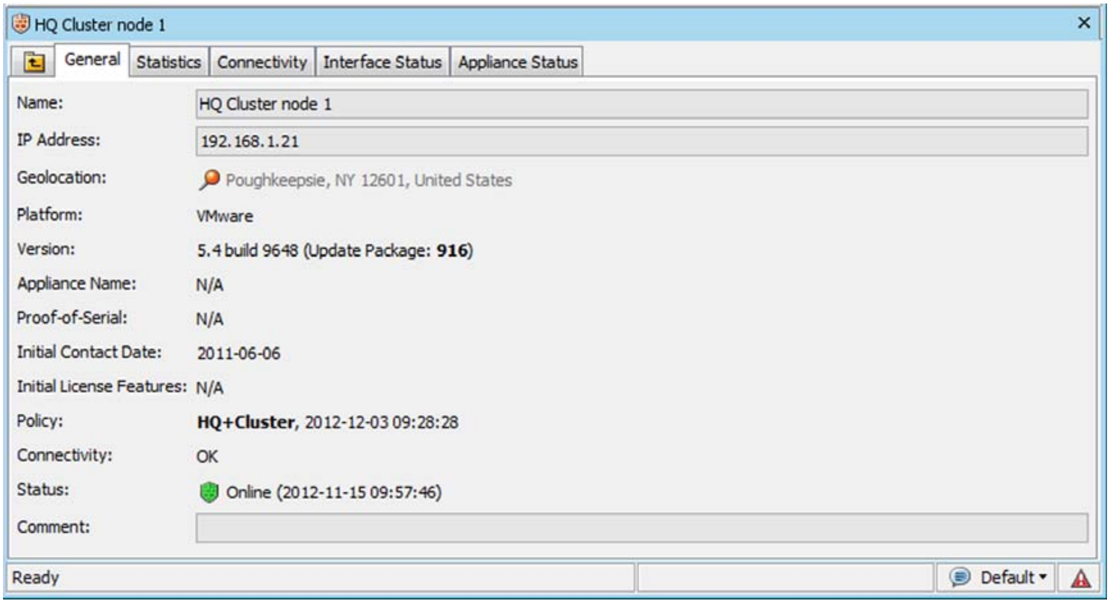

Name:	HQ Cluster node 1
IP Address:	192.168.1.21
Geolocation:	Poughkeepsie, NY 12601, United States
Platform:	VMware
Version:	5.4 build 9648 (Update Package: **916**)
Appliance Name:	N/A
Proof-of-Serial:	N/A
Initial Contact Date:	2011-06-06
Initial License Features:	N/A
Policy:	**HQ+Cluster**, 2012-12-03 09:28:28
Connectivity:	OK
Status:	Online (2012-11-15 09:57:46)
Comment:	

Commands for Monitoring Components

Actions for monitoring a component are available in the Monitoring branch of the element's right-click menu. The available actions depend on the component type.

Illustration 7.6 Monitoring Menu for Components - Firewall/VPN Engine

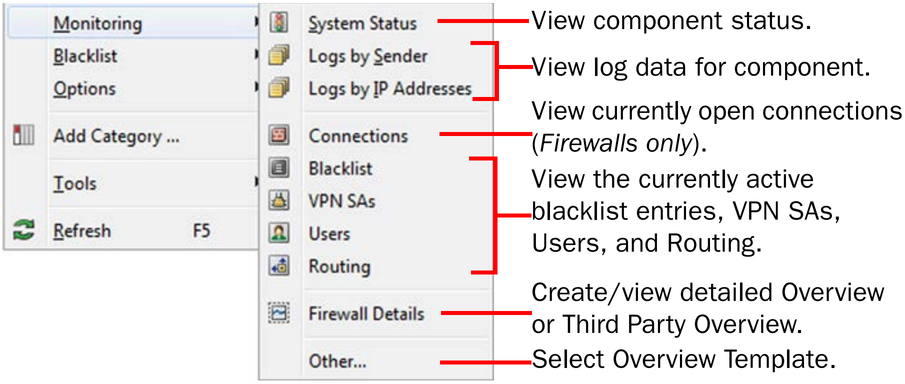

View component status.

View log data for component.

View currently open connections (*Firewalls only*).

View the currently active blacklist entries, VPN SAs, Users, and Routing.

Create/view detailed Overview or Third Party Overview.

Select Overview Template.

Monitoring Menu

All the main actions for monitoring the system are available in the Monitoring menu.

Illustration 7.7 Monitoring Menu

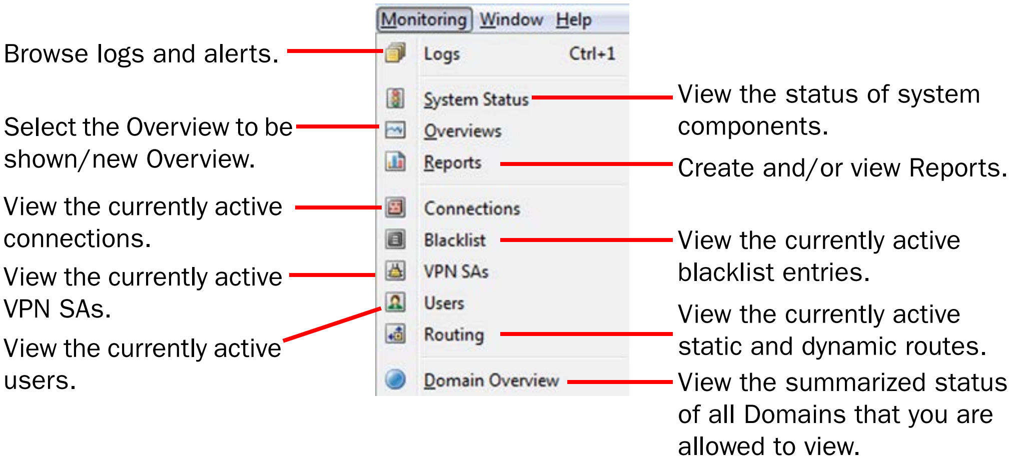

Browse logs and alerts.

Select the Overview to be shown/new Overview.

View the currently active connections.

View the currently active VPN SAs.

View the currently active users.

View the status of system components.

Create and/or view Reports.

View the currently active blacklist entries.

View the currently active static and dynamic routes.

View the summarized status of all Domains that you are allowed to view.

Related Tasks

- Reading Component Statuses
- Getting Started with the Logs View (page 148)
- Creating Overviews (page 104)
- Getting Started with Reports (page 168)
- Monitoring Connections, Blacklists, VPN SAs, Users, and Routing (page 111)
- Using the Domain Overview (page 290)

Reading Component Statuses

The status of the system components and monitored third-party components is indicated by colors in most views where the elements are displayed. In addition, the status of various system communications is shown in color in monitoring diagrams and in the Info panel. Hardware malfunction is indicated with special icons (with an exclamation mark) in the System Status tree. If any problems are indicated, always check the logs and alerts to get a thorough view of what may be causing the problems.

The status colors are explained in the tables below:

- Engine Hardware Malfunction Icons
- Replication Malfunction Icon
- Element Status Colors (page 101)
- Node Status Colors (page 101)
- NetLink Status Colors (page 102)
- VPN Status Colors (page 102)
- Connectivity Status Colors (page 103)

Engine Hardware Malfunction Icons

Engine hardware malfunction is indicated with a special icon on top of the affected engine's icon in the Status tree and on all the top-level branches in the Status tree. For more information about the hardware malfunction, select the engine in the Status tree and switch to the Appliance Status tab in the Info panel.

Table 7.1 Hardware Malfunction Icons

Icon	Hardware Status	Description
Exclamation mark on a yellow triangle ⚠	Warning	A pre-defined Warning level has been reached in hardware monitoring (for example, the remaining file system capacity is less than 15%). The system also generates a log entry.
Exclamation mark on a red circle ❗	Alert	A pre-defined Alert level has been reached in hardware monitoring (for example, the remaining file system capacity is less than 5%). The system also generates an alert entry.

Replication Malfunction Icon

In an environment with more than one Management Server, the configuration data is replicated to all Management Servers that are online. If the replication of configuration data among the Management Servers fails, a yellow triangle with an exclamation mark in the center (see Table 7.1 for the Warning status icon) is displayed on top of the Management Server's icon in the Status tree for each Management Server that is not synchronized with the other Management Servers.

Element Status Colors

The element-level status gives an overview of the status of all engine nodes that are represented by the element (also shown for single-node components).

Table 7.2 Element-Level Status

Color	Element Status	Description
Green	All OK	All nodes have a normal status (online or standby).
Yellow	Warning	Some of the nodes have an abnormal status or have been commanded offline, but are still sending status updates normally.
Red	Alert	All of the nodes have an abnormal status, there are one or more nodes that have not sent expected status updates, or all nodes have been commanded offline.
Gray	Unknown status	No policy has been installed on any of the nodes.
White	Not monitored	An administrator has disabled monitoring for all of the nodes.

Node Status Colors

The node status gives more detailed information about individual engines.

Table 7.3 Node-Level Status

Color	Node Status	Description
Green	Node or server online	The node or server is online.
Green (with slot)	Locked online	The node is locked online to prevent automatic status transitions. The node will not change state unless commanded by an administrator.
Cyan	Standby mode	Used with clustered engines when the cluster is in Standby mode. The node is in standby mode. One of the standby nodes goes online when the previous online node goes to a state in which it does not process traffic.
Blue	Node offline	The node is offline and does not process traffic.
Blue (with slot)	Locked offline	The node is locked offline to prevent automatic status transitions. The node will not change state unless commanded by an administrator.
Gray	Timeout or unknown status	The Management Server does not know the status of the node.
White	Not monitored	An administrator has disabled monitoring for the node.

NetLink Status Colors

NetLink status shows the status of the network links in a Multi-Link configuration.

Note – The NetLink elements are queried and the status is displayed only if Probing settings are configured in the NetLink elements and if the Outbound Multi-Link element is included in the engine configuration (an outbound balancing NAT rule is included in the policy).

Table 7.4 NetLink Status Icons

Color	NetLink Status	Description
Green	OK	The NetLink is up.
Gray	Unknown status	The Management Server does not know the status of the NetLink.
White	Not monitored	An administrator has disabled monitoring for the NetLink.

VPN Status Colors

The VPN status shows the health of the VPN tunnels.

Table 7.5 VPN Status

Color	Cluster Status	Description
Green	Tunnels up	All tunnels have a normal status (online or standby) and there is traffic.
Yellow	Warning	An error was detected, at least for some traffic, but the tunnels are usable in principle, and some other traffic may be getting through.
Red	Error	Some or all tunnels are down.
Blue	Idle	The tunnels are valid, but there has not been recent traffic.
White	Not configured	The VPN has no valid tunnels, because the VPN configuration is not complete or does not allow any valid tunnels.

Connectivity Status Colors

Element diagrams and the Connectivity tab in the Info panel shows the current status of the connectivity between system elements. See the tooltip for the status color for more details.

Table 7.6 Connectivity Status

Color	Status	Explanation
Green	OK	The connection is active (there have been communications within the last 2 minutes) and no problems have been reported.
Red	Error	The Management Server received a report that connection attempts have failed.
Cyan	Idle	Connection between components is still open but there is a pause in communications.
Yellow	Warning	There are problems with heartbeat or state synchronization between nodes in a Firewall Cluster. Only one of the interfaces used for heartbeat or state synchronization functions properly. This does not affect how the cluster functions.
Blue	Closed	The connection was closed by one of the components.
Gray	Timeout, Unknown	The Management Server does not know the status of the connection. The connection may or may not be working.

Creating Overviews

Prerequisites: None

Customizable Overviews can contain information on the system's status, shortcuts to views you use often (such as logs filtered with specific criteria) and statistical charts on the system's operation (such as engine load), and the traffic flow.

You can create new Overviews on an empty template or start your customization based on one of the default Overview templates in the system.

Illustration 7.8 Example Overview with System Summary and Statistics

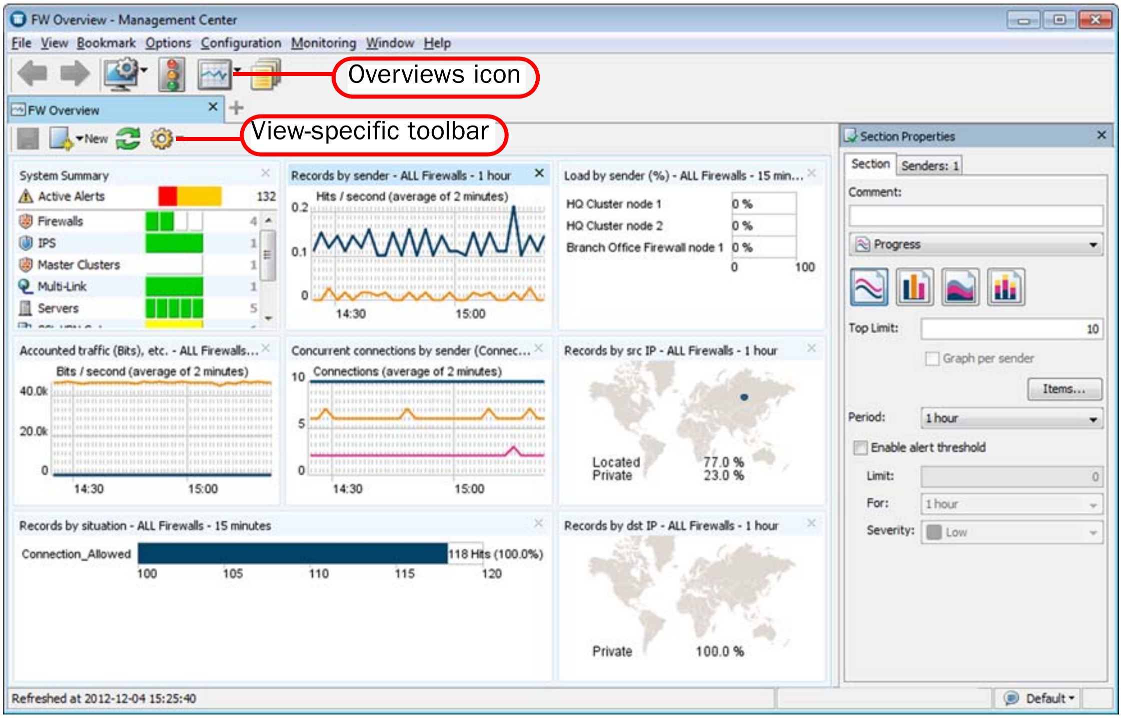
FW Overview - Management Center
File View Bookmark Options Configuration Monitoring Window Help
FW Overview
New
System Summary
Active Alerts 132
Firewalls 4
IPS 1
Master Clusters 1
Multi-Link 1
Servers 5
Records by sender - ALL Firewalls - 1 hour
Hits / second (average of 2 minutes)
0.2
0.1
0
14:30
15:00
Load by sender (%) - ALL Firewalls - 15 min...
HQ Cluster node 1 0 %
HQ Cluster node 2 0 %
Branch Office Firewall node 1 0 %
0
100
Accounted traffic (Bits), etc. - ALL Firewalls...
Bits / second (average of 2 minutes)
40.0k
20.0k
0
14:30
15:00
Concurrent connections by sender (Connec...
Connections (average of 2 minutes)
10
5
0
14:30
15:00
Records by src IP - ALL Firewalls - 1 hour
Located 77.0 %
Private 23.0 %
Records by situation - ALL Firewalls - 15 minutes
Connection_Allowed
118 Hits (100.0%)
100
105
110
115
120
Records by dst IP - ALL Firewalls - 1 hour
Private 100.0 %
Section Properties
Section
Senders: 1
Comment:
Progress
Top Limit: 10
Graph per sender
Items...
Period: 1 hour
Enable alert threshold
Limit: 0
For: 1 hour
Severity: Low
Refreshed at 2012-12-04 15:25:40
Default

Overviews icon

View-specific toolbar

What's Next?

- ▶ Creating a New Overview
  - ▶ Adding a New System Summary Section to an Overview (page 105)
  - ▶ Adding a New Statistics Section to an Overview (page 106)

Creating a New Overview

▼ To create an Overview

1. Select **Monitoring→Overviews**.
2. Right-click **Overviews** and select **New Overview**. The Overview Properties dialog opens.
3. Select the Overview template:
   - If you want to add your own Overview sections to an empty grid, select **Empty Overview** and click **OK**.
   - To use one of the predefined Overview templates, select the template from the list and click **OK**. The selected Overview template opens.

What's Next?

- ▶ To add a summary of the system status to the Overview, see Adding a New System Summary Section to an Overview.
- ▶ To add statistics to the Overview, see Adding a New Statistics Section to an Overview (page 106).

Adding a New System Summary Section to an Overview

The system summary is shown in the default start view (see System Summary (page 96)), but you can also add it to your own Overviews. It is possible to add more than one system summary to the same overview, but the information displayed is always the same.

▼ To create a new System Summary Section

1. Open the Overview to which you want to add a new System Summary Section.
2. Click the New icon in the toolbar and select **System Summary Section**. A new System Summary Section is added to the Overview.
3. (*Optional*) Adjust the placement and size of the new section by dragging and dropping the section or its edges. Resizing is based on preset grid positions and you must drag the edge until it snaps to the next position on screen for the resize to work.

What's Next?

- ▶ If you want to add statistics to the overview, see Adding a New Statistics Section to an Overview (page 106).
- ▶ If you have finished working on your Overview, click the Save icon or select **File→Save as** to save the modified Overview under a different name.

Adding a New Statistics Section to an Overview

▼ To add a new Statistics Section

1. Open the Overview to which you want to add a new Statistics Section.
2. Click the New icon in the toolbar and select a section from the list.
    - If you cannot find an appropriate section in the list, select **Other**, and select a section in the **Select Section** dialog.
    - If you want to add a section based on a statistical item, select **Create from Item** and **Select** an item from the Select Item dialog.
3. Define the basic Section properties in the **Section Properties** panel as explained in the table below:

Option		Explanation
Comment (*Optional*)		A comment for your own reference.
Statistics Type	Progress	Produces a chart and/or a table that shows the progress of items as time passes.
	Top Rate	Produces a chart and/or a table highlighting the values with the highest number of occurrences.
	Summary Table	Produces a table with all the values from items included in the report.
Diagram Type		Select the Diagram Type by clicking the appropriate diagram icon (Curve, Bar, Stacked Curve, Stacked Bar, or Map). The available diagram types depend on the type of statistics you have selected. The Summary Table type statistics are always displayed as a table, so this option is disabled for that Diagram Type.
Top Limit (*Only if Statistics Type is Progress or Top Rate*)		Defines the (maximum) number of sources shown in the chart. (Not available if statistics type is Summary Table.)
Graph per sender (*Optional*)		If selected, a separate graph is shown for each separate sender. Otherwise, the graph shows a single combined graph with the average data from all the senders.
Items (*Optional*)		Allows you to add and remove statistical items for a section. See Selecting Statistical Items (page 108).
Period		Defines the time scale for the display of information.
Enable alert threshold (*Optional*)		Allows you to define a limit for automatic tracking of monitored items. See Setting Thresholds for Monitored Items (page 109).

4. *(Optional)* Switch to the **Senders** tab and select which elements are shown in the section.

5. (*Optional*) Adjust the placement and size of the new section by dragging and dropping the section or its edges. Note that resizing is based on preset grid positions and you must drag the edge until it snaps to the next position on the screen for the resize to work.

What's Next?

▶ If you have finished working on your Overview, click the Save icon or select **File→Save as** to save the modified Overview under a different name.

Creating a New Statistics Section

You can save a section you have customized in one Overview as a Statistics Section to allow you to create the same type of section with the same settings in other Overviews.

▼ To create a new Statistics Section

1. Open the Overview to which you want to add a new Statistics Section.
2. Right-click any section in the Overview and select **Save As New Section**. The Section Properties dialog opens.
3. Define the basic Section properties in the General tab of the Section Properties dialog as explained in the table below.

Option	Explanation
Name	Enter a name for the new section.
Comment (*Optional*)	Enter a comment for your reference.
Filter (*Optional*)	Click **Edit** to define a local filter or add a Filter Element for the section. See Creating and Editing Local Filters (page 187).
Log type	Select the log type for the section.
Related Element (*Optional*)	Select a related element for the section.
Export	Select the export format for the data of the section.
Traffic Unit	Select the Traffic Unit for the data in charts and tables.

4. Select the diagram settings in the **Visualization** tab:
   - **Progress** produces a chart and/or a table that shows the progress of items as time passes.
   - **Top Rate** produces a chart and/or a table highlighting the values with the highest number of occurrences.
   - **Drill-down Top Rate** preprocesses the data and then creates a chart and/or table from it.
   - **Summary Table** produces a table with all the values from items included in the section.
   - **System Information** produces a table with information about current system state.
5. Select the diagram type. You can change this selection when you include the section in an Overview. The options available depend on the diagram settings you made in the previous step.

- If you selected **Top Rate** as the diagram type, enter the number of items to be included in the **Top Limit** field.
- If you selected **Progress** as the diagram type, you can select **Graph per sender**.

6. *(Optional)* Switch to the **Items** tab and select and/or remove statistics items for the section.
7. *(Optional)* Switch to the **Senders** tab and select which elements are shown in the section.
8. Click **OK**.

You can now add the new section to other Overviews as instructed in Adding a New Statistics Section to an Overview (page 106).

What's Next?

▶ If you have finished working on your Overview, click the Save icon or select **File→Save as** to save the modified Overview under a different name.

Selecting Statistical Items

Statistics process and visualize data. They help you to focus on the most relevant information in the system (for example, notable quantities and trends in network traffic) and also to find changes and anomalies in network traffic. You can use statistics in Overviews and Reports, and when you browse Logs or Connections. Filters are available to help you find information.

Statistical items count log entries (referred to as *records* in the item names), summaries of log fields included in those entries (like traffic volumes in log data that contain accounting information), or specifically gathered statistical data (counter items). The items are organized based on the component types, as the runtime data they produce is different.

▼ To add a new Statistical Item to a Section

1. In the Section Properties panel, click **Items**. The Properties dialog opens.
2. Click **Add**. The Select Item dialog opens.
3. **Select** the items. The table below explains the items in the Select Item dialog.

Table 7.7 Options in Select Item Dialog

Column		Explanation
Name		Name of the item.
Type	Progress	Produces a chart and/or a table that shows the progress of items as time passes.
	Top Rate	Produces a chart and/or a table highlighting the values with the highest number of occurrences.
Data	Logs	Items that use log data (log events).
	Counter	Items that use pre-counted statistical data.
	Database	Items that collect information regarding the system's operating state.

Table 7.7 Options in Select Item Dialog (Continued)

Column		Explanation
Count		The value that is counted for a specific period in a curve chart or per sector in a pie chart (for example, the number of connections or bytes).
Ranked by		The basis for ordering the items in the summary.
Method	Cumulative	Statistical data presented as cumulative sums.
	Average	Loads and connections presented as average values.
Context		The context that the item belongs to.

4. Back in the Properties dialog, click **OK**. The items are added to the section and their data is displayed in the section.

▼ To remove a Statistical Item from a Section

1. In the Section Properties panel, click **Items**. The Properties dialog opens.
2. Select the item that you want to remove from the section and click **Remove**.
3. Back in the Properties dialog, click **OK**. The item is removed from the section.

What's Next?

- If you have finished working on your Overview, click the Save icon or select **File→Save as** to save the modified Overview under a different name.

Setting Thresholds for Monitored Items

Thresholds activate automatic tracking of monitored items in Overviews. The values of the monitored items are checked once an hour.

For Progress items, the total of the values is compared to the threshold limit. The threshold is considered exceeded if the average level of the curve is above the threshold limit during the tracking period.

For Top Rate and Top Rate Curve items, the highest value is compared to the threshold limit. The threshold is considered exceeded if the highest value is above the threshold limit during the tracking period. If the threshold limit is exceeded, an alert is sent.

▼ To set a threshold for a monitored item

1. Select the Overview section for which you want to set a threshold. The Section Properties of the selected section are displayed.

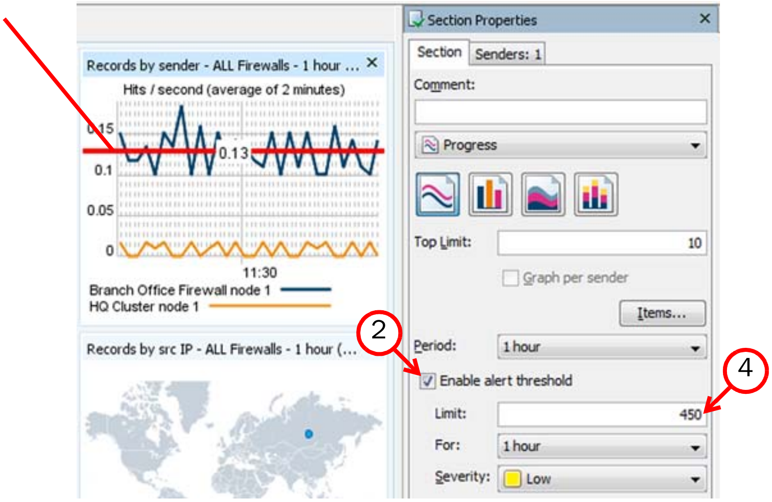

Threshold line

Records by sender - ALL Firewalls - 1 hour ...
Hits / second (average of 2 minutes)
0.15
0.13
0.1
0.05
0
11:30
Branch Office Firewall node 1
HQ Cluster node 1
Records by src IP - ALL Firewalls - 1 hour (...
Section Properties
Section
Senders: 1
2
4

Comment:

Progress

Top Limit: 10

□ Graph per sender

Items...

Period: 1 hour

☑ Enable alert threshold

Limit: 450

For: 1 hour

Severity: Low

2. Select **Enable alert threshold**.
3. Specify the threshold limit in one of the following ways:
   - Drag and drop the threshold line in the overview section.
   - Enter the **Limit** as a number in the Section Properties panel.
4. (*Optional*) Select the tracking period during which monitored items are compared to the threshold limit from the **For** list. By default, 1 hour is selected.
5. (*Optional*) Select the **Severity** of the alert that is sent when the threshold limit is exceeded. By default, Low is selected.

What's Next?

▶ If you have finished working on your Overview, click the Save icon or select **File→Save as** to save the modified Overview under a different name.

Monitoring Connections, Blacklists, VPN SAs, Users, and Routing

Prerequisites: None

Firewalls keep track of allowed connections, active VPN SAs, active users, and routing. Firewall, Layer 2 Firewall, and IPS engines also keep track of combinations of IP addresses, ports, and protocols that are currently included on the blacklist of the Firewall, Layer 2 Firewall, and IPS engines.

Note – To be able to monitor users by name, you must enable the logging of user information in the Firewall IPv4 Access rules. See Defining Access Rule Logging Options (page 684).

You can monitor connections, blacklists, VPN SAs, users, and routing in the following ways:

- You can view currently open connections, enforced blacklist entries, active VPN SAs, active users, and routing. See Checking Connections, Blacklists, VPN SAs, Users, and Routing.
- You can save snapshots of currently open connections, enforced blacklist entries, active VPN SAs, active users, and routing. See Saving Snapshots of Connections, Blacklists, VPN SAs, Users, and Routing (page 113).
- You can view and compare snapshots of currently open connections, enforced blacklist entries, active VPN SAs, active users, and routing. See Comparing Snapshots of Connections, Blacklists, VPN SAs, Users, and Routing (page 115).

Checking Connections, Blacklists, VPN SAs, Users, and Routing

There are several views in which you can monitor the current status of the system.

- The Connections view shows you the currently active connections.
- The Blacklist view shows you the current blacklist entries.
- The VPN SAs view shows you all active VPN tunnels.
- The Users view shows you all currently active users.
- The Routing view shows you current static and dynamic routes.

The lists of connections, blacklisted entries, active VPN SAs, active users, and routes are automatically updated in the Current Events mode. The blacklist entries are added and removed according to their duration.

The Blacklist view does not show whether connections matching the entries are actually blocked by the Firewall, Layer 2 Firewall, or IPS engine, nor can you see any history of entries that have already expired. Use the Logs view to see information on actual connections that are allowed and denied, as well as logs on past blacklist entry creation (depending on logging options selected in the policy). See Getting Started with the Logs View (page 148).

▼ To check connections, blacklist entries, active VPN SAs, active users, or routing

1. Select **Monitoring→Connections**, **Monitoring→Blacklist**, **Monitoring→VPN SAs**, **Monitoring→Users**, or **Monitoring→Routing**. The Select Element dialog opens.
2. Select the element depending on the type of information that you want to view:
   - Select the Firewall engine for which you want to view open connections, active VPN SAs, active users, or routing.
   - Select the Firewall, Layer 2 Firewall, or IPS engine for which you want to view the blacklist.
3. Click **Select**. The Connections view, the Blacklist view, the VPN SAs view, the Users view, or the Routing view opens.

Illustration 7.9 Toolbar and Query Panel in Monitoring Views

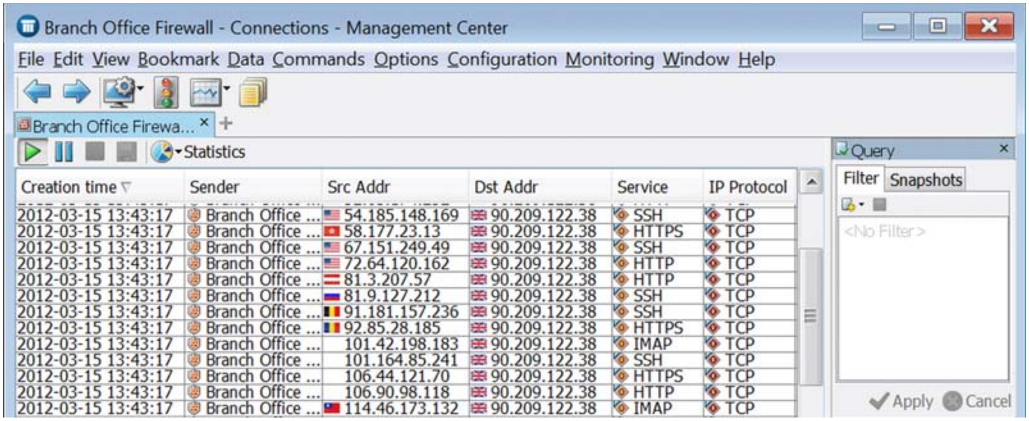
Branch Office Firewall - Connections - Management Center
File Edit View Bookmark Data Commands Options Configuration Monitoring Window Help
Branch Office Firewa...
Statistics
Query
Filter
Snapshots
<No Filter>
Apply
Cancel

Creation time	Sender	Src Addr	Dst Addr	Service	IP Protocol
2012-03-15 13:43:17	Branch Office ...	54.185.148.169	90.209.122.38	SSH	TCP
2012-03-15 13:43:17	Branch Office ...	58.177.23.13	90.209.122.38	HTTPS	TCP
2012-03-15 13:43:17	Branch Office ...	67.151.249.49	90.209.122.38	SSH	TCP
2012-03-15 13:43:17	Branch Office ...	72.64.120.162	90.209.122.38	HTTP	TCP
2012-03-15 13:43:17	Branch Office ...	81.3.207.57	90.209.122.38	HTTP	TCP
2012-03-15 13:43:17	Branch Office ...	81.9.127.212	90.209.122.38	SSH	TCP
2012-03-15 13:43:17	Branch Office ...	91.181.157.236	90.209.122.38	SSH	TCP
2012-03-15 13:43:17	Branch Office ...	92.85.28.185	90.209.122.38	HTTPS	TCP
2012-03-15 13:43:17	Branch Office ...	101.42.198.183	90.209.122.38	IMAP	TCP
2012-03-15 13:43:17	Branch Office ...	101.164.85.241	90.209.122.38	SSH	TCP
2012-03-15 13:43:17	Branch Office ...	106.44.121.70	90.209.122.38	HTTPS	TCP
2012-03-15 13:43:17	Branch Office ...	106.90.98.118	90.209.122.38	HTTP	TCP
2012-03-15 13:43:17	Branch Office ...	114.46.173.132	90.209.122.38	IMAP	TCP

There are many ways to browse the data:

- Click the appropriate icon in the toolbar:
  - Click Play to view current entries.
  - Click Pause to create a temporary snapshot of entries.
  - Click Save to save a snapshot.
  - Click Statistics to view statistical items.
  - Click Visualization to view a diagram of the data (only available in the Routing view).
- Select the **Filter** tab or the **Snapshots** tab to filter data or open/compare snapshots.
- You can adjust the view and the displayed data in a similar way as in the Logs view. See Browsing Log Data (page 155).
- To combine entries based on entries that have the same value for a given column type, right-click the heading of the corresponding column and select **Aggregate by** <column name>. Repeat the action if you want to aggregate entries according to other column types.
- To view data as charts, click the **Statistics** button in the toolbar, and select one of the pre-defined statistical items, or select **Other** to create a custom statistical element. For more information, see Selecting Statistical Items (page 108).
  - The **Statistics** button is not available in the Users and Routing views.
  - The **Other** statistical items are not available in the VPN SAs view.
- Select one or more entries in the table and right-click for a menu of actions you can take:
  - Select **Show Referenced Events** to view more information on related log events.
  - Select **New Blacklist Entry** or **New Entry** to blacklist connection(s) manually. See Blacklisting Connections Manually (page 228).
  - Select **Terminate** to terminate a connection in the Connections view.

- Select **Remove Entry** to remove a blacklist entry in the Blacklist view.
- Select **Delete** to force an SA to renegotiate in the VPN SAs view.
- Select **Delete** to close the end-user’s session in the Users view.

Related Tasks

▶ Getting Started with Blacklisting (page 816)

Saving Snapshots of Connections, Blacklists, VPN SAs, Users, and Routing

You can save snapshots of open connections, blacklist entries, active VPN SAs, active users, and routing in the appropriate views. The saved snapshots are listed under **Other Elements→Monitoring Snapshots** in the Monitoring view.

The saved snapshots are stored in the following directories:

Table 7.8 Snapshot Storage Directories

Snapshot Type	Server	Directory
Blacklist snapshots	Log Server	*<installation directory>*`/data/storage/snapshots/blacklist`
Connection snapshots	Log Server (Firewall engines 5.2 or higher)/ Management Server (all earlier Firewall engine versions)	*<installation directory>*`/data/storage/snapshots/connections`
VPN SA snapshots	Log Server	*<installation directory>*`/data/storage/snapshots/VPN SAs`
User snapshots	Log Server	*<installation directory>*`/data/storage/snapshots/users`
Routing snapshots	Log Server	*<installation directory>*`/data/storage/snapshots/routing`

Note – If you installed the Management Server in the `C:\Program Files\Stonesoft\Management Center` **directory in Windows, some program data may be stored in the** `C:\ProgramData\Stonesoft\Management Center\data` **directory.**

▼ To save snapshots

1. Select **Monitoring→Blacklist**, **Monitoring→Connections**, **Monitoring→Logs**, **Monitoring→Routing**, **Monitoring→Users**, or **Monitoring→VPN SAs**. The Select Element dialog opens.
2. Select the element depending on the snapshot that you want to save:
   - Select the Firewall engine for which you want to view open connections, active VPN SAs, active users, or routing.
   - Select the Firewall, Layer 2 Firewall, or IPS engine for which you want to view the blacklist.
3. Click the Pause icon in the toolbar to select the entries for the snapshot.
   - When you click Pause, the system automatically creates a temporary snapshot of the currently displayed entries. The name of the temporary snapshot is displayed on the Snapshots tab in the Query panel. The temporary snapshot is automatically deleted.
4. Click the Save icon in the toolbar. The Snapshot Properties dialog opens.
5. Enter a **Name** for the snapshot and click **OK**. The name of the snapshot is displayed on the Snapshots tab in the Query panel.

Exporting Snapshots of Connections, Blacklists, VPN SAs, Users, and Routing

You can export snapshots from the Monitoring view to save stored snapshots elsewhere (for example, on your local workstation).

▼ To export snapshots

1. Select **Configuration→Configuration→Monitoring**. The Monitoring view opens.
2. Browse to one of the following branches under **Other Elements→Monitoring Snapshots** in the Monitoring tree:
   - **Blacklist→Log Server**
   - **Connections→Management Server** (snapshots saved for Firewall engine version 5.1 or earlier).
   - **Connections→Log Server**
   - **Logs→Management Server**
   - **Routing→Log Server**
   - **Users→Log Server**
   - **VPN SAs→Log Server**
3. Right-click a snapshot and select **Export**.
4. Select the location for saving the snapshot and click **Select**.

Related Tasks

▶ Comparing Snapshots of Connections, Blacklists, VPN SAs, Users, and Routing (page 115)

Viewing Snapshots of Connections, Blacklists, VPN SAs, Users, and Routing

The snapshots of connections, blacklists, VPN SAs, users, and routing are listed in the Monitoring view. You can also open snapshots saved on your local workstation.

▼ To view a snapshot stored on the Log Server or the Management Server

1. Select **Configuration→Configuration→Monitoring**. The Monitoring view opens.
2. Browse to one of the following branches under **Other Elements→Monitoring Snapshots** in the Monitoring tree:
   - **Blacklist→Log Server**
   - **Connections→Management Server** (snapshots saved for Firewall engine version 5.1 or earlier).
   - **Connections→Log Server**
   - **Logs→Management Server**
   - **Routing→Log Server**
   - **Users→Log Server**
   - **VPN SAs→Log Server**
3. Right-click a snapshot and select **Open**. The snapshot opens for viewing.

▼ To view a snapshot from your own workstation

1. Select **Configuration→Configuration→Monitoring**. The Monitoring view opens.
2. Browse to the **Other Elements→Monitoring Snapshots**.
3. Right-click **Connections, Blacklist, VPN SAs, Users**, or **Routing** and select **Open Local Snapshot**. The Select Snapshot File dialog opens.
4. Select the snapshot and click **Open**. The snapshot opens for viewing.

Related Tasks

▶ Monitoring Connections, Blacklists, VPN SAs, Users, and Routing (page 111)

Comparing Snapshots of Connections, Blacklists, VPN SAs, Users, and Routing

You can compare snapshots of connections, blacklists, VPN SAs, users, or routing with another snapshot of the same type. You can also compare a snapshot with the current blacklist, connections, VPN SAs, users, or routing entries.

The saved snapshots are listed under **Other Elements→Monitoring Snapshots** in the Monitoring view.

▼ To compare snapshots

1. Open the Blacklist view, Connections view, Routing view, Users view, or VPN SAs view as explained in Monitoring Connections, Blacklists, VPN SAs, Users, and Routing (page 111).

2. Switch to the Snapshots tab in the Query panel.

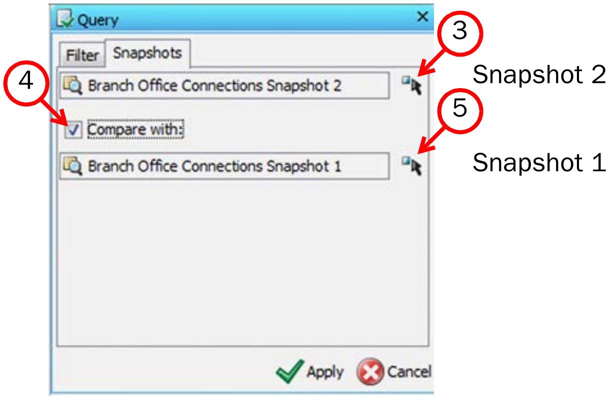

Branch Office Connections Snapshot 2

☑ Compare with:

Branch Office Connections Snapshot 1

3. Select the first snapshot for comparison:
   - If you want to compare the current entries with a snapshot, make sure that the view is in Current Events mode and that Current is displayed in the first snapshot selection field.
   - If you want to compare a temporary snapshot of the current entries with a saved snapshot, click **Pause** in the toolbar to create a temporary snapshot. The name of the temporary snapshot is displayed in the first snapshot selection field.
   - Otherwise, click the Select icon next to the first snapshot selection field and browse to the correct branch to select the first snapshot for comparison.
4. Select **Compare with** to enable comparing snapshots.
5. Click the Select icon next to the second snapshot selection field and browse to the correct branch to select the second snapshot for comparison.
   - You can use drag and drop to change the order of the snapshots in the comparison.
6. Click **Apply** to start the comparison. The results of the comparison are highlighted in the list of entries.

Illustration 7.10 Snapshot Comparison Results between Snapshots 1 and 2

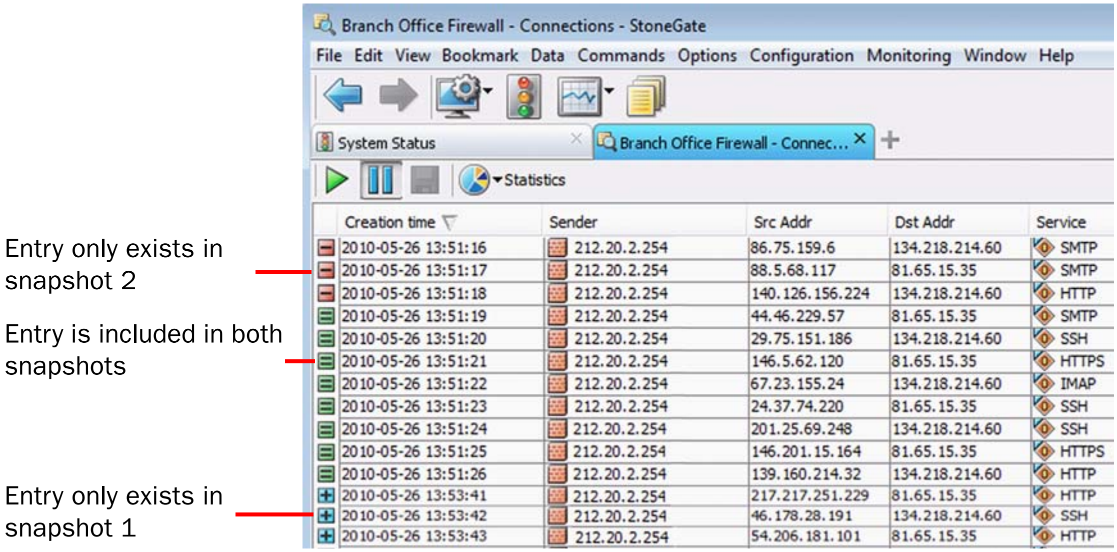
Branch Office Firewall - Connections - StoneGate
File Edit View Bookmark Data Commands Options Configuration Monitoring Window Help
System Status
Branch Office Firewall - Connec...
Statistics
Entry only exists in snapshot 2
Entry is included in both snapshots
Entry only exists in snapshot 1

	Creation time	Sender	Src Addr	Dst Addr	Service
−	2010-05-26 13:51:16	212.20.2.254	86.75.159.6	134.218.214.60	SMTP
−	2010-05-26 13:51:17	212.20.2.254	88.5.68.117	81.65.15.35	SMTP
−	2010-05-26 13:51:18	212.20.2.254	140.126.156.224	134.218.214.60	HTTP
=	2010-05-26 13:51:19	212.20.2.254	44.46.229.57	81.65.15.35	SMTP
=	2010-05-26 13:51:20	212.20.2.254	29.75.151.186	134.218.214.60	SSH
=	2010-05-26 13:51:21	212.20.2.254	146.5.62.120	81.65.15.35	HTTPS
=	2010-05-26 13:51:22	212.20.2.254	67.23.155.24	134.218.214.60	IMAP
=	2010-05-26 13:51:23	212.20.2.254	24.37.74.220	81.65.15.35	SSH
=	2010-05-26 13:51:24	212.20.2.254	201.25.69.248	134.218.214.60	SSH
=	2010-05-26 13:51:25	212.20.2.254	146.201.15.164	81.65.15.35	HTTPS
=	2010-05-26 13:51:26	212.20.2.254	139.160.214.32	134.218.214.60	HTTP
+	2010-05-26 13:53:41	212.20.2.254	217.217.251.229	81.65.15.35	HTTP
+	2010-05-26 13:53:42	212.20.2.254	46.178.28.191	134.218.214.60	SSH
+	2010-05-26 13:53:43	212.20.2.254	54.206.181.101	81.65.15.35	HTTP

Illustration 7.11 Example of Snapshot Comparison with Aggregated Entries - Aggregated by Service

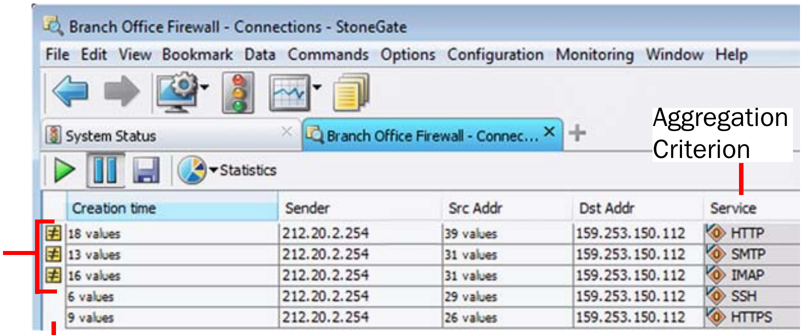
Branch Office Firewall - Connections - StoneGate
File Edit View Bookmark Data Commands Options Configuration Monitoring Window Help
System Status
Branch Office Firewall - Connec...
Statistics

	Creation time	Sender	Src Addr	Dst Addr	Service
≠	18 values	212.20.2.254	39 values	159.253.150.112	HTTP
≠	13 values	212.20.2.254	31 values	159.253.150.112	SMTP
≠	16 values	212.20.2.254	31 values	159.253.150.112	IMAP
	6 values	212.20.2.254	29 values	159.253.150.112	SSH
	9 values	212.20.2.254	26 values	159.253.150.112	HTTPS

Entries that match the **Aggregate by** <column name> selection are included in both snapshots but the entries are otherwise different.

Aggregation Criterion

(*No icon*) There are more than 100 entries that match the **Aggregate by** <column name> selection in the snapshots. No further comparison can be done.

You can alternatively open a snapshot for comparison directly in the Monitoring view by right-clicking the selected snapshot(s) and selecting **Compare to Current**.

Related Tasks

- Monitoring Connections, Blacklists, VPN SAs, Users, and Routing (page 111)
- Browsing Log Data (page 155)

Viewing and Comparing Element Snapshots

Prerequisites: None

You can view earlier configurations of an element and compare them to the current configuration with Element Snapshots. You can also restore earlier configurations of an element.

Element Snapshots are automatically generated and saved in Audit logs each time element properties are saved. An Element Snapshot contains all the properties of an element saved in the Properties dialog.

▼ To view and compare Element Snapshots

1. Select **Monitoring→Logs**.
2. Select **Audit** in the Query panel and click **Apply**. All Element Snapshots generated during the defined time range are displayed in the Snapshot column of the log entry table.

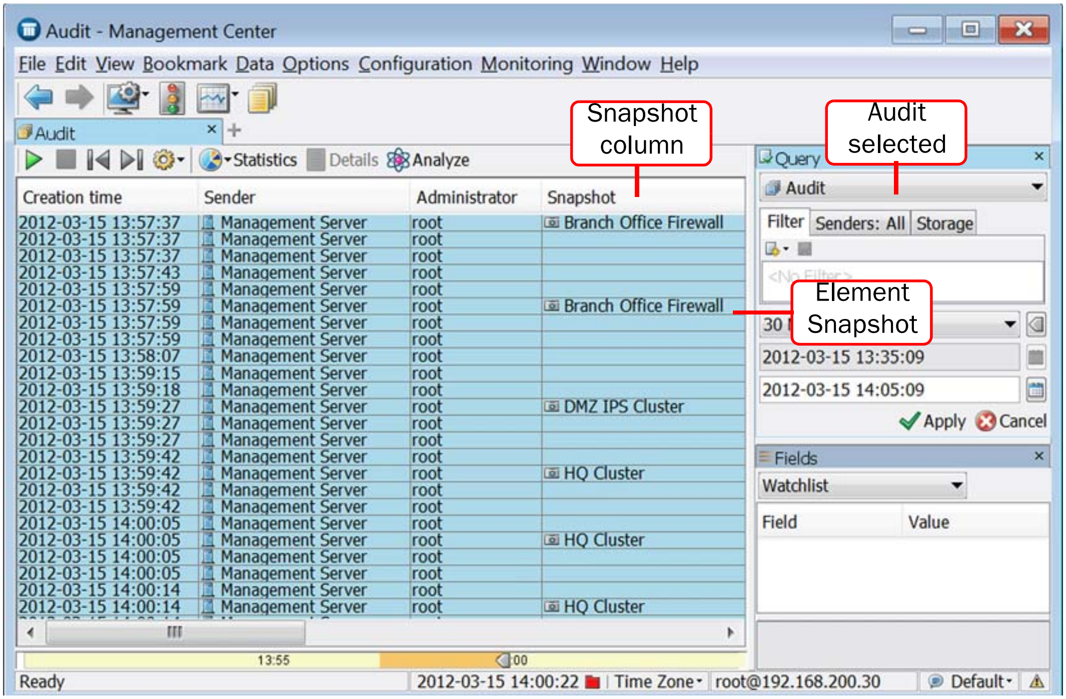

Creation time	Sender	Administrator	Snapshot
2012-03-15 13:57:37	Management Server	root	Branch Office Firewall
2012-03-15 13:57:37	Management Server	root	
2012-03-15 13:57:37	Management Server	root	
2012-03-15 13:57:43	Management Server	root	
2012-03-15 13:57:59	Management Server	root	
2012-03-15 13:57:59	Management Server	root	Branch Office Firewall
2012-03-15 13:57:59	Management Server	root	
2012-03-15 13:57:59	Management Server	root	
2012-03-15 13:58:07	Management Server	root	
2012-03-15 13:59:15	Management Server	root	
2012-03-15 13:59:18	Management Server	root	
2012-03-15 13:59:27	Management Server	root	DMZ IPS Cluster
2012-03-15 13:59:27	Management Server	root	
2012-03-15 13:59:27	Management Server	root	
2012-03-15 13:59:42	Management Server	root	
2012-03-15 13:59:42	Management Server	root	HQ Cluster
2012-03-15 13:59:42	Management Server	root	
2012-03-15 13:59:42	Management Server	root	
2012-03-15 14:00:05	Management Server	root	
2012-03-15 14:00:05	Management Server	root	HQ Cluster
2012-03-15 14:00:05	Management Server	root	
2012-03-15 14:00:05	Management Server	root	
2012-03-15 14:00:14	Management Server	root	
2012-03-15 14:00:14	Management Server	root	HQ Cluster

Query: Audit

30

2012-03-15 13:35:09

2012-03-15 14:05:09

Fields: Watchlist

Ready | 2012-03-15 14:00:22 | Time Zone | root@192.168.200.30 | Default

Tip – **If the Snapshot column is not displayed, you can add it to the log entry table by selecting Vinculum Selection. See Changing Data Columns in the Log Entry Table (page 162).**

3. To view an Element Snapshot in more detail, right-click the Audit entry and select **View Element Snapshot**. The Element Properties dialog opens.

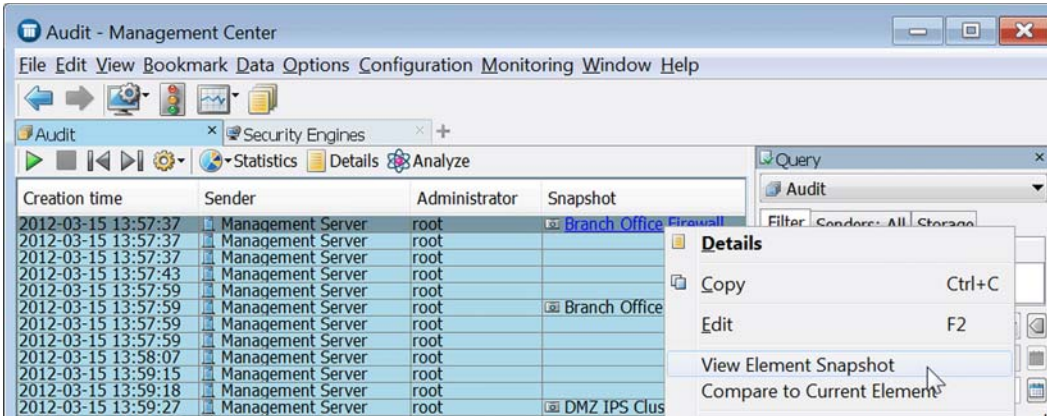

Creation time	Sender	Administrator	Snapshot
2012-03-15 13:57:37	Management Server	root	Branch Office Firewall
2012-03-15 13:57:37	Management Server	root	
2012-03-15 13:57:37	Management Server	root	
2012-03-15 13:57:43	Management Server	root	
2012-03-15 13:57:59	Management Server	root	
2012-03-15 13:57:59	Management Server	root	Branch Office
2012-03-15 13:57:59	Management Server	root	
2012-03-15 13:57:59	Management Server	root	
2012-03-15 13:58:07	Management Server	root	
2012-03-15 13:59:15	Management Server	root	
2012-03-15 13:59:18	Management Server	root	
2012-03-15 13:59:27	Management Server	root	DMZ IPS Clus

4. To compare an Element Snapshot to the current configuration of the same element, right-click an Audit entry and select **Compare to Current Element**. The Compare Elements dialog opens.
   - If the Element Snapshot properties differ from the current element properties, a red border is displayed around the Audit Log Version (the snapshot) and Current Version of the element.

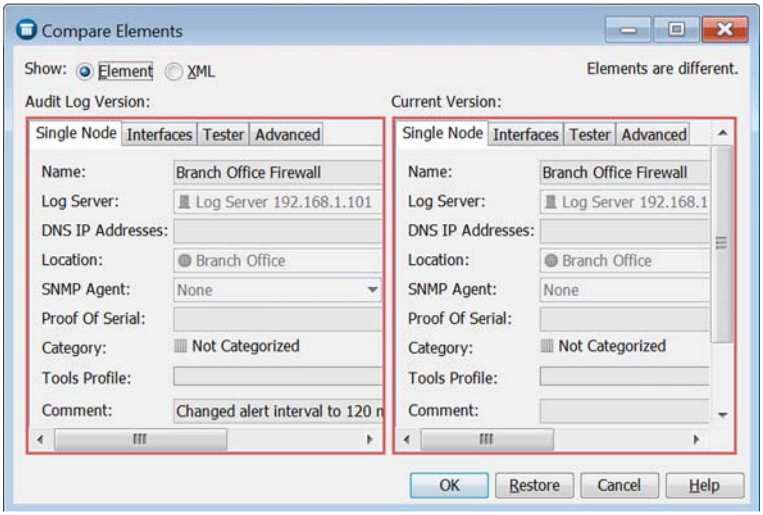

Show: ● Element ○ XML

Elements are different.

Audit Log Version:

Name:	Branch Office Firewall
Log Server:	Log Server 192.168.1.101
DNS IP Addresses:	
Location:	Branch Office
SNMP Agent:	None
Proof Of Serial:	
Category:	Not Categorized
Tools Profile:	
Comment:	Changed alert interval to 120 n

Current Version:

Name:	Branch Office Firewall
Log Server:	Log Server 192.168.1
DNS IP Addresses:	
Location:	Branch Office
SNMP Agent:	None
Proof Of Serial:	
Category:	Not Categorized
Tools Profile:	
Comment:	

- Select **Show: XML** to display all the values of the snapshot and the current element in XML format with differences indicated in red.

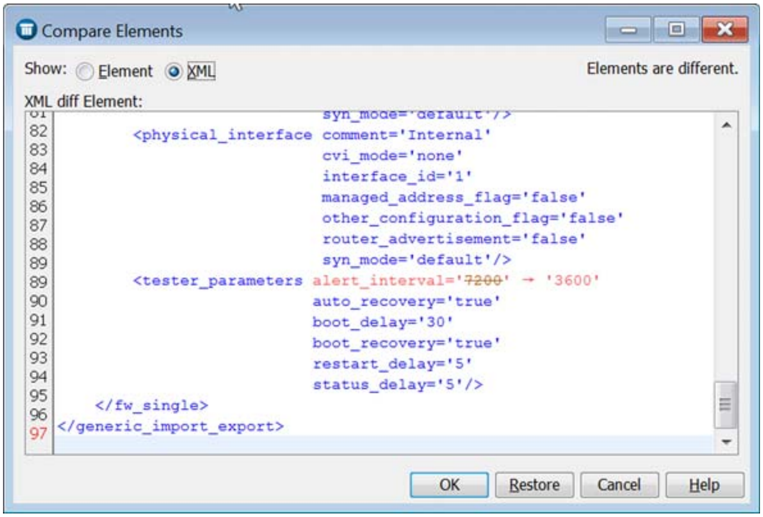

Compare Elements

Show: ○ Element ● XML

Elements are different.

XML diff Element:

```
                              syn_mode='default'/>
        <physical_interface comment='Internal'
                              cvi_mode='none'
                              interface_id='1'
                              managed_address_flag='false'
                              other_configuration_flag='false'
                              router_advertisement='false'
                              syn_mode='default'/>
        <tester_parameters alert_interval='~~7200~~' → '3600'
                             auto_recovery='true'
                             boot_delay='30'
                             boot_recovery='true'
                             restart_delay='5'
                             status_delay='5'/>
    </fw_single>
</generic_import_export>
```

5. Click **OK** to close the Compare Elements dialog.

Related Tasks

- ▶ Restoring Elements From Element Snapshots (page 86)
- ▶ Exporting, Importing, and Restoring Elements (page 79)

Monitoring Connections on a Map

Prerequisites: None

Geolocation elements can be used to define the physical locations of network elements, such as Hosts. Geolocations enable you to see on a world map where Hosts (for example, attackers) are located. You can also see, for example, how much traffic they create.

The location information for public IP addresses is based on an internal Geolocation database on the Management Server. The IP addresses in the Geolocation database are updated when the Management Server is updated.

You can also define Geolocations manually in the SMC. Geolocations based on internal IP addresses must be configured manually, as these addresses cannot be found in the Geolocation database based on public IP addresses. Defining a Geolocation manually for a public IP address overrides location data found for the address in the Geolocation database. See Defining a New Geolocation (page 121).

Geolocation maps are available in Reports, Statistics, and Overviews. Geolocation elements can also be displayed on a world map in the System Status view.

Illustration 7.12 Geolocations in the System Status view

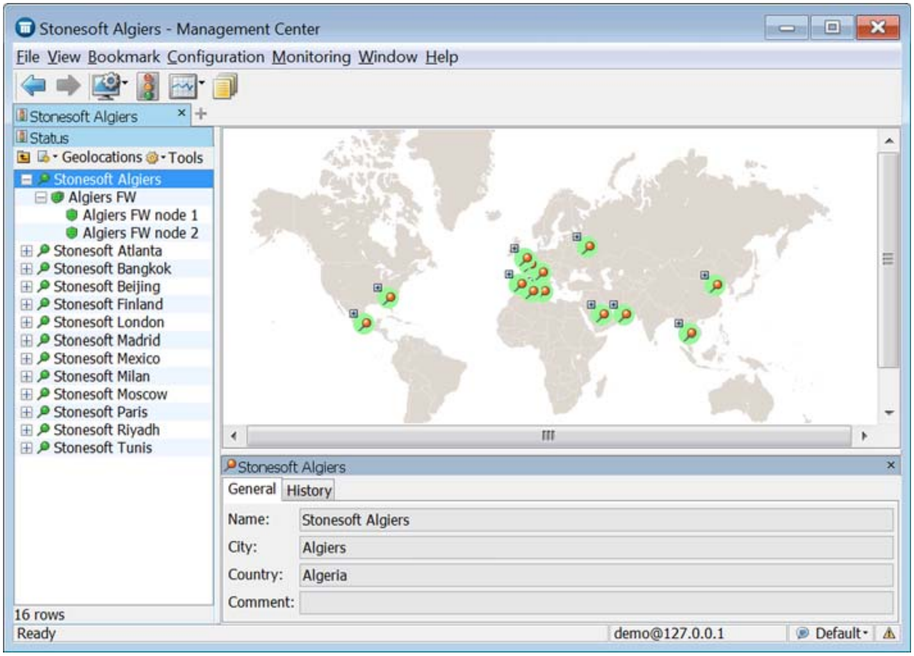

Name: Stonesoft Algiers

City: Algiers

Country: Algeria

Comment:

▼ To view Geolocation elements in the System Status view

1. Click the Up icon above the System Status tree and select **Geolocations**.
   - A Geolocation is shown in the Status tree only if it has been set for an element. See Setting a Geolocation for an Element (page 123).
2. Select a Geolocation element to view its location on the world map.

Tip – You can also see the actual location of a Geolocation or IP address in Google Maps. See Viewing Geolocations and IP Addresses in Google Maps (page 123).

Defining a New Geolocation

▼ To define a new Geolocation

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Expand **Other Elements**.
3. Right-click **Geolocations** and select **New Geolocation**. The Geolocation Properties dialog opens.

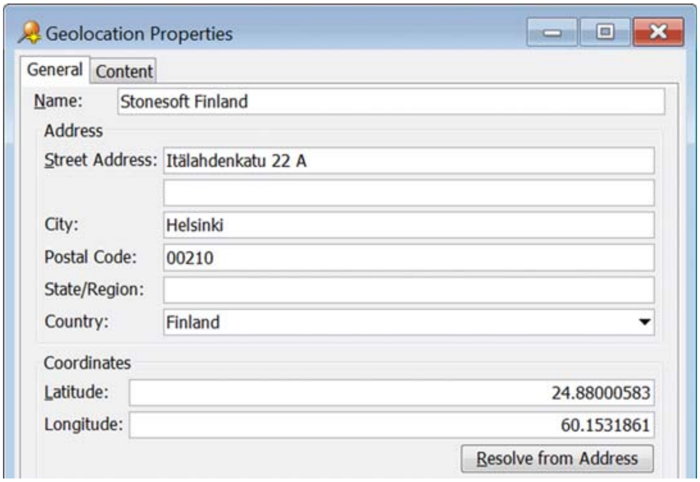

Name: Stonesoft Finland

Street Address: Itälahdenkatu 22 A

City: Helsinki

Postal Code: 00210

State/Region:

Country: Finland

Latitude: 24.88000583

Longitude: 60.1531861

4. Enter the **Name** and **Address**.
5. Define the Coordinates in one of the following ways:
   - Click **Resolve from Address** to automatically resolve the Geolocation coordinates.
   - Enter the **Latitude** and **Longitude** in Decimal Degrees format (for example, latitude 49.5000° and longitude -123.5000°). Switch to the **Content** tab to define which elements belong to the Geolocation.

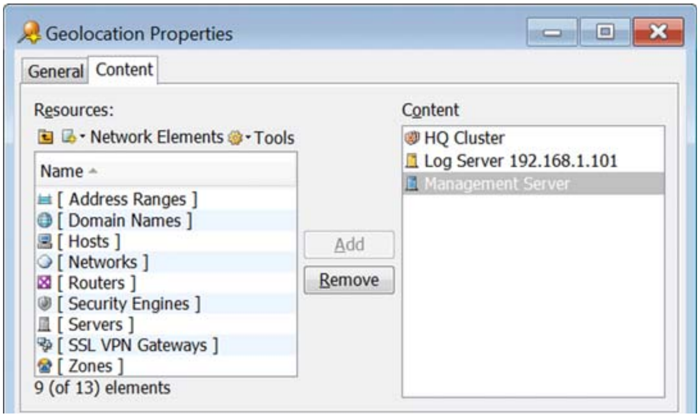

Resources:

- [Address Ranges]
- [Domain Names]
- [Hosts]
- [Networks]
- [Routers]
- [Security Engines]
- [Servers]
- [SSL VPN Gateways]
- [Zones]

9 (of 13) elements

Content:

- HQ Cluster
- Log Server 192.168.1.101
- Management Server

6. Switch to the **Content** tab and select the elements that belong to the Geolocation.
7. Click **OK**.

Related Tasks

- ▶ Setting a Geolocation for an Element (page 123)

Setting a Geolocation for an Element

Prerequisites: Defining a New Geolocation

▼ To set a Geolocation for an element

1. Right-click an element and select **Tools→Set Geolocation**. The Select Geolocation dialog opens.
2. Select a Geolocation for the element and click **Select**.

Related Tasks

- ▶ Defining a New Geolocation (page 121)
- ▶ Viewing Geolocations and IP Addresses in Google Maps

Viewing Geolocations and IP Addresses in Google Maps

Prerequisites: Defining a New Geolocation

You can view the actual location of a Geolocation element or an IP address in more detail in Google Maps. You can use the Show in Google Maps option in:

- Geolocation maps
- The Logs view
- The Whois Information dialog.

Viewing Geolocation Element Locations in Overviews and Reports

Prerequisites: Defining a New Geolocation

Geolocation elements are displayed in Google Maps based on the location data that was entered for them.

▼ To view the location of a Geolocation element in an Overview or Report

↪ Right-click a location in a Geolocation map in an Overview or Report section and select **Show in Google Maps**. The location is opened in Google Maps in your default browser.

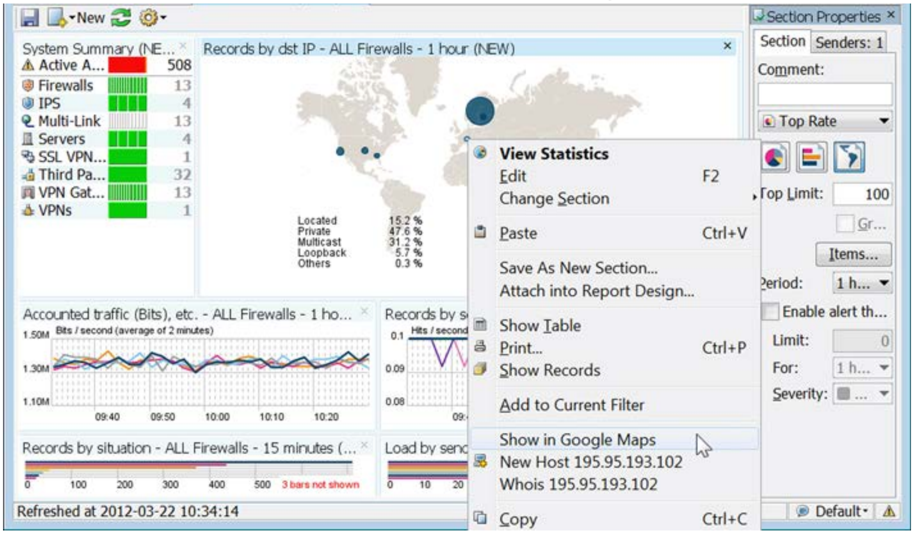

Related Tasks

- ▶ Creating Overviews (page 104)
- ▶ Creating and Modifying Report Designs (page 169)

Viewing IP Address Locations in the Logs view

Prerequisites: Defining a New Geolocation

Only IP addresses associated with a location (for example, a city or street address) can be displayed in Google Maps. In the Logs view these IP addresses are indicated with a country flag icon next to the IP address.

▼ To view the location of an IP address in the Logs view

1. Select **Monitoring→Logs**. The Logs view opens.
2. Select a log entry with an IP Address that has a flag icon associated with it.
3. Right-click the IP address in the Fields panel and select **Show in Google Maps**. The location is opened in Google Maps in your default browser.

Viewing IP Address Locations from the Whois Information Dialog

Prerequisites: Checking Whois Records for IP Addresses in Logs

Only IP addresses associated with a location (for example, a city or street address) can be displayed in Google Maps.

▼ To view the physical location of an IP address from a WHOIS record

↳ Right-click the Whois Information dialog and select **Show in Google Maps**. The location is opened in Google Maps in your default browser.

Monitoring Configurations and Policies

Prerequisites: None

The engines receive their configuration when a policy is installed. You can monitor the policies and configurations installed on the engines in the following ways:

- You can check which policy is currently being enforced and when it was last installed as explained in Checking the Currently Installed Policy (page 654).
- You can quickly view the most recent version of the installed policy as explained in Previewing the Currently Installed Policy (page 654).
- You can view the configurations that were transferred in each past policy installation and compare them to each other or the current policy stored on the Management Server, as explained in Checking and Comparing Policy Versions (page 654).
- You can quickly check if there are changes in the configurations on the Management Server that have not been transferred yet as explained in Checking for Untransferred Configuration Changes (page 656).

Related Tasks

- ▶ Creating and Managing Policy Elements (page 645).

Monitoring Administrator Actions

Prerequisites: None

A record of administrator actions is maintained in the system. The records can only be viewed by administrators who are allowed to view Audit logs. They can be browsed like any other logs in the Logs view.

▼ To view audit records

1. Select **File→System Tools→Audit** from the main menu. The Logs view opens with the Audit logs selected for viewing.
2. Browse and filter the logs as explained in section Browsing Logged Data (page 147).

Monitoring Task Execution

Prerequisites: Creating New Task Definitions, Scheduling Tasks, Starting Tasks Manually

You can check the status of running tasks and executed tasks (for example, upgrade tasks and system tasks) in the **Tasks→History** branch in the Administration Configuration view. Running tasks are shown automatically. Executed tasks are only shown when you enable the Show Executed Tasks option in the right-click menu. System tasks that run automatically are not visible in the History branch.

▼ To show Running and Executed Tasks

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Expand the **Tasks** branch.
3. Right-click **History** and select **Show Executed Tasks**.

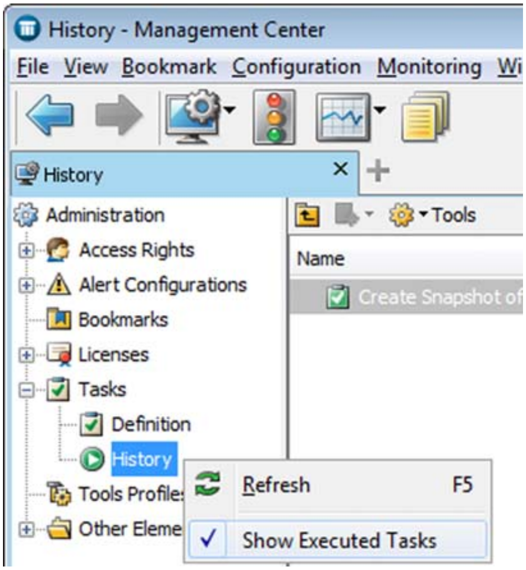

To check Task status

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Browse to **Tasks→History**. The details of the Tasks open in the right panel. The following columns may be of the most interest:
   - The **Progress** column shows the progress of a running Task.
   - The **Info** column shows additional details about the execution of a Task.
   - The **State** column shows the status of the Task. The table below explains the different Task states.

Table 7.9 Task State Information

State	Explanation
Started	The Task has started running.
Aborted	The Task was commanded to stop and stopped correctly.
Failed	An error caused the Task to stop.
Finished	The Task was completed successfully.
Partial	Part of the Task was completed successfully, but other parts had errors. This can occur, for example, when a Task has many sub-tasks and some of the sub-tasks fail.
Timeout	The Task stopped because it did not complete before the end of the time limit.

Related Tasks

- Stopping Task Execution (page 1041)

Taking a Traffic Capture

Prerequisites: None

You can capture network traffic data for network troubleshooting purposes. This data helps you to analyze network traffic to and from security engines. It is also often useful to have this data available when contacting Stonesoft Support.

Traffic capture creates a .zip file that contains a tcpdump CAP file, which is compatible with standard “sniffer” tools such as tcpdump, WinDump, or Wireshark. You can select whether to include full packet information or only IP headers in the tcpdump. You can also include a free-form description and information on your system configuration and trace files in the traffic capture .zip file.

The data can be archived and analyzed later, as the traffic capture .zip file is saved either on the Management Server or to a directory on your local workstation.

Traffic captures can only be taken on nodes that are online and have a policy uploaded.

Note – You must have permissions to send Advanced Commands to be able to take traffic captures. See Defining Administrator Roles (page 249).

You can stop or abort a traffic capture at any point once it has been launched.

- If you stop a traffic capture, all captured tcpdump data is compressed and sent to the Management Server or to your local workstation.
- If you abort a traffic capture, all captured tcpdump data is deleted.

▼ To take a traffic capture

1. Select **Monitoring→System Status**.
2. Right-click a Security Engine and select **Tools→Capture Traffic**. The Select Engine Interface dialog opens.

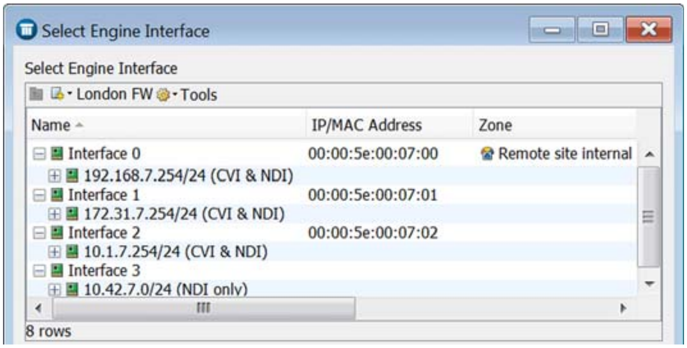

Name	IP/MAC Address	Zone
⊟ Interface 0	00:00:5e:00:07:00	Remote site internal
⊞ 192.168.7.254/24 (CVI & NDI)		
⊟ Interface 1	00:00:5e:00:07:01	
⊞ 172.31.7.254/24 (CVI & NDI)		
⊟ Interface 2	00:00:5e:00:07:02	
⊞ 10.1.7.254/24 (CVI & NDI)		
⊟ Interface 3		
⊞ 10.42.7.0/24 (NDI only)		

8 rows

3. Select the interface(s) whose traffic you want to capture and click **Select**. The Traffic Capture Task Properties dialog opens.

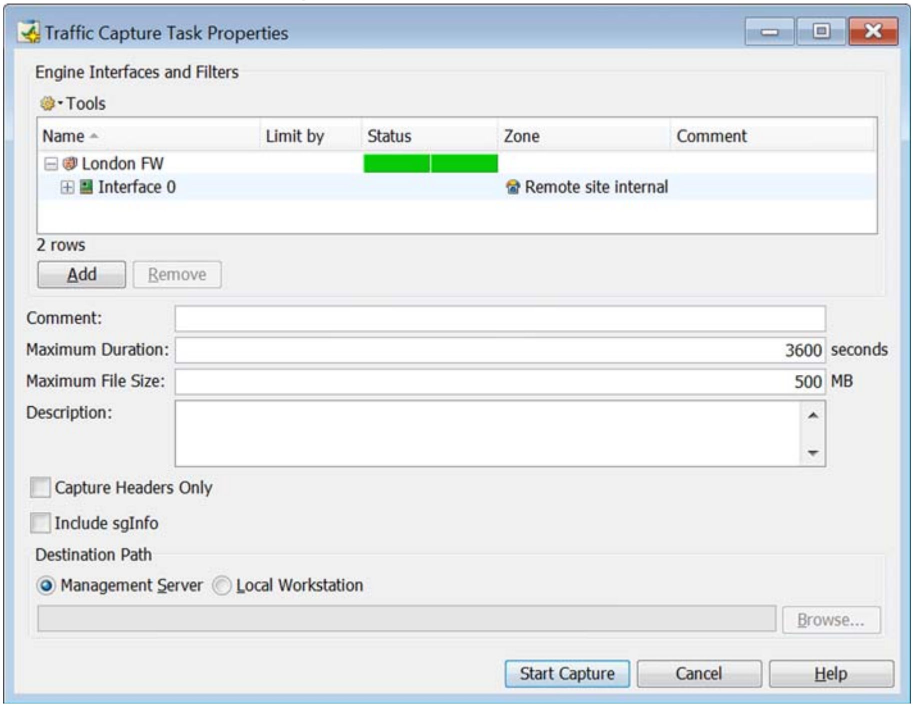

Name	Limit by	Status	Comment	Zone	Comment
⊟ London FW		🟩🟩			
⊞ Interface 0				Remote site internal	

2 rows

Comment:

Maximum Duration: 3600 seconds

Maximum File Size: 500 MB

Description:

□ Capture Headers Only

□ Include sgInfo

● Management Server ○ Local Workstation

4. (*Optional*) Click **Add** to add more interfaces to the traffic capture. You can also add interfaces from other types of security engines.
   - You can create tcpdump files for several different interfaces in the same Traffic Capture task. The Traffic Capture .zip file contains a separate CAP file for each interface included in the capture.

5. (*Optional*) Click the **Limit by** field and enter an IPv4 or IPv6 address to limit the scope of the traffic capture. The IP address must match either the source or destination of the packets included in the capture.
6. Define other traffic capture options as explained in the table below.

Table 7.10 Other Traffic Capture Options

Option		Explanation
Comment (*Optional*)		Add a comment to be appended to the .zip file name. The maximum length is 60 characters.
Maximum Duration		Define the maximum duration of the traffic capture. The duration is applied to all interfaces selected for the capture. The creation of the tcpdump file stops automatically once the maximum duration has been reached.
Maximum File Size		Define the maximum size of the tcpdump file. The creation of the tcpdump file stops automatically once the maximum file size has been reached.
Description (*Optional*)		Add a description of the traffic capture. This description is included as a separate file in the traffic capture .zip file.
Capture Headers Only (*Optional*)		Select this option to include only IP headers in the tcpdump file(s). Do not select this if you want to include full packets in the capture.
Include sgInfo (*Optional*)		Select this option to include system configuration files and system trace files in the traffic capture .zip file. It is important to include this information if you send the traffic capture to Stonesoft Support.
Destination Path	Management Server	Select this option to save the traffic capture .zip file in the `<installation directory>/data/TrafficCapture` directory on the Management Server. **Note!** If you installed the Management Server in the C:\Program Files\Stonesoft\Management Center directory in Windows, some program data may be stored in the C:\ProgramData\Stonesoft\Management Center\data directory.
	Local Workstation	Select this option and **Browse** to the location on your local workstation where you want to save the file.

7. Click **Start Capture**. The traffic capture is launched. The traffic capture .zip file is generated and saved in the selected location.

Related Tasks

- Troubleshooting Alert, Log, and Error Messages (page 1081)
- Troubleshooting Engine Operation (page 1101)

Checking Maintenance Contract Information

Prerequisites: None

You can view maintenance contract and support level information for your licenses in the Management Client by allowing your Management Server to contact Stonesoft servers. This information is available for each license in the **Licenses→All Licenses** branch of the Administration Configuration view if the Management Server is able to contact the Stonesoft servers.

To enable viewing maintenance contract and support level information permanently for your licenses, you need to allow the Management Server to connect to the Stonesoft servers. This option is found on the **Updates** tab in the Management Server Properties dialog. See Getting Started with Automatic Updates and Engine Upgrades (page 242).

Viewing Maintenance Contract Information

When contract checking is enabled, you can view information on your support contract in the License table in the Management Client.

▼ To view maintenance contract information

1. Select **Configuration→Configuration→Administration**.
2. Expand the **Licenses** tree.
3. Browse to a component-specific branch or to the **All Licenses** branch. The table of Licenses opens in the right panel.
4. Select the license whose information you want to view. If the information is not available in the Maintenance Contract Expires or Support Level columns, enable the checking of maintenance contract information in one of the following ways:
   - Check the information as explained in Fetching Maintenance Contract Information.
   - Enable the checking of maintenance contract information as explained in Getting Started with Automatic Updates and Engine Upgrades (page 242).

Related Tasks

- ▶ Getting Started with Automatic Updates and Engine Upgrades (page 242)

Fetching Maintenance Contract Information

If contacting Stonesoft servers is allowed, the SMC checks the maintenance contract information for licenses automatically. See Getting Started with Automatic Updates and Engine Upgrades (page 242). Otherwise, you can fetch the information manually.

▼ To manually fetch maintenance contract information for licenses

1. Select **Configuration→Configuration→Administration**.
2. Expand the **Licenses** branch.
3. Right-click one of the branches under **Licenses** and select **Check Maintenance Contract**. You are prompted to confirm that you want to send proof of license codes to Stonesoft.
4. Click **Yes**. If the Management Server is able to connect to Stonesoft servers, the Management Client displays the maintenance contract and support level information for the licenses. The information is available in the Management Client until the Management Server is restarted.

Checking When Internal Certificates or Internal CAs Expire

Prerequisites: None

You can check the status of internal certificates used in system communications and the status of the SMC's Internal Certificate Authority that automatically signs the internal certificates. The Internal Certificate Authority is valid for ten years. It is renewed automatically. The SMC does not accept certificates signed by an expired Internal Certificate Authority. All components must receive new certificates signed by the new Internal Certificate Authority before the old Internal Certificate Authority expires. By default, internal certificates are renewed automatically for engines and internal gateways. For all other components, you must generate the certificates manually. See Replacing Expired/Missing Certificates (page 1094).

▼ To check when internal certificates or Internal Certificate Authorities expire

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Browse to **Other Elements→Internal Certificates** or **Other Elements→Internal Certificate Authorities**. The existing internal certificates or Internal Certificate Authorities are displayed in the right panel.
   - See the **Expiration Date** column for information on the expiration of the certificate or the Internal Certificate Authority.
   - To view detailed information, right-click a certificate or Internal Certificate Authority and select **Properties**. The Properties dialog for the certificate or the Internal Certificate Authority opens and shows the Status of the Internal Certificate Authority:

Status	Explanation
Ready to Use	A new Internal Certificate Authority has been created and signs only Management Server certificates.
Active	All system components use certificates signed by this Internal Certificate Authority. When there are multiple Management Servers, all Management Servers are using the new Internal Certificate Authority.
Renewal Started	This is a new Internal Certificate Authority that the system has created automatically. The process of renewing internal certificates for the system components has begun.
Expires Soon	A new Internal Certificate Authority has been created but some components may still use certificates signed by this Internal Certificate Authority.
Inactive	This Internal Certificate Authority has expired or no system component uses a certificate signed by this Internal Certificate Authority.

▶ Troubleshooting Certificates (page 1091)
▶ Checking When Gateway Certificates Expire (page 981)
▶ Checking When an Internal VPN CA Expires (page 982)

CHAPTER 8

MONITORING THIRD-PARTY DEVICES

The Stonesoft Management Center can be configured to log and monitor other manufacturers' devices in much the same way as system components are monitored.

The following sections are included:

Getting Started with Third-Party Device Monitoring

Prerequisites: None

What the Third-Party Device Monitoring Feature Does

You can configure Log Servers for a full range of monitoring features for third-party devices:

- Log Servers can receive a syslog stream and store the information in Stonesoft log format. The stored logs can then be processed just like logs generated by Stonesoft components. For example, the data can be included in reports you generate.
- Log Servers can receive SNMP statistics information and NetFlow data from third-party devices. You can view this information as part of your Overviews or create reports based on the received data.
- Log Servers can probe devices through several alternative methods. You can monitor the device status in the Management Client in the same way as for the system components.

Limitations

Each Log Server can monitor a maximum of 200 third-party devices. This limit is enforced automatically.

Statistics for third-party devices are limited to amount of free and used memory, CPU load, and interface statistics.

Your Management Server license may limit the number of managed components. Each monitored third-party device is counted as a fifth of a managed unit.

Configuration Overview

1. (*Optional*) Define syslog reception settings for each particular type of third-party device. See Converting Logs From External Devices (page 133).
2. (*Optional*) Define the status and statistics monitoring settings for each type of third-party device. See Monitoring the Status of Third-Party Devices (page 141).
3. Activate monitoring by adding the monitoring settings in the properties of each third-party element (Router, Host, Active Directory Server, LDAP Server, RADIUS Authentication Server, or TACACS+ Authentication Server). See Activating Monitoring of a Third-Party Device (page 145).
4. Depending on features used, configure the third-party device for monitoring. See Configuring a Third-Party Device for Monitoring (page 145).

Related Tasks

- ▶ Changing the Ports for Third-Party Device Monitoring (page 146)
- ▶ Activating/Deactivating Third-Party Status Monitoring Alerts (page 146)

Converting Logs From External Devices

Prerequisites: None

You can set up most external devices to send logs to the Log Server in syslog format. The Log Server can convert incoming syslog entries to Stonesoft log entries. You can define a *Logging Profile* element to determine how the field values are selected from a syslog entry and how those field values are inserted into a Stonesoft log entry. A Logging Profile must have at least one *logging pattern*. Logging patterns determine how the fields from syslog entry are parsed to the appropriate log fields in a Stonesoft log entry.

There are two alternative modes that you can use to create logging patterns:

- *Ordered Fields* can be used when the fields in the syslog message are not separated by keywords and the type of field can only be deduced from its position. The received syslog entries are parsed in a sequence that you define in the Logging Profile. If the incoming logs vary in structure, you must define a different sequence for each type of structure. You can define several patterns in one Logging Profile.
- *Key-Value Pairs* can be used when the syslog message contains keywords that describe the type of field. The received syslog entries are parsed based on key values that you define in the Logging Profile. You can define the key values in any order. A single definition can be used even if logs vary in structure.

It is easier to configure a pattern using key-value pairs. We recommend that you use key-value pairs if a third-party device formats the relevant parts of the syslog packet as key-value pairs. Ordered fields can be used to process all syslog data regardless of its format, but it is more difficult to configure a pattern as ordered fields.

If a basic match is found, the system simply converts a syslog entry to a Stonesoft log field. This is appropriate for most data. You can define *Field Resolvers* for more complex operations. For more information on Field Resolvers, see Adding Field Resolvers (page 138).

What's Next?

▶ Creating a Logging Profile Element (page 134).

Related Tasks

▶ Adding Field Resolvers (page 138).

▶ Changing the Ports for Third-Party Device Monitoring (page 146)

▶ Some Logging Profile samples are available at http://stoneblog.stonesoft.com/stoneblog-community/. You can import these as elements (**File→Import→Import Elements**).

Creating a Logging Profile Element

A syslog packet consists of three parts: <PRI>, HEADER, and MSG. In a Logging Profile element you define patterns for converting the MSG part of the syslog packet to a Stonesoft log entry. A Logging Profile parses the data in a syslog message to the corresponding Stonesoft log fields when the syslog entry is converted to a Stonesoft log entry. The parts of the syslog packet are explained in more detail in the table below.

Table 8.1 Parts of the Syslog Packet

Section	Description
<PRI>	Contains facility and priority information. The Log Server automatically extracts the Facility value from the <PRI> part and converts it to the Syslog Facility field in Stonesoft logs. You do not define patterns for mapping this section in the Logging Profile view.
HEADER	Contains a timestamp and a hostname or IP address of a device. The Log Server automatically extracts the data in the HEADER part. This section is optional in syslog packets, so not all devices send this data.
MSG	Contains the text of the syslog message. In the Logging Profile view, you define the mapping for parsing this part of the syslog packet.

▼ To create a Logging Profile element

1. Select **Configuration→Configuration→Monitoring**.
2. Right-click **Third Party Devices** and select **New→Logging Profile**. The Logging Profile Properties dialog opens.
3. Enter a **Name** for the Logging Profile.
4. Click **OK**. The new Logging Profile opens in the Logging Profile view.

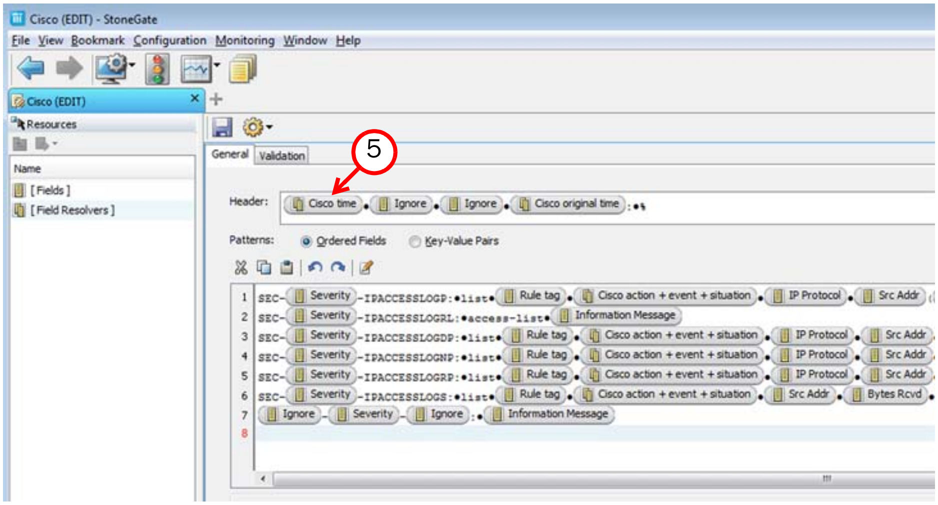

Header: Cisco time • Ignore • Ignore • Cisco original time :•%

Patterns: ● Ordered Fields ○ Key-Value Pairs

```
SEC-Severity-IPACCESSLOGP:•list•Rule tag•Cisco action + event + situation•IP Protocol•Src Addr
SEC-Severity-IPACCESSLOGRL:•access-list•Information Message
SEC-Severity-IPACCESSLOGDP:•list•Rule tag•Cisco action + event + situation•IP Protocol•Src Addr
SEC-Severity-IPACCESSLOGNP:•list•Rule tag•Cisco action + event + situation•IP Protocol•Src Addr
SEC-Severity-IPACCESSLOGRP:•list•Rule tag•Cisco action + event + situation•IP Protocol•Src Addr
SEC-Severity-IPACCESSLOGS:•list•Rule tag•Cisco action + event + situation•Src Addr•Bytes Rcvd•
Ignore-Severity-Ignore:•Information Message
```

5. (*Optional*) Drag-and-drop items to the **Header** field from the **Fields** branch in the left panel or use the type-ahead search to insert fields.

Tip – Press Ctrl-space to activate a type-ahead search at the current cursor position.

- You can add fields that are the same for all the logging patterns that you define in the **Patterns** panel.
- To omit a portion of data add an **Ignore** field.

Caution – You must type or copy-paste from the syslog message any number of tokens that appear before and after the field values. If you do not insert the appropriate tokens the data cannot be parsed.

Example In the illustration above, the header of the syslog entry contains the following data common for all patterns <Cisco time><space><Ignore><space><Ignore><space><Cisco original time>, so the header contains the following data <Sep 21 04:04:56> <cisco-example.stonesoft.com> <1815452:> <Sep 21 04:04:55> %. The **Ignore** field is used for <cisco-example.stonesoft.com> and <1815452:>, so the values are not converted to Stonesoft log entry format.

6. Select the mode for Patterns based on the type of data this Logging Profile handles:
   - **Ordered Fields**: the syslog entries are parsed in the specified order. If the incoming logs vary in structure, you must define several patterns.
   - **Key-Value Pairs**: the syslog entries are parsed on the basis of key-value pairs that you define. You can add key-value pairs in any order. You can use one pattern for all logs even if the logs vary in structure.

What's Next?

▶ Defining Ordered Field Logging Patterns.

▶ Defining Key-Value Pair Logging Patterns (page 137).

Defining Ordered Field Logging Patterns

Each received syslog entry is converted to Stonesoft log entry. The field values that match a specified pattern are copied as-is to a Stonesoft log field. Additionally, you can create Field Resolvers to convert specific values in the syslog data to specific values in Stonesoft logs.

A pattern that you define in a Logging Profile must be an exact match of the incoming syslog entry. If incoming logs vary in structure, you must define a different pattern for each type of entry. If several patterns match, the system uses the pattern with the most matching entries.

If an arriving entry does not match any of the defined patterns, you can add the data into a single Stonesoft log field or discard it.

▼ To define logging patterns as ordered fields

1. Make sure **Ordered Fields** is selected as the mode for parsing syslog data.

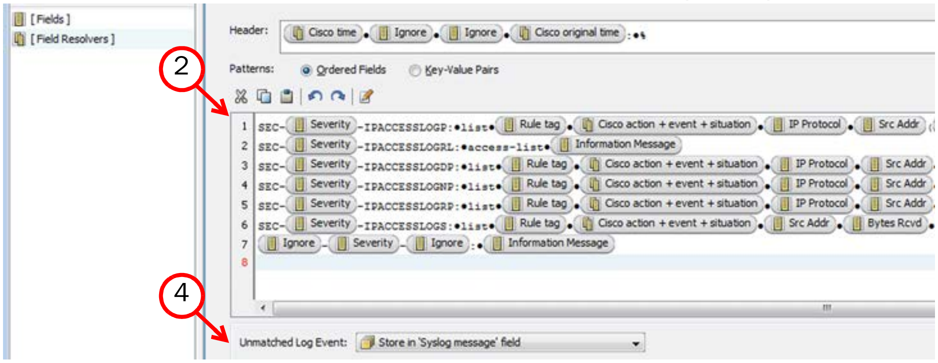

Header: Cisco time • Ignore • Ignore • Cisco original time :•

Patterns: ● Ordered Fields ○ Key-Value Pairs

```
SEC- Severity -IPACCESSLOGP:•list• Rule tag • Cisco action + event + situation • IP Protocol • Src Addr
SEC- Severity -IPACCESSLOGRL:•access-list• Information Message
SEC- Severity -IPACCESSLOGDP:•list• Rule tag • Cisco action + event + situation • IP Protocol • Src Addr
SEC- Severity -IPACCESSLOGNP:•list• Rule tag • Cisco action + event + situation • IP Protocol • Src Addr
SEC- Severity -IPACCESSLOGRP:•list• Rule tag • Cisco action + event + situation • IP Protocol • Src Addr
SEC- Severity -IPACCESSLOGS:•list• Rule tag • Cisco action + event + situation • Src Addr • Bytes Rcvd •
Ignore - Severity - Ignore :• Information Message
```

Unmatched Log Event: Store in 'Syslog message' field

2. Drag-and-drop items to the **Patterns** field from the **Fields** branch in the left panel or use the type-ahead search to insert the field values.

Tip – Press Ctrl-space to activate a type-ahead search at the current cursor position.

- Alternatively, you can define a Field Resolver and add it to the pattern instead of a log field. For more information, see Adding Field Resolvers (page 138).
- To omit a portion of data add an **Ignore** field.
- The pattern you define is used for parsing syslog data. The pattern must be an exact match of the data sequence in the syslog message.
- Right-click the empty space in the Patterns field and select **Toggle Comment** to comment out a pattern. This is a useful tool for testing different patterns.

Caution – You must type or copy-paste from the syslog message any number of tokens that appear before and after the field values. If you do not insert the appropriate tokens the data is not parsed.

3. (*Optional*) If some incoming log entries have a different structure, press `Enter` and define a new pattern.

4. In the **Unmatched Log Event** section, select the action for handling syslog data that does not match any defined logging patterns:
   - **Store in 'Syslog message' field**: a log entry is created and all data is inserted in the Syslog Message log field. The log entry is stored on the Log Server.
   - **Display in 'Syslog message' field (Current mode only)**: a log entry is created and all data is inserted in the Syslog Message log field. The log entry is displayed in the Current Events mode in the Logs view but it is not stored.
   - **Ignore**: the data is discarded.

What's Next?

▶ Validating a Logging Profile (page 139).

▶ Activating Monitoring of a Third-Party Device (page 145).

Defining Key-Value Pair Logging Patterns

You can define key-value pairs for converting syslog data. In this case, the Log Server parses each received syslog entry data based on key-value pairs you define. The data in the incoming syslog message must be formatted as key-value pairs. If you want to parse data that is not formatted as key-value pairs, see Defining Ordered Field Logging Patterns (page 135).

If an arriving entry does not match any of the defined patterns, you can add the data into a single Stonesoft log field or discard it.

▼ To define logging patterns as key-value pairs

1. Make sure **Key-Value Pairs** is selected as the mode for matching log patterns.

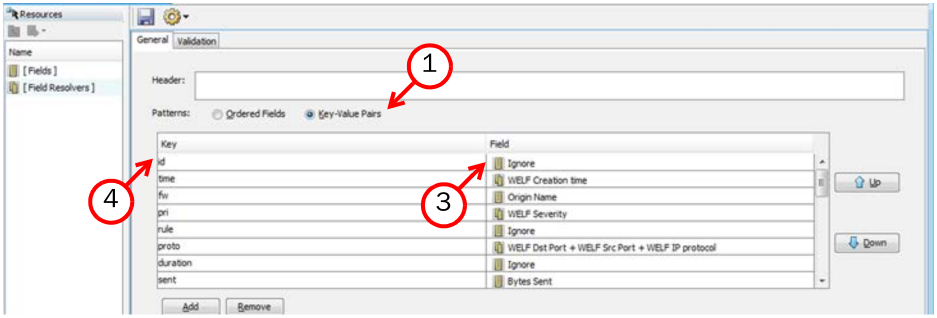

Header:

Patterns: ○ Ordered Fields ● Key-Value Pairs

Key	Field
id	Ignore
time	WELF Creation time
fw	Origin Name
pri	WELF Severity
rule	Ignore
proto	WELF Dst Port + WELF Src Port + WELF IP protocol
duration	Ignore
sent	Bytes Sent

2. Click **Add** to add a new row to the Patterns table. By default, the **Ignore** field is added in the **Field** column.

3. Drag-and-drop Stonesoft log fields from the **Fields** branch to the **Field** column.
   - Alternatively, you can define a Field Resolver and add it to the pattern instead of a log field. For more information, see Adding Field Resolvers (page 138).
   - To omit a portion of data add an **Ignore** field.

4. Double-click the **Key** column for the log field that you added and type the corresponding key value as it appears in the syslog message (for example `time`).

5. *(Optional)* Repeat steps 2 to 4 to add more key-value pairs.
   - The key values can be added in any order.
   - The key values are converted to Stonesoft log entries based on the key values alone.

6. In the **Unmatched Log Event** section, select the action for handling syslog data that does not match any defined logging patterns:
   - **Store in 'Syslog message' field**: a log entry is created and all data is inserted in the Syslog Message log field. The created log entry is stored on the Log Server.
   - **Ignore**: the data is discarded.

What's Next?

- ▶ Validating a Logging Profile (page 139).
- ▶ Activating Monitoring of a Third-Party Device (page 145).

Adding Field Resolvers

Field Resolvers convert values in incoming syslog fields to different values in Stonesoft logs. You can use field resolvers in the following cases:

- To convert timestamps: Different external devices use different date and time formats. A Field Resolver for each different incoming format maps the times and dates correctly to the Stonesoft log format. The date and time syntax in Field Resolvers follows the Java standard.
- To convert product-specific values: Some concepts in the log data have a pre-set range of values on the external devices and in the SMC, but the possible values are different. A Field Resolver can convert this data. For example, you can map a range of alert severities in the original data to similar alert severities in Stonesoft logs (Info/Low/High/Critical).
- To convert one value to several log fields: In some cases, a single value may have several corresponding log fields in Stonesoft logs. A Field Resolver can parse a single value into multiple Stonesoft log fields. For example, Stonesoft components set an Action, a Situation, and an Event for traffic filtering decisions. If the external component notifies a “permitted” action, the Field Resolver can set the corresponding Stonesoft log values for all three log fields. See Defining a Field Resolver for Multiple Values.

▼ To add a Field Resolver

1. Click the **Field Resolvers** branch in the left panel of the Logging Profile view to display the available elements.
2. Add a Field Resolver:
   - Select a Field Resolver and drag-and-drop it to the Header field or to the Patterns panel.
   - If the Field Resolver that you want to select is not listed, create a new Field Resolver element through the New icon at the top of the left panel.

What’s Next?

- ▶ Defining a Field Resolver for Multiple Values
- ▶ Defining a Field Resolver for Date and Time (page 139)

Defining a Field Resolver for Multiple Values

▼ To define a multi-valued field resolver

1. Enter a **Name** for the new Field Resolver.
2. Select **Multi-valued** as **Field Type**.

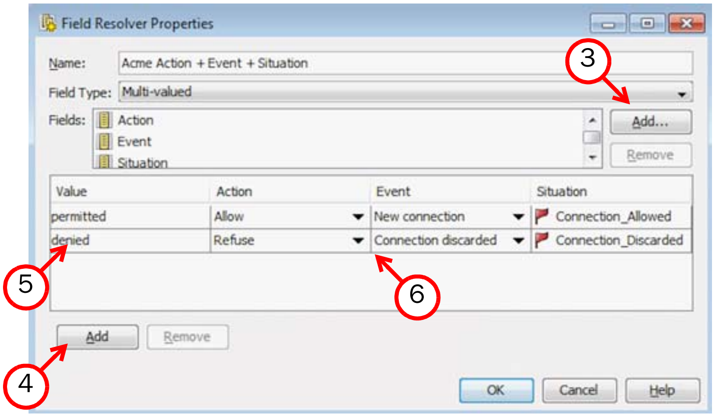

Name: Acme Action + Event + Situation

Field Type: Multi-valued

Fields: Action, Event, Situation

Value	Action	Event	Situation
permitted	Allow	New connection	Connection_Allowed
denied	Refuse	Connection discarded	Connection_Discarded

3. In **Fields**, click the **Add** button and select the appropriate Stonesoft log field(s). The incoming syslog data is inserted into the log fields you select.
4. Click **Add** under the **Value** to add a new row to the Value panel.
5. Enter the value that is used on the third-party device in the **Value** cell.
6. Enter or select the value you want to use for each selected Stonesoft log field depending on the type of log field.
7. *(Optional)* Repeat steps 1 to 4 to add more rows.
8. Click **OK**.
9. Drag-and-drop the Field Resolver to the **Header** or **Pattern** field of your custom Logging Profile.

Defining a Field Resolver for Date and Time

▼ **To define a date field resolver**

1. Click **Select** next to **Time Field** and select the log field for which you want to define a timestamp.
2. In the **Time Format** field, enter the format that the third-party device uses for the timestamp.
   - Type the format according to Java standards. For syntax, see http://java.sun.com/j2se/1.5.0/docs/api/java/text/SimpleDateFormat.html.

Example Enter *MMM dd HH:mm:ss* to map the log timestamp as *Jan 30 13:23:12*.

3. Click **OK**.
4. Drag-and-drop the Field Resolver to the **Header** or **Pattern** field of your custom Logging Profile.

Validating a Logging Profile

You can test a Logging Profile that you created to verify that the syslog data is converted correctly to Stonesoft log fields.

Note – If the Logging Profile is not already open, you have to first open it in the Logging Profile editor and then click Edit.

▼ **To validate a Logging Profile**

1. Browse to **Configuration→Configuration→Monitoring→Third Party Devices→Logging Profiles**.
2. Right-click the Logging Profile that you want to validate and select **Edit Logging Profile**.
3. Switch to the **Validation** tab in the Logging Profile view.

4. Click **Browse** to select a file with syslog data.

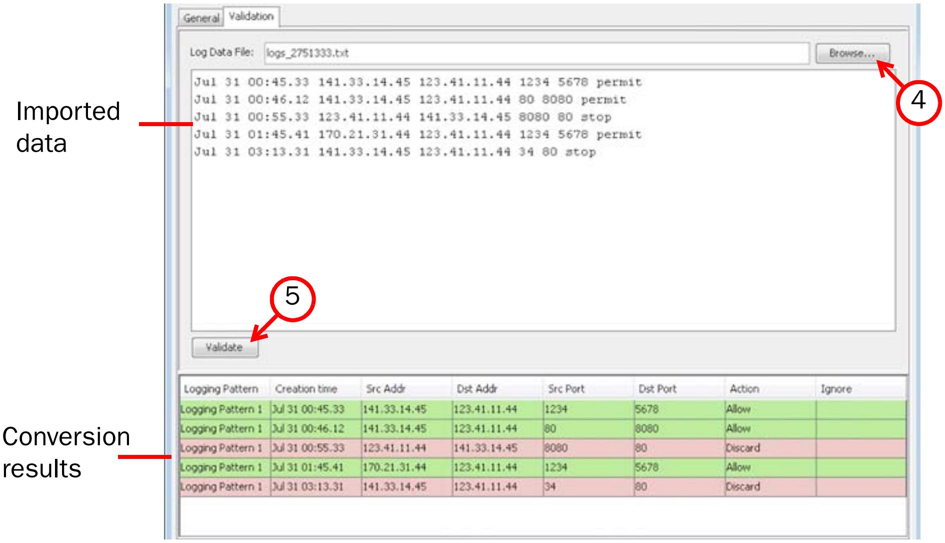

Log Data File: logs_2751333.txt

```
Jul 31 00:45.33 141.33.14.45 123.41.11.44 1234 5678 permit
Jul 31 00:46.12 141.33.14.45 123.41.11.44 80 8080 permit
Jul 31 00:55.33 123.41.11.44 141.33.14.45 8080 80 stop
Jul 31 01:45.41 170.21.31.44 123.41.11.44 1234 5678 permit
Jul 31 03:13.31 141.33.14.45 123.41.11.44 34 80 stop
```

Logging Pattern	Creation time	Src Addr	Dst Addr	Src Port	Dst Port	Action	Ignore
Logging Pattern 1	Jul 31 00:45.33	141.33.14.45	123.41.11.44	1234	5678	Allow	
Logging Pattern 1	Jul 31 00:46.12	141.33.14.45	123.41.11.44	80	8080	Allow	
Logging Pattern 1	Jul 31 00:55.33	123.41.11.44	141.33.14.45	8080	80	Discard	
Logging Pattern 1	Jul 31 01:45.41	170.21.31.44	123.41.11.44	1234	5678	Allow	
Logging Pattern 1	Jul 31 03:13.31	141.33.14.45	123.41.11.44	34	80	Discard	

Imported data

Conversion results

5. Click **Validate**. The imported data is displayed in the first panel. The validation results are displayed in the second panel. The first column of the results panel shows which logging pattern is used to convert each syslog entry.

6. Click **Save** to save the Logging Profile.

What's Next?

- If you want to view statistics and status information on a third party device in the Management Client, see Creating a Probing Profile (page 143).
- To activate the Logging Profile, see Activating Monitoring of a Third-Party Device (page 145).

Monitoring the Status of Third-Party Devices

Prerequisites: None

The Log Server can actively probe the status of third-party components using one of several alternative methods:

- SNMPv1.
- SNMPv2c.
- SNMPv3.
- Pings.
- TCP connection probes.

When one of the SNMP status probing methods is used, you can also set up statistics reception for the device. Statistics reception relies on SNMPv1 traps sent by the third-party device.

The SMC supports statistical monitoring of the following details:

- Amount of free and used memory.
- CPU load.
- Interface statistics.

What's Next?

- ▶ If you want to receive statistics from a new type of device, proceed to Importing MIBs (page 142).
- ▶ To set up status monitoring without statistics for a new type of device, proceed to Creating a Probing Profile (page 143).
- ▶ To activate status and statistics monitoring for a previously configured type of device, proceed to Activating Monitoring of a Third-Party Device (page 145).

Related Tasks

- ▶ Changing the Ports for Third-Party Device Monitoring (page 146).

Importing MIBs

You can import third-party MIBs (management information bases) to support third-party SNMP monitoring. The OIDs (object identifiers) allows resolving the SNMP traps when they appear in log entries. If the OIDs are not resolved, they appear in the logs using a more difficult to read dotted notation. Only SNMPv1 Trap Reception is supported.

▼ To import a MIB

1. Select **Configuration→Configuration→Monitoring**. The Monitoring Configuration view opens.
2. Expand the **Third Party Devices** branch.
3. Right-click **MIBs** and select **Import MIB**. The Import MIB Files window opens.
4. Browse for the MIB file to import and click **Import**. A new view opens showing the progress of the import.
5. (*Optional*) Click **Close** when the import is finished.
6. Browse to **MIBs→All MIBs** or **MIBs→By Vendor** in the Monitoring view to see the MIB you imported.
7. To view the contents of the MIB, right-click it and select **Properties**. The MIB Browser opens.
   - The General tab shows the content of a MIB file “as is”.
   - Switch to the **OID Tree** tab to view the objects that the MIB file defines: the object identifiers, their OIDs in dot notation, and possibly a description of the object.

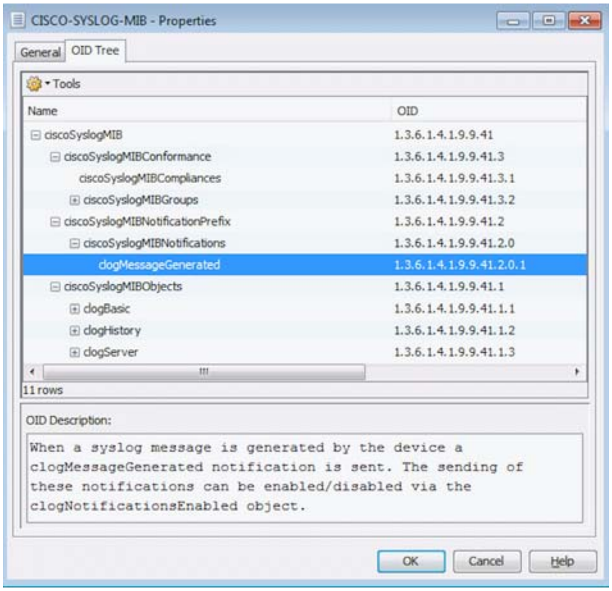

Name	OID
⊟ ciscoSyslogMIB	1.3.6.1.4.1.9.9.41
⊟ ciscoSyslogMIBConformance	1.3.6.1.4.1.9.9.41.3
ciscoSyslogMIBCompliances	1.3.6.1.4.1.9.9.41.3.1
⊞ ciscoSyslogMIBGroups	1.3.6.1.4.1.9.9.41.3.2
⊟ ciscoSyslogMIBNotificationPrefix	1.3.6.1.4.1.9.9.41.2
⊟ ciscoSyslogMIBNotifications	1.3.6.1.4.1.9.9.41.2.0
clogMessageGenerated	1.3.6.1.4.1.9.9.41.2.0.1
⊟ ciscoSyslogMIBObjects	1.3.6.1.4.1.9.9.41.1
⊞ clogBasic	1.3.6.1.4.1.9.9.41.1.1
⊞ clogHistory	1.3.6.1.4.1.9.9.41.1.2
⊞ clogServer	1.3.6.1.4.1.9.9.41.1.3

11 rows

OID Description:

```
When a syslog message is generated by the device a
clogMessageGenerated notification is sent. The sending of
these notifications can be enabled/disabled via the
clogNotificationsEnabled object.
```

8. If the MIB is correct, click **OK**.

Creating a Probing Profile

Probing Profiles define the monitoring of third-party device status (using SNMPv1/SNMPv2c/SNMPv3/Ping/TCP) and the settings for receiving statistics information from them (using SNMPv1).

▼ To create a Probing Profile

1. Select **Configuration→Configuration→Monitoring**. The Monitoring Configuration view opens.
2. Right-click **Third Party Devices** and select **New→Probing Profile**. The Probing Profile Properties dialog opens.

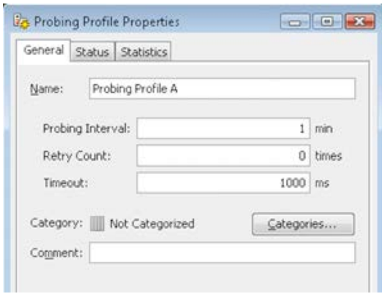

Name: Probing Profile A

Probing Interval: 1 min

Retry Count: 0 times

Timeout: 1000 ms

Category: Not Categorized

Comment:

3. Define the **General** probing profile settings:

Table 8.2 General Probing Profile Settings

Setting	Explanation
Name	The name of the probing profile.
Probing Interval	Defines how often the probing is done.
Retry Count	Defines how many times probing is attempted if the connection is not established.
Timeout	Defines how long the third-party device is probed.

4. Define the probing profile status settings on the **Status** tab

Setting	Explanation
Probing Method	Defines the probing method to be used.
Port	Defines the port to be probed (only visible when SNMP (default port 161) or TCP (default port 80) has been defined as the probing method).
SNMP Community	Defines the SNMP community (only with SNMP and SNMPv2c probing methods.

5. Additionally, if SNMPv3 is selected as the probing method, the following additional SNMPv3-specific security settings must be made:

Setting	Explanation
Security Name	Defines the default security name used for SNMPv3 requests.
Context Name	Defines the default context name used for SNMPv3 requests.
Security Level	Defines the authentication protocol and privacy protocol to be used. NoAuthNoPriv: No authentication protocol and no privacy protocol are used. AuthNoPriv: Authentication protocol is used but no privacy protocol is used. AuthPriv: Both authentication protocol and privacy protocol are used.
Authentication Protocol	Defines whether the authentication protocol used is MD5 or SHA. Enter the password to be used to the Password field.
Privacy Protocol	Defines the privacy protocol to be used (DES, AES, AES192, AES256, 3DES). Enter the password to be used in the Password field.

6. (*Optional, SNMP/SNMPv2c/SNMPv3 probing only*) Switch to the **Statistics** tab and define the settings for statistics reception using SNMPv1 traps:

Table 8.3 Statistics Options

Setting	Explanation
Memory Used	Activates statistics for the amount of used memory. Enter the OID of the third-party device or click **Select** to pick the OID from a list of imported MIBs.
Memory Free	Activates statistics for the amount of free memory. Enter the OID of the third-party device or click **Select** to pick the OID from a list of imported MIBs.
CPU Usage	Activates statistics for the processor load. Enter the OID of the third-party device or click **Select** to pick the OID from a list of imported MIBs.
Interface Statistics	Activates interface statistics reception.
IP Address	Enter the IP address of the third-party device. Click **Test** to launch an SNMP query and display the result in the panel below. Note that the actual statistics reception requires that the third-party device actively sends SNMP traps; the Log Server does not automatically query the device for statistics.

7. Click **OK** to save the Probing Profile.

Activating Monitoring of a Third-Party Device

Prerequisites: Converting Logs From External Devices, Creating a Probing Profile

▼ To activate status monitoring and/or log reception for a third-party device

1. Open the Properties dialog for the element that represents the third-party device (Router, Host, Active Directory Server, LDAP Server, RADIUS Authentication Server, or TACACS+ Authentication Server element).
   - For more information on creating these elements, see Getting Started with Defining IP Addresses (page 720).
2. Switch to the **Monitoring** tab.
3. Select the **Log Server** to which the logs from the third-party device are sent.
4. *(Optional)* To receive status information, select the **Status Monitoring** option and select the **Probing Profile**. See Creating a Probing Profile (page 143).
5. *(Optional)* To receive logging information, select the **Log Forwarding** option and select the **Logging Profile**. See Converting Logs From External Devices (page 133).
6. *(Optional)* Select **SNMP Trap Forwarding** if you want to receive SNMP traps from the third-party device.
7. *(Optional)* Select **NetFlow Forwarding** if you want to receive NetFlow data from the third-party device.
8. Click **OK**.

Related Tasks

▶ Monitoring the System (page 93)

▶ Browsing Logged Data (page 147)

Configuring a Third-Party Device for Monitoring

Prerequisites: Activating Monitoring of a Third-Party Device

For any type of monitoring, make sure that the connections between the third-party component and the Log Server are allowed through any possible traffic filtering in your network.

Ping and TCP status monitoring do not usually require additional configuration on the target device.

SNMP-based polling usually requires that the target device is specifically configured to respond to the Log Server's SNMP queries.

Statistics sending (as SNMPv1 traps) must always be specifically configured on the third-party device. For instructions on these tasks, consult the documentation of the third-party device.

If necessary, you can change the ports that the Log Server uses to listen for syslog and SNMP transmissions as explained in Changing the Ports for Third-Party Device Monitoring (page 146).

Changing the Ports for Third-Party Device Monitoring

Prerequisites: None

The default listening ports are:

- On Windows, the Log Server listens to syslog on port 514 and SNMP traps on port 162
- On Linux, the Log Server listens to syslog on port 5514 and SNMP traps on port 5162.

If necessary, you can change the ports, but the port number on Linux must always be higher than 1024 as explained in Changing Log Server Configuration Parameters (page 320).

If it is not possible to reconfigure the third-party device to send data to the correct port, you can either redirect traffic to a different port using an intermediate network device or on the Log Server, for example, using `iptables` in Linux:

```
iptables -t nat -A PREROUTING -p udp -m udp --dport 514 -j REDIRECT --to-ports
5514.
```

Activating/Deactivating Third-Party Status Monitoring Alerts

Prerequisites: Activating Monitoring of a Third-Party Device

You can optionally activate the status surveillance feature, which generates an alert if a monitored component's status remains unknown for 15 minutes.

▼ To enable/disable third-party device status surveillance

1. Select **Monitoring→System Status**.
2. Expand the **Third Party Devices** branch.
3. Right-click an element and select/deselect **Options→Status Surveillance**.

CHAPTER 9

BROWSING LOGGED DATA

You can view log, alert, and audit entries through the same unified tool. You can view data from Management Center servers, all types of engines, and from third-party components that are configured to send data to the Stonesoft Management Center.

The following sections are included:

Getting Started with the Logs View

Prerequisites: None

The Logs view displays all log, alert, and audit entries for the entire system. You can view any types of entries from any number of components together or individually.

Overview

For an overview, start by Opening the Logs View as explained below and proceed in order through introductions to the different arrangements. These include:

1. Default (Records) Arrangement, Panels, and Tools (page 148)
2. Details Arrangement (page 151)
3. Statistics Arrangement (page 152)
4. Log Analysis Arrangement (page 154)

Opening the Logs View

There are several ways to access the Logs view, for example:

- Select **Monitoring→Logs** or click the corresponding toolbar icon.
- Right-click an element that produces logs and select a log-related item in the **Monitoring** submenu (to view logs sent by that component).
- Click the alert indicator in the bottom right corner (to view active alerts).
- Create different bookmarks to open the Logs view with different filtering criteria.

Default (Records) Arrangement, Panels, and Tools

The default **Records** arrangement is optimized for efficient browsing of many entries.

Illustration 9.1 Logs View in the Records Arrangement

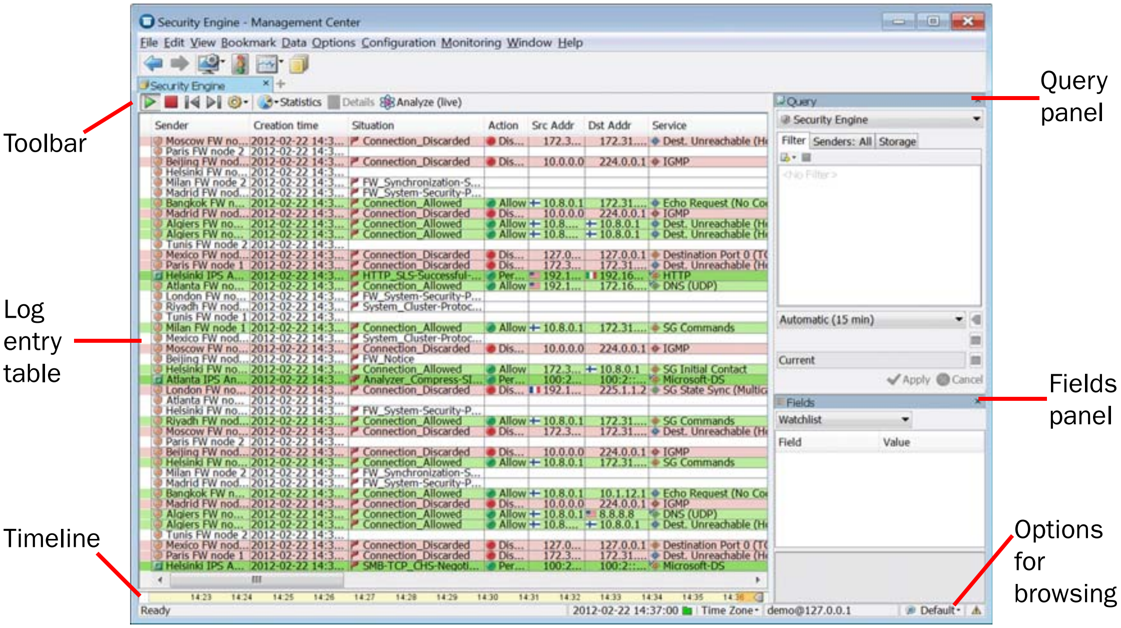

Sender	Creation time	Situation	Action	Src Addr	Dst Addr	Service
Moscow FW no...	2012-02-22 14:3...	Connection_Discarded	Dis...	172.3...	172.31....	Dest. Unreachable (H
Paris FW node 2	2012-02-22 14:3...					
Beijing FW nod...	2012-02-22 14:3...	Connection_Discarded	Dis...	10.0.0.0	224.0.0.1	IGMP
Helsinki FW no...	2012-02-22 14:3...					
Milan FW node 2	2012-02-22 14:3...	FW_Synchronization-S...				
Madrid FW nod...	2012-02-22 14:3...	FW_System-Security-P...				
Bangkok FW n...	2012-02-22 14:3...	Connection_Allowed	Allow	10.8.0.1	172.31....	Echo Request (No Co
Madrid FW nod...	2012-02-22 14:3...	Connection_Discarded	Dis...	10.0.0.0	224.0.0.1	IGMP
Algiers FW no...	2012-02-22 14:3...	Connection_Allowed	Allow	10.8....	10.8.0.1	Dest. Unreachable (H
Algiers FW no...	2012-02-22 14:3...	Connection_Allowed	Allow	10.8....	10.8.0.1	Dest. Unreachable (H
Tunis FW node 2	2012-02-22 14:3...					
Mexico FW nod...	2012-02-22 14:3...	Connection_Discarded	Dis...	127.0...	127.0.0.1	Destination Port 0 (T
Paris FW node 1	2012-02-22 14:3...	Connection_Discarded	Dis...	172.3...	172.31....	Dest. Unreachable (H
Helsinki IPS A...	2012-02-22 14:3...	HTTP_SLS-Successful-...	Per...	192.1...	192.16...	HTTP
Atlanta FW no...	2012-02-22 14:3...	Connection_Allowed	Allow	192.1...	172.16....	DNS (UDP)
London FW no...	2012-02-22 14:3...	FW_System-Security-P...				
Riyadh FW nod...	2012-02-22 14:3...	System_Cluster-Protoc...				
Tunis FW node 1	2012-02-22 14:3...					
Milan FW node 1	2012-02-22 14:3...	Connection_Allowed	Allow	10.8.0.1	172.31....	SG Commands
Mexico FW nod...	2012-02-22 14:3...	System_Cluster-Protoc...				
Moscow FW no...	2012-02-22 14:3...	Connection_Discarded	Dis...	10.0.0.0	224.0.0.1	IGMP
Beijing FW nod...	2012-02-22 14:3...	FW_Notice				
Helsinki FW no...	2012-02-22 14:3...	Connection_Allowed	Allow	172.3...	10.8.0.1	SG Initial Contact
Atlanta IPS An...	2012-02-22 14:3...	Analyzer_Compress-SI...	Per...	100:2...	100:2::...	Microsoft-DS
London FW no...	2012-02-22 14:3...	Connection_Discarded	Dis...	192.1...	225.1.1.2	SG State Sync (Multic
Atlanta FW no...	2012-02-22 14:3...					
Helsinki FW no...	2012-02-22 14:3...	FW_System-Security-P...				
Riyadh FW nod...	2012-02-22 14:3...	Connection_Allowed	Allow	10.8.0.1	172.31....	SG Commands
Moscow FW no...	2012-02-22 14:3...	Connection_Discarded	Dis...	172.3...	172.31....	Dest. Unreachable (H
Paris FW node 2	2012-02-22 14:3...					
Beijing FW nod...	2012-02-22 14:3...	Connection_Discarded	Dis...	10.0.0.0	224.0.0.1	IGMP
Helsinki FW no...	2012-02-22 14:3...	Connection_Allowed	Allow	10.8.0.1	172.31....	SG Commands
Milan FW node 2	2012-02-22 14:3...	FW_Synchronization-S...				
Madrid FW nod...	2012-02-22 14:3...	FW_System-Security-P...				
Bangkok FW n...	2012-02-22 14:3...	Connection_Allowed	Allow	10.8.0.1	10.1.12.1	Echo Request (No Co
Madrid FW nod...	2012-02-22 14:3...	Connection_Discarded	Dis...	10.0.0.0	224.0.0.1	IGMP
Algiers FW no...	2012-02-22 14:3...	Connection_Allowed	Allow	10.8.0.1	8.8.8.8	DNS (UDP)
Algiers FW no...	2012-02-22 14:3...	Connection_Allowed	Allow	10.8....	10.8.0.1	Dest. Unreachable (H
Tunis FW node 2	2012-02-22 14:3...					
Mexico FW nod...	2012-02-22 14:3...	Connection_Discarded	Dis...	127.0...	127.0.0.1	Destination Port 0 (T
Paris FW node 1	2012-02-22 14:3...	Connection_Discarded	Dis...	172.3...	172.31....	Dest. Unreachable (H
Helsinki IPS A...	2012-02-22 14:3...	SMB-TCP_CHS-Negoti...	Per...	100:2...	100:2::...	Microsoft-DS

Illustration 9.2 Toolbar in the Records Arrangement

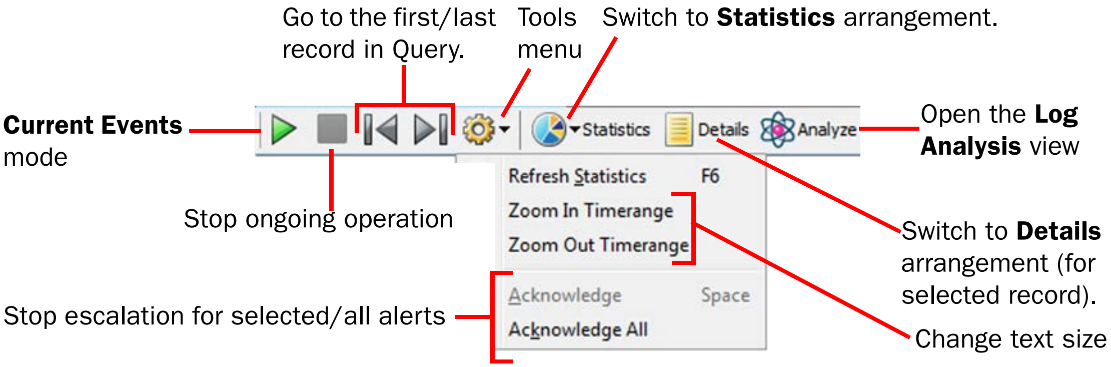

Go to the first/last record in Query.

Tools menu

Switch to **Statistics** arrangement.

Current Events mode

Open the **Log Analysis** view

Stop ongoing operation

Switch to **Details** arrangement (for selected record).

Stop escalation for selected/all alerts

Change text size

Illustration 9.3 Status Bar Options for Log Browsing

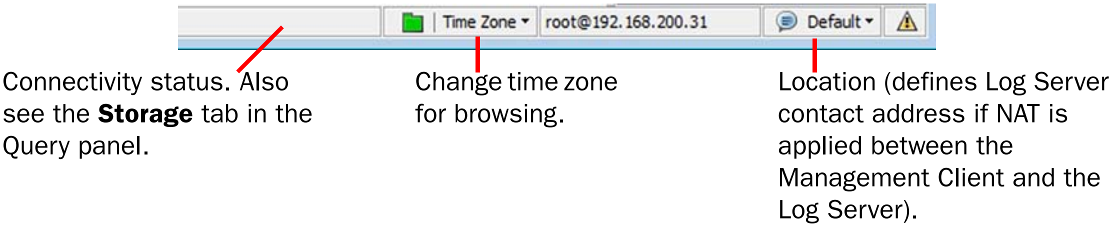

Connectivity status. Also see the **Storage** tab in the Query panel.

Change time zone for browsing.

Location (defines Log Server contact address if NAT is applied between the Management Client and the Log Server).

Note – You must select the correct Location for your Management Client to be able to see the logs if NAT is applied between your Management Client and a Log Server. See Configuring System Communications (page 63) for more information.

Logs View Panels

You can select and deselect panels through **View→Panels** in the main menu.

The following panels are available in most arrangements:

- **Fields** panel: Provides quick access to categorized log entry details.
- **Query** panel: The most important tool in the Logs view. The Query panel allows you to filter the log entries so that you can find the information you need.
- **Alert Events** panel: Details on alert reception and escalation.
- **Task Status** panel: Only available in the Records view. Displays the status of log-related tasks, such as a log export that you launch from the Logs view.
- **Hex** panel: Displays traffic recordings generated by the Excerpt logging option of an Inspection rule (other recordings are viewed using an external viewer).
- **Summary** panel: Textual explanation of the event that generated the record.
- **Event Visualization** panel: A graphic showing the main information about the event.
- **Info** panel: Displays detailed information on a selected log entry.

Timeline

The Timeline is a visual navigation tool to the log data. It provides you a reference point to how the current view relates to the Query you are browsing. It also allows you to quickly move within the time range.

In the default Records, Details, and Statistics views, the chart in the timeline is hidden. You can view it by expanding the panel upwards. See Browsing Log Entries on a Timeline (page 159) for more information.

Illustration 9.4 Timeline Expanded

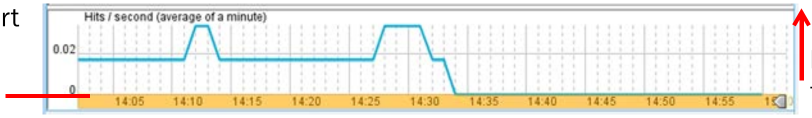

The lower part shows browsing progress.

Drag the border up to reveal chart.

Log Entry Table (Records Arrangement)

The *log entry table* in the default Records arrangement is the primary view to the logs. You can freely select which details are shown and the order of the columns. Different types of entries contain different types of information, so none of the entries use all columns.

When you right-click a cell in a log entry, the menu that opens allows you to select various actions related to the log entry. The actions vary slightly depending on the information that the cell contains. For example, right-clicking an element adds general element-specific actions (such as Properties). The actions include, but are not limited to, the following:

- **Details**: Switch to the Details view of the selected record.
- **Copy**: Copy the entry details to the clipboard.
- **View Rule**: View the rule that generated the log entry (if applicable).
- **Create Rule**: Create a new rule based on the entry. See Creating Rules From Logs (page 166).
- **Whois**: Look up the selected IP address in the online Whois database. See Checking Whois Records for IP Addresses in Logs (page 160).
- **Export**: Export records or attach records to an Incident case. See Exporting Extracts of Log Data (page 163), Exporting IPS Traffic Recordings (page 164), or Attaching Logs to Incident Cases (page 165).
- **Filter Connections**: Add basic details from the current selection to the Query panel. See Filtering Logs in the Logs View (page 156).
- **Show Connection**: Add basic details from the selected connection to the Query panel.
- **Search Related Events**: Some special events are parts of a larger chain of events. This option shows other events related to the selected log entry.
- **New Blacklist Entry**: Blacklist connections that match the entry's details. See Getting Started with Blacklisting (page 816) for more information about blacklisting.
- **Add Filter**: Add the selected detail to the filter in the Query panel. See Specifying Filters for a Query (page 156).
- **New Filter**: Create a new filter based on the selected detail. See Specifying Filters for a Query (page 156).

Details Arrangement

The **Details** arrangement of the Logs view provides an overview to an individual event. It is particularly useful for browsing alerts and records generated by Inspection rule matches.

Illustration 9.5 Logs View in the Details Arrangement

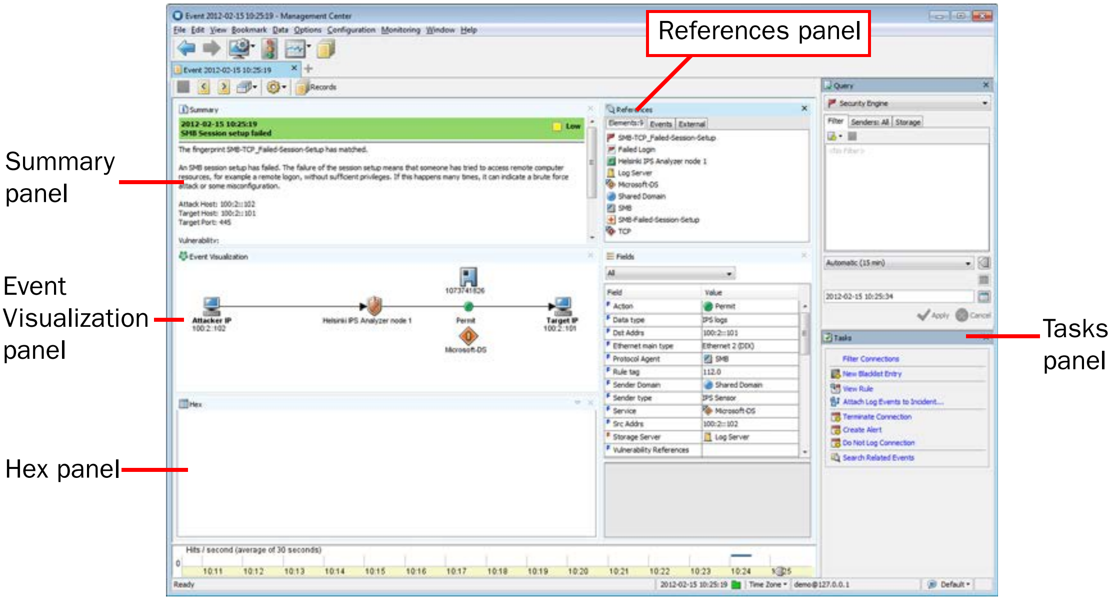

2012-02-15 10:25:19
SMB Session setup failed Low

The fingerprint SMB-TCP_Failed-Session-Setup has matched.

An SMB session setup has failed. The failure of the session setup means that someone has tried to access remote computer resources, for example a remote logon, without sufficient privileges. If this happens many times, it can indicate a brute force attack or some misconfiguration.

Attack Host: 100:2::102
Target Host: 100:2::101
Target Port: 445

Vulnerability:

Field	Value
Action	Permit
Data type	IPS logs
Dst Addrs	100:2::101
Ethernet main type	Ethernet 2 (DIX)
Protocol Agent	SMB
Rule tag	112.0
Sender Domain	Shared Domain
Sender type	IPS Sensor
Service	Microsoft-DS
Src Addrs	100:2::102
Storage Server	Log Server
Vulnerability References	

Illustration 9.6 Toolbar in the Details Arrangement

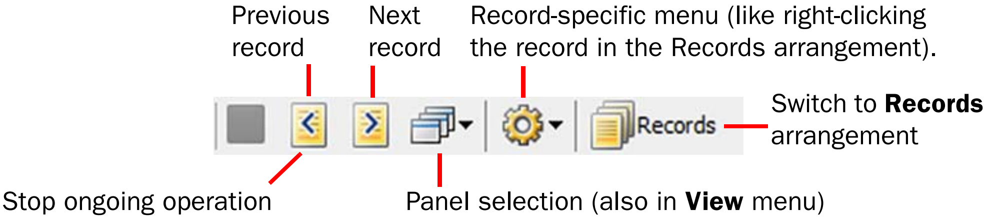

Previous record
Next record
Record-specific menu (like right-clicking the record in the Records arrangement).
Switch to **Records** arrangement
Records
Stop ongoing operation
Panel selection (also in **View** menu)

The Details arrangement has the following additional panels:

- The **References** panel (shown by default): Displays a list of elements that correspond to the details in the record and possibly additional information on related records for some special records that are part of a larger event.
- **Tasks** panel: Shortcuts to configuration tasks that you can launch based on the displayed entry (as in the Records arrangement in the right-click menu for entries).

Statistics Arrangement

In the **Statistics** arrangement of the Logs view, you can view charts of multiple events interactively. You can create a quick report of the logs that match the active Query and then further refine the Query by viewing logs that correspond to a chart segment.

Illustration 9.7 Logs View in the Statistics Arrangement

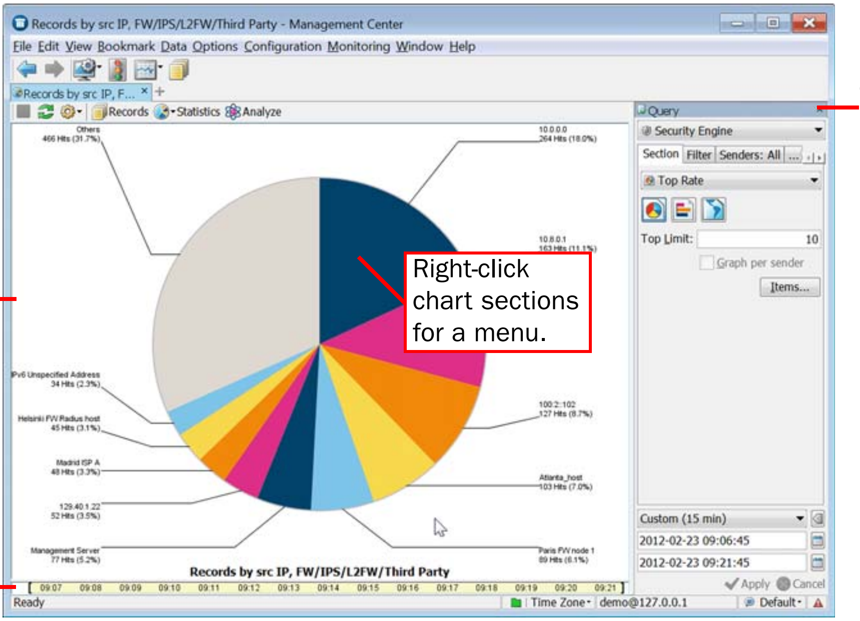

Top Limit: 10

Custom (15 min)

2012-02-23 09:06:45

2012-02-23 09:21:45

Query panel

Chart area

Right-click chart sections for a menu.

Timeline

The Query panel in the Statistics arrangement includes an additional **Section** tab. You can use the tabs to control the statistical display. The data can also be filtered in the same way as in the other Logs view arrangements.

Illustration 9.8 Toolbar in the Statistics Arrangement

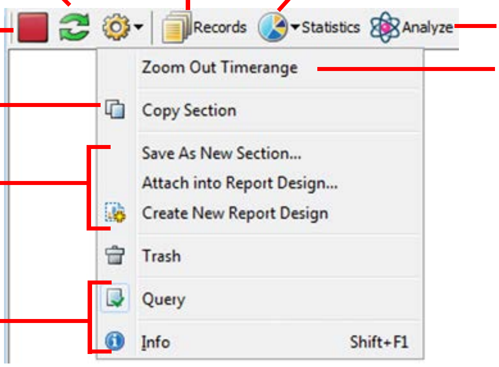

Refresh

Switch to **Records** arrangement

Shortcuts to some statistics items. More items are available in the Query panel.

Stop ongoing operation

Open **Log Analysis** view.

Generate a new chart with a wider time range.

Copy the section.

Options for creating new sections and reports.

Shortcuts to show/ hide panels.

You can also add statistical items to a section in the Statistics view.

▼ To add statistical items to a section of a Statistics view

1. In the Query panel, click **Items**. The Properties dialog opens.
2. Click **Add**.
3. Select one or more items from the list (Shift-click or Ctrl-click to select multiple items) and click **Select**.
4. Back in the Properties dialog, click **OK**.
5. In the Query panel, click **Apply** to update the view.

The chart area in the Statistics arrangement can contain a pie chart, a bar chart, a line chart, stacked line chart, or a map chart (based on an internal geolocation database). The available options depend on the chart type that is selected:

- **Top rate** charts can be displayed as a pie chart, bar chart, or a map. A top rate chart shows the total numbers of records that match the query.
- **Progress** charts can be displayed as a line chart, stacked line chart, bar chart, or stacked bar chart. A progress chart illustrates the numbers of records plotted over time (similar to the timeline, but in more detail).

When a chart is generated, you can right-click for a menu of actions related to the section and possibly the element that the section corresponds to. The actions available vary by section. Some of the most important actions are listed below:

- **Show Records**: Opens the Records arrangement filtered to show the entries that comprise the chart section you right-clicked.
- **Add to Current Filter**: Allows you to use sections to filter data by adding the section in question to the Filter tab of the Query panel.
- **Statistics** item shortcuts: Drill down to create a new chart from data that matches the previous chart section.

Log Analysis Arrangement

The **Log Analysis** arrangement provides various tools to analyze logs, alerts, and audit entries.

Illustration 9.9 Log Analysis Arrangement

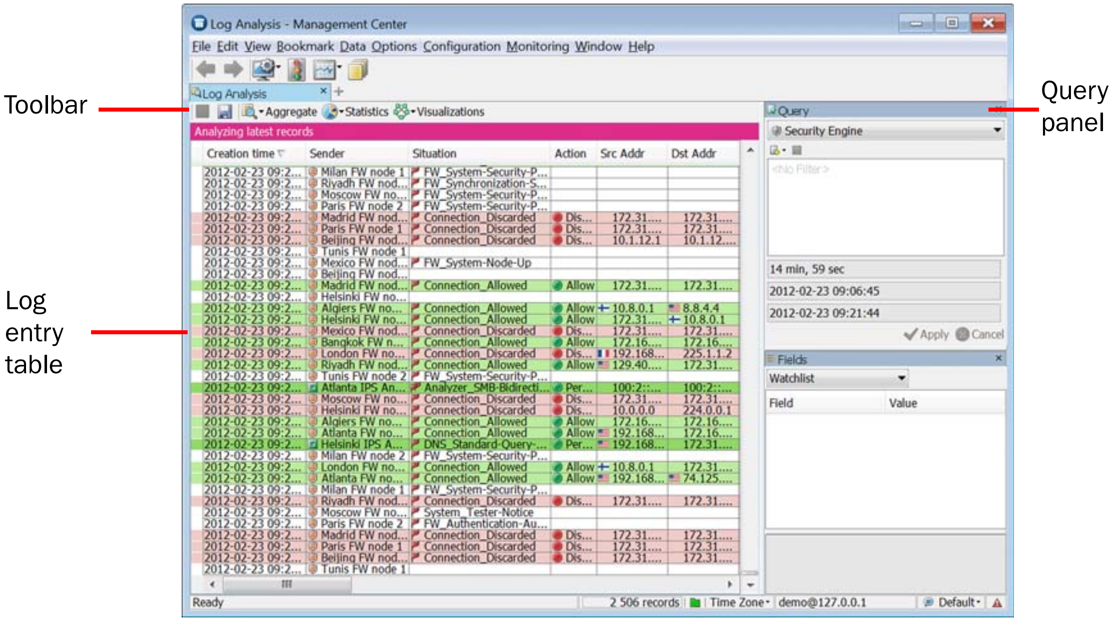

Creation time	Sender	Situation	Action	Src Addr	Dst Addr
2012-02-23 09:2...	Milan FW node 1	FW_System-Security-P...			
2012-02-23 09:2...	Riyadh FW nod...	FW_Synchronization-S...			
2012-02-23 09:2...	Moscow FW no...	FW_System-Security-P...			
2012-02-23 09:2...	Paris FW node 2	FW_System-Security-P...			
2012-02-23 09:2...	Madrid FW nod...	Connection_Discarded	Dis...	172.31....	172.31....
2012-02-23 09:2...	Paris FW node 1	Connection_Discarded	Dis...	172.31....	172.31....
2012-02-23 09:2...	Beijing FW nod...	Connection_Discarded	Dis...	10.1.12.1	10.1.12....
2012-02-23 09:2...	Tunis FW node 1				
2012-02-23 09:2...	Mexico FW nod...	FW_System-Node-Up			
2012-02-23 09:2...	Beijing FW nod...				
2012-02-23 09:2...	Madrid FW nod...	Connection_Allowed	Allow	172.31....	172.31....
2012-02-23 09:2...	Helsinki FW no...				
2012-02-23 09:2...	Algiers FW no...	Connection_Allowed	Allow	10.8.0.1	8.8.4.4
2012-02-23 09:2...	Helsinki FW no...	Connection_Allowed	Allow	172.31....	10.8.0.1
2012-02-23 09:2...	Mexico FW nod...	Connection_Discarded	Dis...	172.31....	172.31....
2012-02-23 09:2...	Bangkok FW n...	Connection_Allowed	Allow	172.16....	172.16....
2012-02-23 09:2...	London FW no...	Connection_Discarded	Dis...	192.168...	225.1.1.2
2012-02-23 09:2...	Riyadh FW nod...	Connection_Allowed	Allow	129.40....	172.31....
2012-02-23 09:2...	Tunis FW node 2	FW_System-Security-P...			
2012-02-23 09:2...	Atlanta IPS An...	Analyzer_SMB-Bidirecti...	Per...	100:2::...	100:2::...
2012-02-23 09:2...	Moscow FW no...	Connection_Discarded	Dis...	172.31....	172.31....
2012-02-23 09:2...	Helsinki FW no...	Connection_Discarded	Dis...	10.0.0.0	224.0.0.1
2012-02-23 09:2...	Algiers FW no...	Connection_Allowed	Allow	172.16....	172.16....
2012-02-23 09:2...	Atlanta FW no...	Connection_Allowed	Allow	192.168...	172.16....
2012-02-23 09:2...	Helsinki IPS A...	DNS_Standard-Query-...	Per...	192.168...	172.31....
2012-02-23 09:2...	Milan FW node 2	FW_System-Security-P...			
2012-02-23 09:2...	London FW no...	Connection_Allowed	Allow	10.8.0.1	172.31....
2012-02-23 09:2...	Atlanta FW no...	Connection_Allowed	Allow	192.168...	74.125....
2012-02-23 09:2...	Milan FW node 1	FW_System-Security-P...			
2012-02-23 09:2...	Riyadh FW nod...	Connection_Discarded	Dis...	172.31....	172.31....
2012-02-23 09:2...	Moscow FW no...	System_Tester-Notice			
2012-02-23 09:2...	Paris FW node 2	FW_Authentication-Au...			
2012-02-23 09:2...	Madrid FW nod...	Connection_Discarded	Dis...	172.31....	172.31....
2012-02-23 09:2...	Paris FW node 1	Connection_Discarded	Dis...	172.31....	172.31....
2012-02-23 09:2...	Beijing FW nod...	Connection_Discarded	Dis...	172.31....	172.31....
2012-02-23 09:2...	Tunis FW node 1				

14 min, 59 sec

2012-02-23 09:06:45

2012-02-23 09:21:44

2 506 records

demo@127.0.0.1

- To combine logs by Service or Situation, select **Aggregate→Aggregate by Service** or **Aggregate→Aggregate by Situation**.
- To sort logs by column type, select **Aggregate→Sort by Column**, and click the heading of the corresponding column.
- To view the data as charts, click **Statistics** and select one of the pre-defined statistical items. Select **Other** to select an item from a complete list of statistical items.
- To view the data as a diagram, click **Visualizations** and select one of the visualization options:

Table 9.1 Visualization Options

Option	Explanation
Audit Map	Displays a map based on how users in the system manipulate elements.
Application Usage	Displays a map of users and the applications that they use and/or access. Indicates also allowed and disallowed connections between users and applications.
Service Map	Displays access to services in the network.

You can zoom into the data presented in the visualization diagrams with the mouse wheel. Right-clicking elements in the diagrams opens a pop-up menu with various options to further analyze the elements and add them to filters. You can also, for example, drag and drop objects from the visualization diagram to the Query panel to create a filter.

What's Next?

- For more instructions on how to use the Logs view, see Browsing Log Data. To see how you can customize the Logs view, see Changing How Data Entries Are Displayed.

Browsing Log Data

Prerequisites: None

This section explains the basics of finding and browsing data in the Logs view.

Viewing Log Entry Details in the Side Panel

The Fields panel provides several alternative views to the log entry that is currently selected. The Fields panel is most useful in the Records arrangement. It shows a subset of fields and the information contained in the selected field, so that you can quickly browse the logs table for the exact details you are looking for without scrolling sideways or rearranging the columns.

If the Fields panel is not visible in your Logs view, select **View→Panels→Fields**.

Illustration 9.10 Fields Panel

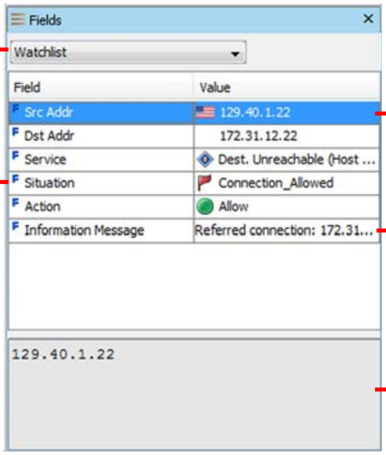

Field	Value
Src Addr	129.40.1.22
Dst Addr	172.31.12.22
Service	Dest. Unreachable (Host ...
Situation	Connection_Allowed
Action	Allow
Information Message	Referred connection: 172.31...

```
129.40.1.22
```

Select the types of fields you want to view.

Right-click a field for a menu of actions.

Right-click an IP address and select **Whois** for more information about the source of the traffic.

If the selected entry contains information in the fields included, the details are shown.

When you select a field, it is shown below as text, allowing you to copy-paste portions of the information.

The **Watchlist** item allows you to create a customized list of fields for your own use. The Watchlist is specific to each Management Client installation.

To look up a selected IP address in the online Whois database, see Checking Whois Records for IP Addresses in Logs (page 160).

If an IP address in a log entry has a country flag icon next to it, a Geolocation (for example, a street, city, or country) has been associated with it. You can view the physical location of these IP addresses in Google Maps. See Viewing IP Address Locations in the Logs view (page 124).

▼ To edit your personal Watchlist of log fields

1. In the Fields panel, select **Watchlist** in the list at the top of the panel.
2. Change the selection to your desired fields:
   - To remove fields, right-click the field and select **Remove** (to remove one field) or **Clear** (to remove all fields).
   - To add additional fields, drag and drop cells from the log entry table into the Fields panel (the value of the field is irrelevant in this case).
   - To add a field to the Watchlist from other views in the Fields panel, right-click a field in the Fields panel and select **Add to Watchlist**.

Filtering Logs in the Logs View

Efficient use of the logs requires that you filter the records displayed in the Logs view. This is done using the Query panel, which contains the following tabs for filtering data:

- **Filter** tab allows you to filter entries based on any information in the entries.
- **Senders** tab allows you to filter entries according to the component that created the entry. Filtering by sender speeds up log browsing when there are many log sending components, but you are only interested in a limited set.
- **Storage** allows you to filter entries according to the servers on which the entries are stored.

Illustration 9.11 Query Panel in the Logs View

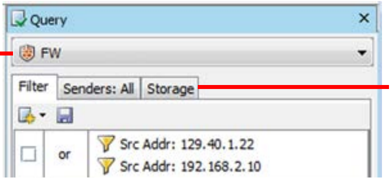

Src Addr: 129.40.1.22
Src Addr: 192.168.2.10

General type for data included.

Options on the three tabs allow setting additional filtering criteria.

What's Next?

▶ The main tools for selecting data are on the Filter tab. See Specifying Filters for a Query.

Specifying Filters for a Query

The **Filter** tab in the Query panel of the Logs view allows you to:

- create quickly local filters by dragging and dropping details from existing entries (from the logs table, Fields panel, or sections of a statistical chart)
- filter the logs according to time
- apply filtering criteria stored in Filter elements.

If the Query panel is not visible, select **View→Panels→Query.**

Illustration 9.12 Query Panel with Different Types of Filtering Criteria Selected

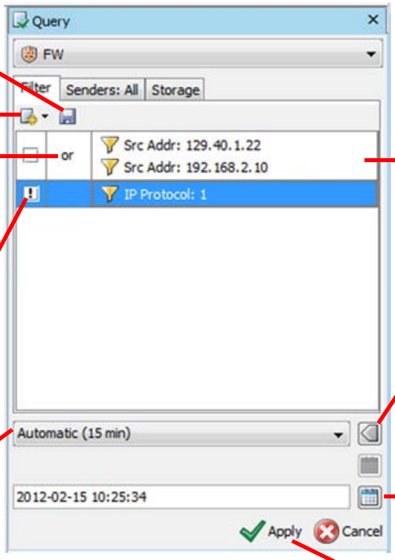

Src Addr: 129.40.1.22
Src Addr: 192.168.2.10
IP Protocol: 1

Automatic (15 min)

2012-02-15 10:25:34

Save the contents as a permanent Filter element.

Create a new filter.

Rows with several details can match any one detail (**or**) or only when all details are found in the same entry (**and**). Click the **and**/**or** cell to toggle.

Negate the row to filter *out* matching entries.

Time range to retrieve log entries: select a specific time range, or no fixed time range.

Drag and drop details from data for quick filters. Drop in empty space to add a row. Edit as needed.

Start browsing from the beginning/end of the selected time range.

Select the date from a calendar.

Click **Apply** after you make changes.

Note – The time selection refers to the entries' creation timestamp (not the reception time at the Log Server, which is also included in many entries). Internally, the system always uses universal time (UTC). The times are displayed according to the time zone selected in your Management Client's status bar.

You can drag and drop any field from the log entries to the Filters tab to create a Filter, select existing Filter elements, or add a new filtering criterion using the toolbar icon and typing in the detail. You can then further modify and utilize the Filters you have created:

- Double-click a detail (such as IP address) to open the Filter Properties dialog, in which you can change the detail manually by typing in a new value.
- Right-click a field in the log entry table or in the Fields panel and select **Add Filter: <field name>** to add the item and its value as a new filter row.
- Right-click an item in the log entry table or in the Fields panel and select **New Filter: <item name>** to define a value for the item and add it as a new filter row.
- You can add a new empty row by right-clicking a filter row or in the empty space and selecting **Row**.
- To make a search based on a word or a string, right-click the Query panel, select **New→Filter: Text,** and type your search string in the Filter Properties dialog.
- To remove a detail, right-click and select **Remove <detail description>**.
- To remove a whole row, right-click something other than a detail on the row you want to remove and select **Remove**.
- Temporarily disable a filter row by right-clicking it and selecting **Disable** from the menu.
- You can save the current filtering criteria as a permanent Filter element by clicking the Save button at the top of the Filter tab in the Query panel (see Illustration 9.12 above).

Remember to click **Apply** after you make any changes to re-filter the logs.

What's Next?

- The **Senders** tab in the Query panel allows you to select specific components for observation. See Viewing Logs From Specific Components (page 158).

Related Tasks

- For information on creating Filter elements, see Filtering Data (page 183).

Viewing Logs From Specific Components

You can filter the logs based on the component(s) that created the entries. If the Senders tab is empty, data from all components is displayed in the Logs view. The Senders tab allows you to maintain the sender filtering independent of changes on the Filter tab. Restricting the included senders makes log browsing faster when there are many components in the system.

▼ To specify senders

1. Switch to the **Senders** tab in the Query panel.
2. Click the **Select** button (arrow) at the top of the Senders tab. The Select Element dialog opens.
3. Select the element(s) you want to use as the sender(s) and click **Select**.
4. Click **Apply**. The log data is refreshed, and only logs from the selected sender(s) are displayed.

What's Next?

- The **Storage** tab in the Query panel allows you to display data based on where it is stored. See Viewing Logs From Specific Servers and Archive Folders (page 158).

Viewing Logs From Specific Servers and Archive Folders

By default, the Logs view fetches data from the active log storage folder of all data storage servers, except those Log Servers that are excluded from log browsing. You view logs from the active storage folder on specific Log and Management Servers. You can also view logs from archives stored on Logs Servers or archives stored locally on the computer where you are using the Management Client. Management Servers store audit data only. Log Servers store all logs and alerts in addition to audit data regarding their own operation.

▼ To specify storage

1. Switch to the **Storage** tab in the Query panel.
2. In the server list, select the servers and storage folders that you want to include.
3. Click **Apply**. The log data is refreshed and filtered according to the selected servers and folders.

Related Tasks

- Changing Log Server Configuration Parameters (page 320)
- Archiving Log Data (page 1018)

Analyzing Logs, Alerts, and Audit Entries

▼ To analyze logs, alerts and audit entries

1. Select **Monitoring→Logs**. The Logs view opens.
2. Click **Analyze**. The Log Analysis view opens.

The Log Analysis view provides various tools to analyze logs, alerts, and audit entries. For more details on the different analysis options, see Log Analysis Arrangement (page 154).

Saving Snapshots of Log, Alert, and Audit Entries

You can save snapshots of log, alert, and audit entries in the Log Analysis view. The snapshots are saved on the Management Server. The saved snapshots are listed the Monitoring view.

▼ To save snapshots of log, alert, and audit entries

1. Select **Monitoring→Logs**. The Logs view opens.
2. Select the entries for the snapshot as explained in the Browsing Log Data (page 155). You can select a maximum of 100000 entries.
3. Click **Analyze**. The Log Analysis view opens.
4. Click the Save icon in the toolbar. The Snapshot Properties dialog opens.
5. Enter a **Name** for the snapshot and click **OK**.

Viewing Snapshots of Log, Alert, and Audit Entries

The snapshots of log, alert, and audit entries are listed in the Monitoring view.

▼ To open a snapshot stored on the Management Server

1. Select **Configuration→Configuration→Monitoring**. The Monitoring view opens.
2. Browse to **Other Elements→Monitoring Snapshots→Logs→Management Server**.
3. Right-click a snapshot and select **Open**. The snapshot opens for viewing.

Browsing Log Entries on a Timeline

You can navigate logs using the Timeline. In the Records and Details view arrangements, part of the Timeline is hidden by default. You can view the full timeline by dragging its upper edge. Depending on your selection, the timeline either allows you to browse freely (the **Automatic** option) or stops when the first or last entry within the specified time range is reached.

Illustration 9.13 Timeline Minimized

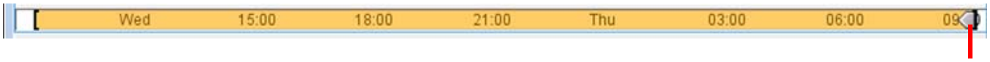

When you are browsing within a set time range, you cannot accidentally browse out of the time range set in the Query panel. Square brackets are shown at each end to mark this.

Drag the arrowhead to browse. The arrow also indicates the selected start position (from the beginning or the end of the time range).

Illustration 9.14 Timeline Expanded

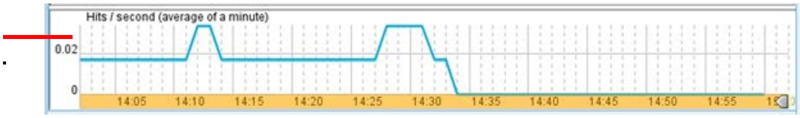

Chart plots the number of matching entries over time.

Viewing Temporary Log Entries

The Logs view has two operating modes. In the normal mode, you can browse entries freely from any time period. When you activate the **Current Events** mode by clicking the Play icon in the toolbar, the log entries update automatically to show the stream of log entries as they arrive at the Log Server(s). Typically, you must filter out some entries to keep the pace of the Current Events mode slow enough that you can keep up with the entries.

The Current Events mode also displays temporary entries that are not stored on the Log Server (Transient log entries and log entries that are pruned out before permanent storing) so you may see more logs than in the normal mode. Temporary entries only exist within the current view and are permanently lost when the view is refreshed or closed. The updates in the Current Events mode are automatically deactivated when you select an entry or start browsing manually.

Note – Under some operating conditions, a small portion of log entries may arrive in mixed order. Because the Current events view attempts to maintain a logical flow of events, out-of-sequence entries are not shown. You may see a notification message if this happens.

Sorting Log Entries

By default, log entries are sorted according to their creation time. You can alternatively sort log entries according to any other column heading. Because large numbers of logs may require significant resources to be sorted, the Logs view may shorten your selected time range if your current Query matches too many records to be efficiently sorted.

▼ To sort log entries according to the column heading

➥ Click the column heading according to which you want to sort the log entries. The log entry table is refreshed with the logs sorted according to the column heading you selected. This may take a while.

Note – The Current events view is always sorted according to entry creation time. Sorting can only use stored data, so any temporary data visible in the Current events view is permanently lost if you change sorting.

Checking Whois Records for IP Addresses in Logs

To get more information about the source of the traffic that triggered a log entry, you can look up the Whois record of IP addresses in log entries. The Whois record contains registration information and related contact details provided at the time of domain registration. The contents of the Whois record vary depending on the information provided by the owner of the domain or network segment. In the case of IP addresses used by customers of an ISP, the information shown in the Whois record is usually the ISP's information.

The Whois information is queried from the relevant Regional Internet Registry (RIR). These include, but are not limited to, the ARIN (American Registry for Internet Numbers), the RIPE NCC (Réseaux IP Européens Network Coordination Centre), and the APNIC (Asia Pacific Network Information Centre). More information about the main RIRs can be found at the following links:

- ARIN at a Glance: https://www.arin.net/about_us/overview.html.
- RIPE Database: http://www.ripe.net/db/index.html.
- About APNIC: http://www.apnic.net/about-APNIC/organization.

The Whois query is performed by the computer running the Management Client. To be able to perform Whois queries, the security policy applied on the computer running the Management Client must meet the following criteria:

1. DNS queries must be allowed so that the Management Client can resolve the relevant RIR server IP address.
2. Opening TCP43 (whois) connections must be allowed.

▼ To check the Whois record for an IP address in a log entry

↳ Right-click an IP address and select **Whois**. The Whois record for the IP address is displayed.

Tip – **You can also view the physical location of an IP address in Google Maps. See Viewing IP Address Locations from the Whois Information Dialog (page 124).**

Related Tasks

- ▶ Monitoring Connections on a Map (page 120)
- ▶ Viewing Geolocations and IP Addresses in Google Maps (page 123)

Changing How Data Entries Are Displayed

Prerequisites: None

Increasing and Decreasing Text Size in Data Entries

You can increase, decrease, and reset (“set normal”) the text size in the log entry table from the **View→Text Size** menu.

Changing the Time Zone for Log Browsing

The system uses universal time (UTC) internally. The times in the Logs view are changed to your selected time zone as they are displayed. By default, this is the local time zone of the computer you are using. Changing the time zone does not affect the actual timestamps of the log entries.

If the times in the log entries seem incorrect, make sure both the time and time zone are set correctly in your operating system, on the Management Server, and on all Log Servers.

Illustration 9.15 Time Zone Selection in the Logs View

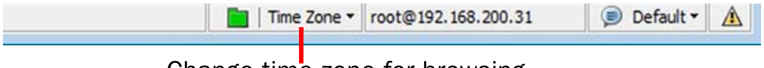

Change time zone for browsing.

Changing Data Columns in the Log Entry Table

You can customize the columns shown in the Logs view. You can add and remove columns and change the order in which the columns are displayed. You can also see subsets of log columns in the Fields panel. See Viewing Log Entry Details in the Side Panel (page 155).

You can arrange the columns in the following ways:

- Drag the column header to a different location to change the order of the columns.
- Double-click the column header to expand the column to the width of its contents.
- Right-click a column header for a menu of actions for adjusting the column widths.

▼ To select columns that are shown in the log entry table

1. Open the Logs view.
2. Select **View→Column Selection**. The Column Selection dialog opens.

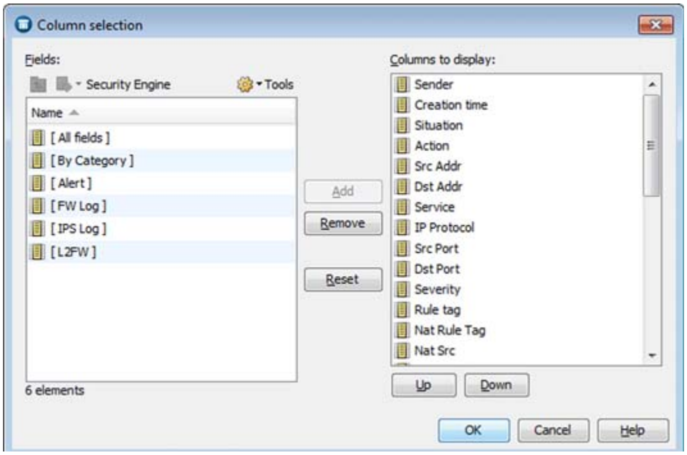

3. Select **Add** and **Remove** to include and exclude selected fields. The Columns to Display list on the right shows your selections. See Log Entry Fields (page 1212) for short descriptions of all log fields.
4. Use **Up** and **Down** to organize selected fields on the Columns to Display list. Fields at the top are shown at the left of the log record table.
5. Click **OK**.

Resolving Log Details to DNS Names or System Elements

To make logs easier to read, you can view the IP addresses, protocols, and senders as DNS names or Management Center elements. Resolving is for display only and does not affect stored log data.

Note – IP address and port resolving works by comparing the information in the logs to internal and external information sources. If the information available is incorrect or ambiguous, the result may not reflect the actual hosts and services involved in the communications. For example, if a detail matches two elements, the first match is used for display even if the other element was used in the corresponding rule.

▼ To enable/disable Sender resolving

↳ Select **Options→Resolve Senders**. When enabled, the IP addresses of engines and SMC servers are resolved using the element definitions.

▼ To enable/disable IP address resolving by DNS

↳ Select **Options→Resolve Addresses by DNS** to resolve IP addresses using DNS (if available).

▼ To enable/disable IP address resolving by elements

↳ Select **Options→Resolve Addresses by Elements** to resolve IP addresses using element definitions.

▼ To enable/disable the display of icons

↳ Select **Options→Show Icons**. When enabled, the data is displayed with the icons of the elements involved, where applicable.

Deactivating/Activating Log Entry Highlighting

Different types of logs can be highlighted with different colors. There is a default color scheme in place, but the log color filters can be customized for all administrators or even individually in each Administrator element's properties.

▼ To turn the log entry colors on/off

↳ Select/deselect **Options→Use Color Filters** in the main menu. See Customizing Log Colors (page 258) for information on how the colors are changed.

Exporting Data from the Logs View

Prerequisites: None

Exporting Extracts of Log Data

Log, alert, and audit data can be exported directly from the Logs view, but note the following:

- You can alternatively use the Log Data Tasks tool to export logs, allowing you to schedule export tasks that are executed automatically. See Getting Started with Log Data Management (page 1016).
- If you want a human-readable output and there is no need for further processing of exported data, we recommend you save the entries as PDF or HTML instead. See Saving Elements, Log Data, Reports, and Statistics (page 51).
- If you want to export traffic recordings, see Exporting IPS Traffic Recordings (page 164).

For a limited number of entries, a simple copy-paste is the quickest export method.

▼ To copy a limited number of logs in CSV format

↳ Select data entries and copy-paste them to the other application, for example a spreadsheet application. The entries are copied with the column titles in the CSV (comma-separated values) format.

For large numbers of entries and other export formats, use the export command instead.

▼ To export data entries from the Logs view

1. (*Optional*) To export only some of the entries that match your current Query, select some data entries in the Records arrangement. Ctrl- or Shift-click to select several entries.
2. Right-click one of the entries and select **Export→Export Log Events**. The Export Logs dialog opens.
3. Select the **File Export Format** that matches your needs.
   - **XML** logs are suitable for conversion to different formats, such as HTML, using external XML conversion tools (you must develop your own conversions that produce the desired results).
   - **CSV** logs are suitable for reading and further processing in spreadsheets and other similar use.
   - **CEF** logs are suitable for converting logs info syslog format.
   - **Archive ZIP** and **Archive** save the logs in the system's internal format for possible later use in this or some other Stonesoft Management Center.
4. In **Export**, select if you only want to export the selected entries or all of the entries that match the filter criteria currently specified in the Query panel.
5. Select the **Destination File**:
   - For Archive, you can select archive directories defined in the Log Server's configuration file (see Changing Log Server Configuration Parameters (page 320)).
   - For all other formats, you must specify a file name and select whether to export to the **Server "export" Directory** (`<installation directory>/data/export/` on the Log or Management Server) or to a location of your choice on your computer (**Local Workstation**).
6. For formats other than Archive, specify what is done if the specified file already exists.
   - (*CSV only*) Select **Append** to add the new logs to the end of the existing file.
   - Select **Overwrite** to replace the existing file with the new one.
   - Select **Use Number in File Name** to add a number to the file name to distinguish it from the existing file name.
   - Select **Fail Task** to abort the export if the file already exists.
7. *(Local Workstation exports only)* Select **Open File After Export** to view the exported file in the operating system's default application for the file type.
8. Click **OK**. The Task Status panel opens and shows the progress of the export.

Related Tasks

- ▶ Getting Started with Log Data Management (page 1016)

Exporting IPS Traffic Recordings

You can set IPS Inspection rules to record network traffic as a logging option in both the Exceptions and the Rules tree as explained in Editing Inspection Policies (page 699). These recordings are stored on the Log Servers. Recordings generated by the "Excerpt" option are shown directly in the Logs view (select **View→Panels** in the Hex panel). Longer recordings, however, are meant to be viewed in an external application and are not directly viewable.

To view the recording, you can retrieve the recording through the log entry as explained below or define a Task for exporting recordings, for example, from a specific time range as explained in Exporting IPS Traffic Recordings (page 164).

▼ To retrieve traffic recordings for selected log entries

1. Select log entries that are associated with recordings.
   - To browse for more entries that have a recording, change the selection in the Fields panel to **Full Capture** (available when an entry that has an associated recording is selected). The Record ID field is displayed with an identification number for entries that are associated with a recording.
2. Right-click a selected entry and select **Export→Export IPS Recordings**. The Export Selected IPS Recordings dialog opens.
3. Select the **File Export Format** that matches your needs (**PCAP** or **SNOOP**).
4. Select the **Destination File**:
   - Select the **Server "export" Directory** to export to a file on the Log Server (`<`*`installation directory`*`>/data/export/`)
   - Select **Local Workstation** to save on your computer.
   - You must specify the file name in both cases.
5. Select one of the following if the specified file already exists (**Append** is not supported for traffic recordings):
   - **Overwrite**: replace the existing file with the new one.
   - **Use Number in File Name**: add a number to the file name to distinguish it from the existing file name.
   - **Fail Task**: abort the export if the file already exists.
6. Click **OK**. The Task Status panel opens showing the progress of the export.

Attaching Logs to Incident Cases

You can attach logs, alert entries, and audit entries to incident cases directly from the Logs view. You can attach single logs, a group of logs, or all logs that match your filtering criteria. You can attach logs to an existing Incident Case, or you can create a new Incident Case and attach logs to it. See Attaching Logs and Audit Entries to Incident Cases (page 210) for more information about attaching logs to Incident Cases. For more information about creating new Incident Cases, see Creating a New Incident Case (page 209).

Creating Rules From Logs

Prerequisites: None

You can use log entry details to generate new rules. To convert a log entry to a rule, the log entry must be created based on a rule (the entry contains a rule tag). Creating a rule this way allows you to make quick exceptions to the current policy. You can create the following types of rules:

- A rule that allows a connection from an entry that logs stopped traffic.
- A rule that stops a connection from an entry that logs allowed traffic.
- A rule that changes the log level or stops the logging of matching connections.

▼ To create a rule based on one or more log entries

1. Select the log entries you want to include in the operation. You can select multiple log entries to create several rules in the same operation. Do not include incompatible entries in the selection:
   - If you select multiple log entries, the **Sender** of all entries must be the same component.
   - All selected entries must have a value in the **Rule Tag** field (entries must be created by rules in a policy).
2. Right-click a (selected) log entry and select one of the options under **Create Rule** in the menu that opens. The selection determines how the handling of matching connections is changed. The New Rule Properties dialog opens.
3. (*Optional*) Click **Select** and change the policy to which the new rule is added (for example, to insert the rule in a Sub-Policy instead of the main policy).
4. (*Optional*) Edit the **Comment** (this is added to the rule's Comment cell).
5. Select the **Action**. All actions create the displayed rules at the beginning of the first insert point in the selected policy. You can also optionally install the policy with the new rule or open the policy for editing (with the new rule highlighted for you).

Note – You cannot edit the rule in this dialog. Select Add Rules and Edit the Policy to edit the rule.

6. Click **OK**.

What's Next?

- ▶ If you selected to refresh the policy, the policy installation dialog opens. See Installing Policies (page 652) for more instructions.
- ▶ If you selected to edit the policy, the policy opens for editing. See Using the Policy Editing View (page 661) for more instructions.

CHAPTER 10

REPORTS

You can process data from logs and engine statistics into easy-to-read diagrams, charts, and tables. The reports you generate are always based on a Report Design, which can be one of the predefined designs, a modified version of it, or a custom design you create yourself.

The following sections are included:

Getting Started with Reports

Prerequisites: None

Reports allow you to gather together and visualize the data that interests you the most in an easy-to-read format to provide an overview of what is happening in the network. Reports are configured and generated in the Monitoring view. You can view reports as graphs, charts, tables, and geolocation maps.

What You Can Do with Reports

You can generate reports based on two types of data:

- *Log data* consists of distinct events (for example, a connection opening or closing). It contains all the details of the events including the exact time when the events occurred.
- *Counter data* consists of pre-processed statistics summaries that are based on sums or averages of events or traffic units within a certain period. Counter data that is older than an hour is consolidated to an accuracy of one hour.

You can generate reports based on predefined Report Designs and Report Sections, or on Report Designs that you have created yourself. You can use your own Style Template for the PDF creation to give the reports a unified corporate look.

You can view the reports in the Management Client and in the Web Portal.

To export your reports, you can print the generated reports into PDF or HTML documents for convenient sharing and printing. You can also directly e-mail reports as they are generated.

In addition to creating and generating reports based on Report Designs, you can also create simple reports in the Logs view (see Statistics Arrangement (page 152)).

You can also generate special System reports from data about system configuration and events to help you provide auditing information in compliance with regulatory standards.

Configuration Overview

The general workflow for creating and generating reports is as follows:

1. Customize an existing Report Design or create a new Report Design. See Creating and Modifying Report Designs (page 169).
2. Customize or add Report Sections or Items. See Creating and Modifying Report Sections (page 172).
3. Generate a report. See Generating and Viewing Reports (page 175).

What's Next?

- ▶ To generate a report based on an existing Report Design, proceed to Generating and Viewing Reports (page 175).
- ▶ Otherwise, begin by Creating and Modifying Report Designs (page 169).

Creating and Modifying Report Designs

Prerequisites: None

Report Designs determine how to process the data and how the results are displayed. They can also determine which template is used for PDF exports and which charts appear on them. Ready-made Report Designs serve as a useful guide for constructing your own Report Designs. You can also create new custom Report Designs.

Illustration 10.1 Report Design

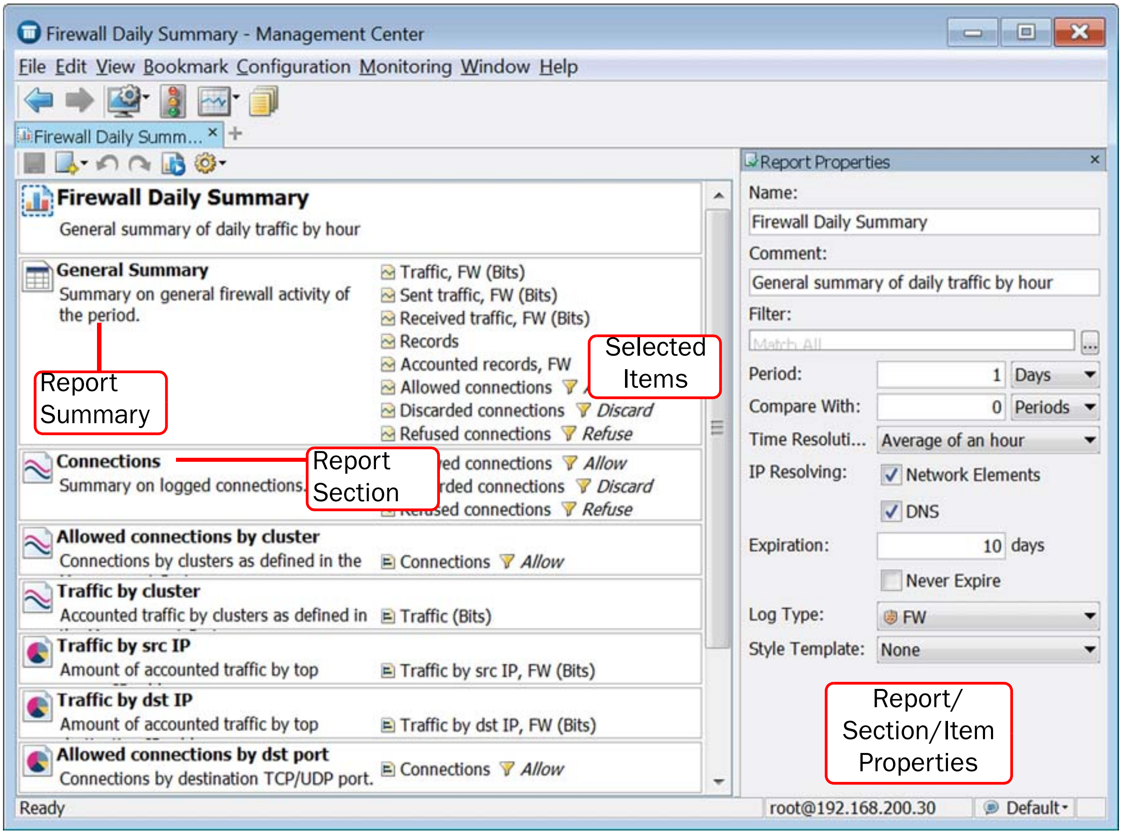

Firewall Daily Summary - Management Center
File Edit View Bookmark Configuration Monitoring Window Help
Firewall Daily Summ...
Report Summary
Report Section
Selected Items
Report/ Section/Item Properties

Firewall Daily Summary
General summary of daily traffic by hour

General Summary
Summary on general firewall activity of the period.
Traffic, FW (Bits)
Sent traffic, FW (Bits)
Received traffic, FW (Bits)
Records
Accounted records, FW
Allowed connections
Discarded connections *Discard*
Refused connections *Refuse*

Connections
Summary on logged connections.
...ed connections *Allow*
...rded connections *Discard*
...used connections *Refuse*

Allowed connections by cluster
Connections by clusters as defined in the
Connections *Allow*

Traffic by cluster
Accounted traffic by clusters as defined in
Traffic (Bits)

Traffic by src IP
Amount of accounted traffic by top
Traffic by src IP, FW (Bits)

Traffic by dst IP
Amount of accounted traffic by top
Traffic by dst IP, FW (Bits)

Allowed connections by dst port
Connections by destination TCP/UDP port.
Connections *Allow*

Report Properties

Name: Firewall Daily Summary
Comment: General summary of daily traffic by hour
Filter: Match All
Period: 1 Days
Compare With: 0 Periods
Time Resoluti...: Average of an hour
IP Resolving: ☑ Network Elements ☑ DNS
Expiration: 10 days
□ Never Expire
Log Type: FW
Style Template: None

Ready root@192.168.200.30 Default

What's Next?

- If you want to modify an existing Report Design, proceed to Modifying Report Designs (page 170).
- Otherwise, proceed to Creating New Report Designs (page 170).

Modifying Report Designs

▼ To edit a Report Design

1. Select **Configuration→Configuration→Monitoring** from the menu. The Monitoring Configuration view opens.
2. Right-click the Report Design that you want to use and select **Edit Report Design**. The Report Design opens for editing.

What's Next?

▶ Continue by Creating and Modifying Report Sections (page 172)

Creating New Report Designs

You create a new Report Design in two steps: first you create a Report Design, and then you add items to it.

▼ To create a new Report Design

1. Select **Configuration→Configuration→Monitoring** from the menu. The Monitoring Configuration view opens.
2. Right-click **Report Designs** and select **New Report Design**. The Report Design Properties dialog opens.
3. Select a template for the new Report Design and click **OK**. The template opens for editing with the Report Properties panel on the right.

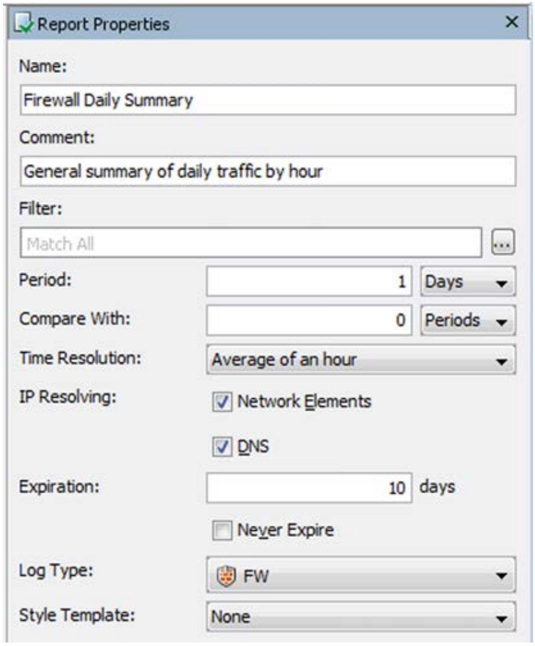

Report Properties

Name: Firewall Daily Summary

Comment: General summary of daily traffic by hour

Filter: Match All

Period: 1 Days

Compare With: 0 Periods

Time Resolution: Average of an hour

IP Resolving: ☑ Network Elements ☑ DNS

Expiration: 10 days

□ Never Expire

Log Type: FW

Style Template: None

4. Type in a **Name** for the new Report Design and an optional **Comment**.

5. *(Optional)* Select a **Filter**. For instructions on how to create a new filter, see Creating and Editing Local Filters (page 187).

Note – When you generate a report, all the filters defined in the report task properties, in the Report Design, Report Sections, and individual Report Items are used to filter the data. If the filters do not intersect, empty Report Sections may be generated in the report.

6. Adjust the other properties as needed.

Option	Explanation
Period	Defines the default time frame for this report. This affects the dates offered by default when creating a report using this design. The longer period you choose, the more data is included in the report. The level of detail in the charts may have to be reduced in order to keep them legible by adjusting the time resolution. This is done automatically, but you can also change it by hand. All report items are not compatible with the shortest period length.
Compare With	Allows you to include data from a previous period of the same length in your reports to make comparison easier.
Time Resolution	The level of detail in the progress charts and tables. A small time resolution increases the level of detail, but having too much detail may make the produced charts difficult to read. The time resolution and the available choices are automatically adjusted when you change the Period.
IP Resolving	Activates IP address resolving in the reports, using Network Elements and/or DNS queries. Network Elements or DNS addresses are shown instead of IP addresses when available.
Expiration	Defines the number of days after which the reports generated based on this Report Design are automatically deleted. If you select **Never Expire**, you must manually delete all the reports generated using this Report Design when you no longer need them.
Log Type	Allows you to filter the data included in log-based report items by the type of log (alert, audit log, firewall log, IPS log, SSL VPN log, third-party device log).
Style Template	Allows you to select the Style Template to be used with the PDF printing of the new Report.

What's Next?

▶ Continue by Creating and Modifying Report Sections (page 172).

Creating and Modifying Report Sections

A Report Design consists of one or more Report Sections, which define parameters for all items within it. Each Report Section in the Report Design creates a separate chart and/or a table in the generated report.

To browse predefined Report Sections, select **Configuration→Configuration→Monitoring** and go to **Reports→Sections**. You can modify and create new Report Sections in a Report Design.

Modifying Report Sections

▼ To modify report sections

1. Select a section in the Report Design. The section properties are displayed in the Section Properties panel.
2. Modify the section properties as described in the table below:

Table 10.1 Section Properties

Option	Explanation
Name	Change the name of the section.
Comment (*Optional*)	Add a comment for your own reference.
Filter (*Optional*)	Select a filter to be used with the section. See Defining Filters (page 185).
Log Type	Define what type of logs are used in the section.
Related Element (*Optional*)	Select the element(s) for which data is generated in the section. Used mainly in System Reports.
Export	Select **All** to export the report with data presented in tables and diagrams.
	Select **Diagram** to export the report with data presented in diagrams.
	Select **Table** to export the report with data presented in tables.
Traffic Unit	Define the traffic unit for the data in charts and tables.
Diagram Settings	**Progress** produces a chart and/or a table that shows the progress of items as time passes.
	Top Rate produces a chart and/or a table highlighting the values with the highest number of occurrences.
	Drill-down Top Rate preprocesses the data and then produces a chart and/or table from it.
	Summary Table produces a table with all the values from items included in the report.
	System Information produces a table with information about current system status.

Table 10.1 Section Properties (Continued)

Option	Explanation
Chart Type	Select the appropriate chart type. You can change this selection when you start the report generation manually. The options available depend on the Diagram Settings that you have selected.
Top Limit	If you have selected **Top Rate** or **Top Rate Progress** as the Section Type, enter the number of items to be included.
Graph per sender (*Optional*)	Select **Graph per sender** to show a separate graph for each sender. This option is only available for Progress type items.
Items	Click **Items** to select new statistical items for the Section. For more information, see Creating and Modifying Report Items (page 173).

3. (*Optional*) Change the order of the Sections in the Report Design by dragging and dropping them on top of each other. The changes you make to a section are applied immediately.

Creating New Report Sections

▼ To create new report sections

1. Right-click the Report Design and select **Add New Section**.
2. Select a section from the list. The new section is added to the Report Design and the background of the new section is highlighted. The Section Properties panel opens.
   - If you cannot find an appropriate section in the list, select **Other**, and select a section in the **Select Section** dialog.
   - If you want to add a section based on a statistical item, select **Create from Item** and select an item from the Select Item dialog.
3. Modify the section properties as instructed in Modifying Report Sections (page 172).
4. To add new Sections, repeat as necessary.

What’s Next?

▶ If you are finished customizing the Report Design, select **Save** or **Save as** from the menu to save the Report Design.

▶ To edit the statistical items in the Report Sections, continue by Creating and Modifying Report Items.

Creating and Modifying Report Items

You can add and edit statistical items in Report Sections. Statistical items count the following types of data:

- log entries (referred to as *records* in the item names)
- summaries of some log field included in those entries (such as traffic volumes in log data that contains accounting information)
- specifically gathered statistical data (counter items)
- system information to summarize current configuration information in the Management Server’s internal database.

Creating Report Items

▼ To create new Items in a Report Section

1. Select the Report Section to which you want to add Items.
2. Click **Items** in the Section Properties panel. The Properties dialog opens.
3. Click **Add** under the Items field. The Select Item dialog opens.

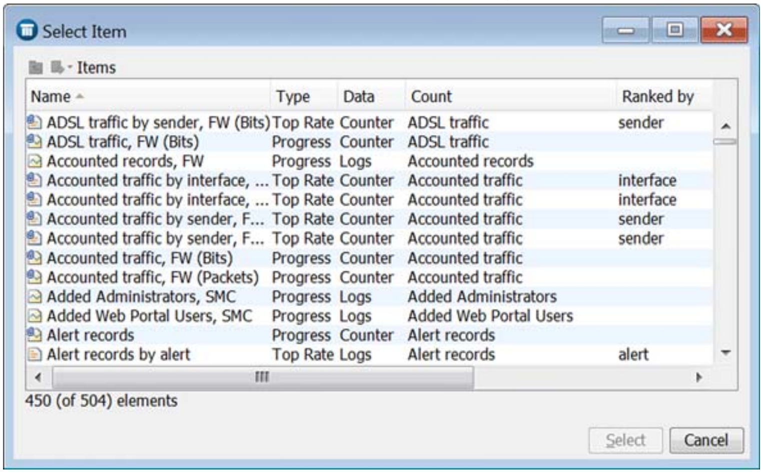

Name	Type	Data	Count	Ranked by
ADSL traffic by sender, FW (Bits)	Top Rate	Counter	ADSL traffic	sender
ADSL traffic, FW (Bits)	Progress	Counter	ADSL traffic	
Accounted records, FW	Progress	Logs	Accounted records	
Accounted traffic by interface, ...	Top Rate	Counter	Accounted traffic	interface
Accounted traffic by interface, ...	Top Rate	Counter	Accounted traffic	interface
Accounted traffic by sender, F...	Top Rate	Counter	Accounted traffic	sender
Accounted traffic by sender, F...	Top Rate	Counter	Accounted traffic	sender
Accounted traffic, FW (Bits)	Progress	Counter	Accounted traffic	
Accounted traffic, FW (Packets)	Progress	Counter	Accounted traffic	
Added Administrators, SMC	Progress	Logs	Added Administrators	
Added Web Portal Users, SMC	Progress	Logs	Added Web Portal Users	
Alert records	Progress	Counter	Alert records	
Alert records by alert	Top Rate	Logs	Alert records	alert

450 (of 504) elements

4. Select the Item(s) that you want to add:
   - Double-click to select a single Item.
   - Ctrl-click or Shift-click to select multiple Items and click **Select**.
5. To add additional Items, repeat from Step 2 as necessary.
6. Click **OK.**
7. To save the changes that you have made to the Report Design, select **File→Save** or **File→Save As**.

What's Next?

- ▶ If you want to edit the Items, proceed to Modifying Report Items.
- ▶ Otherwise, proceed to Generating and Viewing Reports (page 175).

Modifying Report Items

▼ To modify Items in a Report Section

1. Select the Report Section in which you want to modify the Item selection.
2. Click **Items** in the Section Properties panel. The Section Properties dialog opens.
3. Right-click the item that you want to modify and select **Properties**. The Item Properties dialog opens.

4. Modify the Item properties and click **OK**:

Option	Explanation
Name	Enter a new name for the Item.
Log Type	Select the type of log data you want to use for the Item.
Filter	Click **Edit** to select filters to be used with the Item.

5. (*Optional*) Change the order of the added Items by dragging and dropping them on top of each other.
6. Click **OK**.
7. To save the changes that you have made to the Report Design, select **File→Save** or **File→Save As**.

What's Next?

- ▶ Generating and Viewing Reports (page 175)

Generating and Viewing Reports

Prerequisites: Creating New Report Designs

Reports are generated from the Report Designs that are displayed in the Reports tree of the Monitoring Configuration view.

Generating a Report

▼ To generate a report

1. Select **Configuration→Configuration→Monitoring** from the menu. The Monitoring Configuration view opens.
2. Right-click a Report Design and select **Start**. The Report Operation Properties dialog opens.

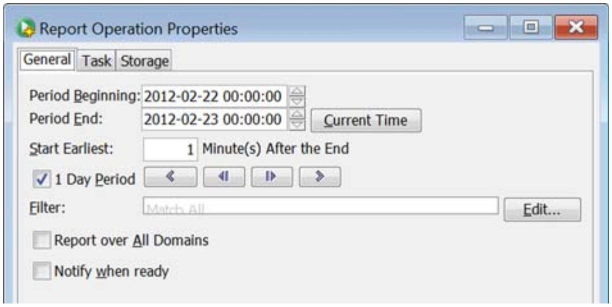

Period Beginning: 2012-02-22 00:00:00

Period End: 2012-02-23 00:00:00

Start Earliest: 1 Minute(s) After the End

☑ 1 Day Period

Filter: Match All

□ Report over All Domains

□ Notify when ready

3. Select the **Period Beginning** and **Period End** for the report. Enter the date as indicated, and enter the time using the 24-hour clock in the time of the workstation you are using to run the Management Client.
   - The **1 Day Period** option *(default)* defines the previous day as the reporting period, beginning at 00:00:00 that day and ending exactly 24 hours later.

4. (*Optional*) Enter the delay in minutes after the end date and time in the **Start Earliest** field.
   - Report generation begins after the specified delay has passed. The delay is to ensure that the system has all relevant data available.
   - If the **Period End** date and time are in the future, the report is generated on the specified date at the specified time once the **Start Earliest** delay had passed.
5. (*Optional*) If you want to restrict the data included in the report, select a filter.

Note – When you generate a report, all the filters defined in the report task properties, in the Report Design, Report Sections, and individual Report Items are used to filter the data. If the filters do not intersect, empty Report Sections may be generated in the report.

6. *(Shared Domain only)* If you want to make a report containing the information concerning all Domains, select **Report over All Domains**.
7. (*Optional*) Select **Notify when ready** to receive a notification in the Management Client when the report generation is complete.

What's Next?

▶ To select how often to automatically repeat the report generation, proceed to Defining the Report Task.

▶ To select the Management and Log Servers from which data is used to generate the report, proceed to Selecting Data Sources (page 178).

▶ Otherwise, click **OK** to start generating the report. Proceed to Viewing Reports (page 179).

Defining the Report Task

▼ To define the Report Task

1. Switch to the **Task** tab.
2. Select how often you want to **Repeat** the report generation (the Report Design determines the available choices).

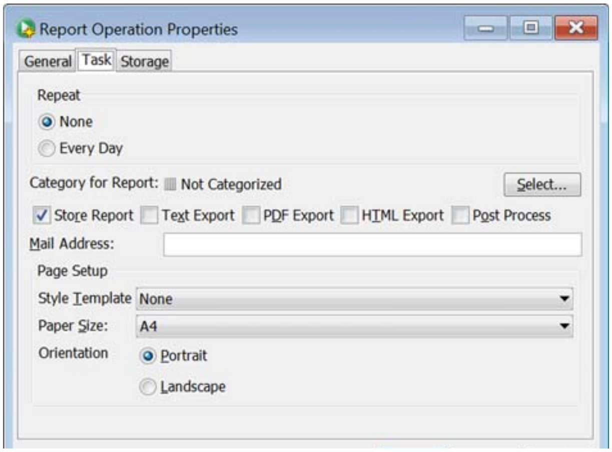

Report Operation Properties

Repeat

● None

○ Every Day

Category for Report: Not Categorized

☑ Store Report □ Text Export □ PDF Export □ HTML Export □ Post Process

Mail Address:

Page Setup

Style Template: None

Paper Size: A4

Orientation: ● Portrait ○ Landscape

3. Select one or more outputs to be produced directly, or leave only **Store Report** selected to view the report before deciding if you want to process it further:
   - **Text Export**: The report is stored as a text file on the Management Server, in the *<installation directory>/data/reports/files/[Report Design name]/* directory. The report is named according to the time range chosen.

Note – If you installed the Management Server in the `C:\Program Files\Stonesoft\Management Center` **directory in Windows, some program data may be stored in the** `C:\ProgramData\Stonesoft\Management Center\data` **directory.**

   - **PDF Export/HTML Export**: The report is stored as a PDF file/HTML file on the Management Server (same directory as above).
   - **Post Process**: The report is generated according to options chosen and then a script is launched, by default `SgPostProcessReport.bat` located on the Management Server in the *<installation directory>/data/reports/bin* directory. See the *SMC Reference Guide* for more information.

4. (*Optional*) Enter the **Mail Address** to which the completed report is sent directly out as e-mail. To add several addresses, separate the addresses with a comma.

Note – The SMTP server for sending reports must be defined in the Management Server's properties. See Modifying a Management Server Element (page 342).

5. *(For PDF exports only)* Define the Page Setup settings by selecting a **Style Template**, **Paper Size** and **Orientation** when exporting the report to PDF format. For information on how to create new Style Templates, see Adding Style Templates for PDF Output (page 52).

What's Next?

▶ To select the Management and Log Servers from which data is used to generate the report, proceed to Selecting Data Sources (page 178).

▶ Otherwise, click **OK** to start generating the report. Proceed to Viewing Reports (page 179).

Selecting Data Sources

By default, the Management and Log Servers are selected as data sources.

▼ To select data sources

1. Switch to the **Storage** tab.

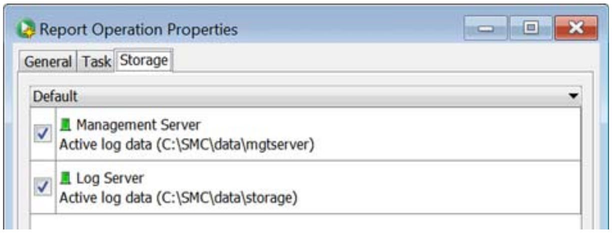

☑ Management Server
Active log data (C:\SMC\data\mgtserver)

☑ Log Server
Active log data (C:\SMC\data\storage)

2. Select the data storage type.
   - **Default**: The Management and Log Servers are used as the data sources.
   - **Primary archive**: Archived data is used as the data source.
   - **Custom**: A combination of archived data and data provided by the Management and Log Servers is used as the data source.
3. Select the Management Server(s) and Log Server(s) from which you want to include data in the report.
4. Click **OK** to start generating the report.

What's Next?

▶ Proceed to Viewing Reports (page 179).

Canceling Ongoing Report Tasks

If the report you are generating includes large amounts of data, generating the report may take a long time. This may be the case when, for example, the time frame is very wide and the data filter in use does not restrict the data sources.

▼ To cancel an ongoing report task

1. Select **Configuration→Configuration→Monitoring** from the menu. The Monitoring Configuration view opens.
2. Select **Reports→Stored Reports**.
3. Click **Today** to see the list of reports that are running, but have not yet finished.
4. Right-click the task you want to cancel and select **Abort** or **Delete**.
   - Aborted tasks can be edited and restarted by double-clicking them.
   - Deleted tasks are permanently removed.

Viewing Reports

After you have generated a report, it is available for viewing in the Stored Reports tree. A report may be automatically deleted according to the expiration setting defined in the Report Design. If you want to keep a generated report permanently, we recommend exporting the report as instructed in Exporting Reports (page 180).

▼ To view a report

1. Select **Configuration→Configuration→Monitoring** from the menu. The Monitoring Configuration view opens.
2. Select **Reports→Stored Reports**. The created reports are automatically grouped according to their creation time (Today, Yesterday, Week, Month).
   - To view all generated reports, click the Tools icon and select **All.**
   - To display the reports created using a particular design, click the Tools icon and select **By Design**.

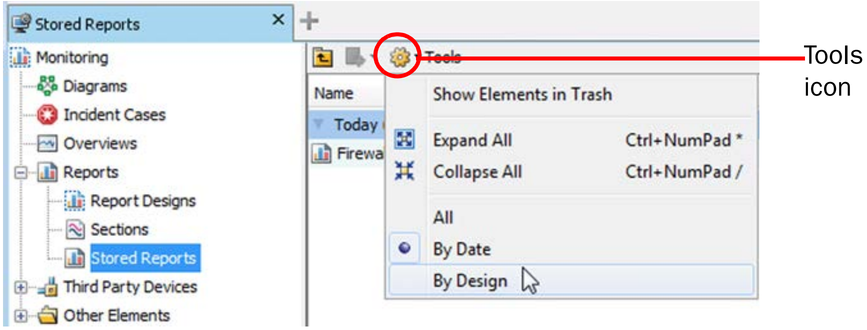

3. Double-click the report you want to view. The contents of the report are shown.

Tip – If you want to see the data of a report section in table format, right-click the section and select **Show Table**. The table is added to the section.

Related Tasks

- ▶ Exporting Reports (page 180)

Changing the Properties of Generated Reports

To change the category and expiration of a report

1. Select **Configuration→Configuration→Monitoring** from the menu. The Monitoring Configuration view opens.
2. Select **Reports→Stored Reports**.
3. Right-click the title of the report you want to modify and select **Properties**. The report properties dialog opens.

Illustration 10.2 Report Properties

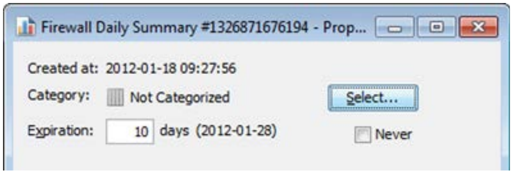

Firewall Daily Summary #1326871676194 - Prop...

Created at: 2012-01-18 09:27:56

Category: Not Categorized

Select...

Expiration: 10 days (2012-01-28)

□ Never

4. *(Optional)* Click **Categories** to select the category for the Report. For more information on Categories, see Using Categories (page 76).
5. *(Optional)* If you want to change the time when the report is automatically deleted, change the value in the **Expiration** field. Select **Never** if you want to delete the report yourself.
6. Click **OK**.

Related Tasks

- Generating and Viewing Reports (page 175)
- Exporting Reports (page 180)

Exporting Reports

Prerequisites: Generating a Report

After you have generated a report, you may want to share it or process the data further. You can export a generated report as a PDF or HTML file.

What's Next?

- To export a previously generated report to a PDF file for viewing or printing, proceed to Exporting a Report as a PDF File (page 181).
- To export a previously generated report to an HTML file for viewing or printing, proceed to Exporting a Report as an HTML File (page 182).
- To send a report directly as e-mail, proceed to E-Mailing Reports (page 182).

Note – You can also export reports as tab-delimited text files while generating reports. For more information, see Defining the Report Task (page 176).

Exporting a Report as a PDF File

When you export a report file as a PDF file, a default template is automatically used as the background for the report. Optionally, you can select a Style Template as the background. If you want to use a Style Template for the report you are about to export, you must have the template ready on the computer you are using.

▼ To export a report as a PDF file

1. Select **Configuration→Configuration→Monitoring** from the menu. The Monitoring Configuration view opens.
2. Select **Reports→Stored Reports**.
3. Double-click the report you want to print to PDF. The Report opens.
4. Select what part of the report to print:
   - To print the whole report as a PDF file, select **File→Print**. The printing dialog opens.
   - To print a section of the report as a PDF file, right-click the section you want to print and select **Print**. The printing dialog opens.
5. Select PDF as the **Format**.
6. Select the printing option:
   - Select **Print to PDF Reader** to send the generated PDF directly to your default PDF reader.
   - Select **Print to File** and browse to the location where you want to save the generated PDF.
7. (*Optional*) Select the **Style Template**, **Paper Size** and **Paper Orientation** for the PDF. For more information on the use of Style Templates, see Adding Style Templates for PDF Output (page 52).
8. Click **OK**. The PDF is generated.

Check the output. You may have to adjust the headers and footers in your template if the report text or charts are placed on top of your template background (see Adding Style Templates for PDF Output (page 52)).

Exporting a Report as an HTML File

Prerequisites:

When you export a report as an HTML file, a default template is automatically used as the background for the report.

▼ To print a report to an HTML file

1. Select **Configuration→Configuration→Monitoring** from the menu. The Monitoring Configuration view opens.
2. Select **Reports→Stored Reports**.
3. Double-click the report you want to print as HTML. The Report opens.
4. Select what part of the report to print:
   - To print the whole report as HTML, select **File→Print**. The printing dialog opens.
   - To print a section of the report as HTML, right-click the section you want to print and select **Print**. The printing dialog opens.
5. Select HTML as the **Format**.
6. Click **Browse** and select the directory where you want to save the HTML file.
7. Click **OK**. The HTML file is saved in the defined directory. The HTML report opens in your default web browser.

E-Mailing Reports

Reports can be e-mailed directly as they are generated. You must define the SMTP server in the Management Server's properties before you can e-mail reports. See Modifying a Management Server Element (page 342).

After defining the SMTP settings, simply enter the e-mail address on the Task tab of the Report Operation Properties dialog to send the selected output as an e-mail. To add several addresses, separate the addresses with a comma.

Creating a System Audit Report

Prerequisites: None

The System Report contains details of Administrator and Web Portal user activity and access settings, system configuration, changes to the Firewall and IPS engines, and the configuration of the Management Server. The System Report is intended to help you provide the required data for auditing in compliance with regulations, such as the Payment Card Industry (PCI) Data Security Standard. The System Report is generated, exported, and edited in the same way as other types of reports: the only difference is the content of the report.

What's Next?

▶ To generate System reports, proceed to Generating and Viewing Reports (page 175).

Chapter 11

Filtering Data

Filters allow you to select data based on values that it contains. Most frequently, you will need filters when viewing logs, but filters can also be used for other tasks, such as exporting logs and selecting data for reports.

The following sections are included:

Getting Started with Filtering Data

Prerequisites: None

What Filters Do

You can use Filters to select data for many operations such as viewing log entries in the Logs view or generating statistical reports. Filters allow you to efficiently manage the large amounts of data that the system generates. Filters select entries by comparing values defined in the Filter to each data entry included in the filtering operation. The operation can use the filter to either include or exclude matching data.

Filters can be used for selecting data in the following tasks:

- Browsing logs, alerts, and audit data. See Browsing Logged Data (page 147).
- Browsing in all session monitoring views. For example, see Monitoring Connections, Blacklists, VPN SAs, Users, and Routing (page 111).
- Creating reports. See Reports (page 167).
- Selecting which logs administrators with restricted accounts or Web Portal User accounts are allowed to view. See Administrator Accounts (page 247).
- Defining how logs are highlighted in the Logs view. See Customizing Log Colors (page 258).
- Exporting logs to an external Syslog server. See Forwarding Log Data to Syslog (page 323).
- Browsing which IP addresses, ports, and protocol are on the security engines' blacklists. See Checking Connections, Blacklists, VPN SAs, Users, and Routing (page 111).
- Pruning log data. See Pruning Log Data (page 1024).
- Exporting and deleting log data and alerts. See Exporting Log Data (page 1026) and Deleting Log Data (page 1021).
- Creating Correlation Situations to analyze engine and Log Server events. See Defining Context Options for Correlation Situations (page 762).

Filter Types

There are two types of filters used in the Management Client:

- Local filters that are specific to a view or an element.
- Permanent Filter elements that you can use anywhere in the Management Client. You can also create new permanent Filters. There are predefined permanent Filters as well.

How Filters Are Created

You can create filters in four basic ways:

- Based on criteria you define: you can create a new local filter or Filter element and define any combination of filtering criteria in the Filter properties, constructing the filter completely yourself.
- Based on other Filters: you can duplicate a Filter element or copy-and-paste parts of filter contents to other filters to create a variation of previously defined filtering criteria.
- Based on existing log entries: you can create local filters in views where you view logs and save them as permanent Filter elements. See Specifying Filters for a Query (page 156).
- Based on element configuration: some local filters are created automatically by your selections in specific views or elements.

Parts of a Filter

Filters are constructed from the following parts (see Basics of Constructing Filters (page 185) for more detailed information):

- The *fields* that you want to match in the data (for example, there are separate fields for source IP address and port in logs). You can filter data according to any field.
- The *values* in those fields that you want to match (for example, the exact port number or IP address you are interested in).
- *Operations* define the way in which the fields and values are matched to data entries (especially if there are several fields included as the filtering criteria).

What's Next?

▶ Defining Filters.

Related Tasks

▶ Organizing Filter Elements (page 193)

Defining Filters

Prerequisites: None

You can create filters in various views in the Management Client. You can define local filters that are specific to the element or view in which the local filters were created. You cannot use a local filter anywhere else in the Management Client. However, you can save a local filter as a permanent Filter element that can be used anywhere in the Management Client.

What's Next?

▶ To understand the basic construction of filters and see basic example filters, proceed to Basics of Constructing Filters (page 185).

▶ To define filters specific to elements or views, proceed to Creating and Editing Local Filters (page 187).

▶ To define Filter elements that can be used anywhere in the Management Client, proceed to Creating and Editing Filter Elements (page 190).

Basics of Constructing Filters

Filters are constructed with *fields*, which define the details that you want to match in the data. The fields usually have a specific *value* to match (such as a particular IP address). An *operation* defines how the fields and values in the filter are matched to fields and values in data.

Each field in a filter is attached to one of these operations:

- Calculations (*bitwise and*, *sum of*)
- Comparisons (for example, *equal to*, *greater than, smaller than*)
- Logicals (*AND*, *NOT*, *OR*).

Example A filter that matches a single source IP address:

`Src Addr` **equal to** `192.168.1.101`

Where “Src Addr” is a field, “equal to” is the operation and the IP address is a value. As in these examples, operations are displayed in boldface in Filters.

You can also match several values at a time, as shown in the examples below.

Example A filter that matches any non-empty value for destination port:

`Dst Port` **is defined**

Example A filter that matches all destination ports between 1024 and 49152:

`Dst Port` **between** `1024` **and** `49152`

Example A filter that matches any of three alternative source ports:

`Dst Port` **in** `51111, 52222, 53333`

To create more complex filters, you can add logical operations NOT, AND, and OR. The NOT operation negates the criteria you set.

Example A filter that matches all destination ports except ports between 1024 and 49152:

NOT

`Dst Port` **between** `1024` **and** `49152`

When you add more than one field to a filter, you must define how the fields are used in relation to each other with either AND (all field values must be present in the same entry) or OR (a data entry matches the filter if any one field value is found).

Example A filter that matches if the destination port is lower than 1024 and the source is a particular IP address:

AND

`Src Addr` **equal to** `192.168.1.101`

`Dst Port` **smaller than** `1024`

Example A filter that matches either of two destination ports:

OR

`Dst Port` **equal to** `80`

`Dst Port` **equal to** `8080`

You can apply the AND and OR operations to other AND and OR statements to create more complex statements. You can also negate whole AND and OR sections with NOT.

What’s Next?

- To create local filters in different views or for elements, see Creating and Editing Local Filters.
- To define Filter elements that can be used anywhere in the Management Client, proceed to Creating and Editing Filter Elements (page 190).

Creating and Editing Local Filters

You can create local filters in various views in the Management Client. Local filters are valid only in the view in which they are created.

Local filters that you create in the following Monitoring views are temporary in the sense that they disappear when you close the view:

- Logs view
- Connections view
- Blacklist view
- Users view
- VPN SAs view
- Routing Monitoring view

Local filters created for the following elements are saved with the element:

- Report Designs
- Administrator elements
- Correlation Situations

You can also save local fllters as permanent Filter elements in all views. These Filter elements can then be used anywhere in the Management Client.

▼ To create a local filter

1. Open a view or element for editing.
2. Create a new local filter in one of the following ways:

View/Dialog	Configuration
Monitoring view (Logs, Connections, Blacklist, Users, VPN SAs, Routing)	Select the type of filter that you want to create in the **Filter** tab of the Query panel. • **New→Filter** to define a filter expression from all available fields. • **New→Filter: Text** to define a search filter based on a word or a string.
Reports view	1. Double-click the Report Item you want to filter. The Report Item Properties dialog opens. 3. Select **Add→New→Filter** or **Add→New→Filter:Text**.
Administrator Properties dialog	1. Switch to the **Permissions** tab of the Administrator Properties dialog. 2. Click **Edit** for Log Filters. The Local Filter Properties dialog opens. 3. Select **Add→New→Filter** or **Add→New→Filter:Text**.

View/Dialog		Configuration
Correlation Situation Properties dialog	Sequence Context	1. Switch to the Context tab. 2. **Select** Sequence as the **Context**. 3. Double-click the **Entry to** or **Exit from** cell. The Local Filter Properties dialog opens. 4. Select **Add**→**New**→**Filter** or **Add**→**New**→**Filter:Text**.
	Match Context	1. Switch to the Context tab. 2. **Select** Match as the **Context**. 3. Click **Edit** for the **Filter** field. The Local Filter Properties dialog opens. 4. Select **Add**→**New**→**Filter** or **Add**→**New**→**Filter:Text**.
	Group Context	1. Switch to the Context tab. 2. **Select** Group as the **Context**. 3. Double-click the **Event match** cell of a **Member**. The Local Filter Properties dialog opens. 4. Select **Add**→**New**→**Filter** or **Add**→**New**→**Filter:Text**.
	Compress Context	1. Switch to the Context tab. 2. **Select** Compress as the **Context**. 3. Click **Edit** for the **Compress filter** field. The Local Filter Properties dialog opens. 4. Select **Add**→**New**→**Filter** or **Add**→**New**→**Filter:Text**.

3. (*Optional*) Enter a **Name** for the new filter in the Filter Properties dialog.
4. Select the setting for **Undefined Value Policy** and define the filtering criteria as explained in Adding and Modifying Filtering Criteria in Filters (page 191).
5. Repeat steps 2-4 to add any further filtering criteria, then click **Apply** in the Filter Properties dialog. The new filter is added to the Local Filter Properties dialog.

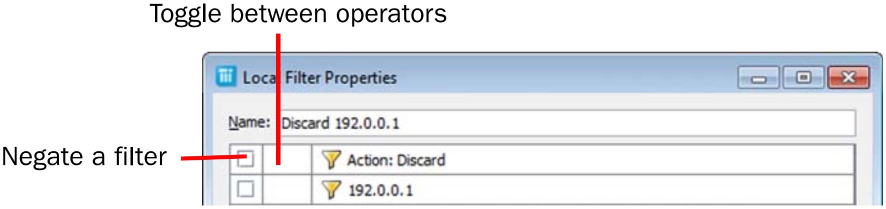

Name: Discard 192.0.0.1

□		Action: Discard
□		192.0.0.1

6. (*Optional*) Modify the filter:
   - To negate a filter row, click its check box. This filters out entries matching the filter.
   - If a row contains more than one item, click the and/or cell to toggle between the **and** and **or** operators.
   - Right-click a filter row to open a list of options to modify, add, and delete filters.
   - To add new filters click **Add**. In a Monitoring view, drag and drop new log entries to the Filter tab.
7. Click **OK**. In a Monitoring view, click **Apply**.

Related Tasks

- Saving Local Filters

Saving Local Filters

Local filters are specific to the view or element for which they have been created. Saving local filters as permanent Filter elements allows you to apply these filters anywhere in the Management Client.

▼ To save a local filter as a permanent Filter element

1. Open a view or element for editing.
2. Save the local filter in one of the following ways:

View/Dialog	Configuration
Monitoring view (Logs, Connections, Blacklist, Users, VPN SAs, Routing)	Click **Save** in the Filter tab of the **Query** panel. The Filter Properties dialog opens.
Reports view	1. Select the section whose filter you want to save filter and double-click the **Filter** field in the Section Properties panel. 2. Click **Save**. The Filter Properties dialog opens.
Administrator Properties dialog	1. Switch to the **Permissions** tab and double-click the **Filter** field. The Local Filter Properties dialog opens. 2. Click **Save**. The Filter Properties dialog opens.
Correlation Situation Properties dialog	1. Switch to the **Context** tab. 2. Double-click the filter. The Local Filter Properties dialog opens. 3. Click **Save**. The Filter Properties dialog opens.

3. Enter a **Name** for the filter and click **OK**. The Local Filter Properties dialog opens.
4. Click **OK**. The filter is saved as a permanent Filter element.

Related Tasks

- ▶ Creating and Editing Filter Elements (page 190)
- ▶ Adding and Modifying Filtering Criteria in Filters (page 191)
- ▶ Removing Filtering Criteria from Filters (page 192)

Creating and Editing Filter Elements

You can create Filter elements in the Monitoring view. You can also save local filters as Filter elements. See Saving Local Filters (page 189). If you want to create a simple filter and you can easily find a matching entry in the Logs view, it may be more convenient to create a local filter and then save it as a permanent Filter element. See Filtering Logs in the Logs View (page 156).

▼ To create a Filter element

1. Select **Configuration→Configuration→Monitoring**. The Monitoring Configuration view opens.
2. Expand the **Other Elements** branch.
3. Right-click **Filters** in the element tree and select **New→Filter**. The Filter Properties dialog opens.
4. Enter a **Name** for the new Filter.
5. Select the setting for **Undefined Value Policy** and Define the filtering criteria as explained in Adding and Modifying Filtering Criteria in Filters (page 191).
6. Click **OK**. The new filter element is created.

Editing Filter Elements

▼ To edit Filter elements

1. Select **Configuration→Configuration→Monitoring**. The Monitoring Configuration view opens.
2. Expand **Other Elements→Filters→All Filters** in the tree.
3. Right-click the filter and select **Properties**. The Filter Properties dialog opens.
4. Define the filtering criteria as explained in Adding and Modifying Filtering Criteria in Filters (page 191).
5. Click **OK**.

Adding and Modifying Filtering Criteria in Filters

The steps below explain the general workflow for constructing filters. You can also modify existing criteria using the tools explained in this workflow. Often, you can utilize existing definitions in new filters, for example, if you copy-and-paste criteria from other existing filters or save a Logs view Query as a Filter element.

▼ To add criteria to a Filter

1. Select the setting for **Undefined Value Policy**. This setting defines how data entries are matched when a field is included in the filter, but is not present in the entry (for example, if a filter defines a range of destination ports, but the operation encounters a log entry of traffic over a protocol that does not use ports, such as ICMP).
   - **False by comparison** (*default*): The AND operations are *false*. As a result, also the OR operation is *false*. The event does not match the Filter.
   - **False by filter**: The AND operations are *undefined* (neither *true* nor *false*). As a result, also the OR operation is *undefined*. The setting interprets the *undefined* result as *false*. The event does not match the Filter.
   - **True by filter**: The AND operations are *undefined* (neither *true* nor *false*). As a result, also the OR operation is *undefined.* The setting interprets the *undefined* result as *true*. The event matches the Filter.
   - **Undefined**: The AND operations are *undefined* (neither *true* nor *false*). As a result, the OR operation is also *undefined*. The Undefined setting passes the *undefined* value to the component that uses the log data, which interprets the *undefined* result as *false*. The event does not match the Filter.

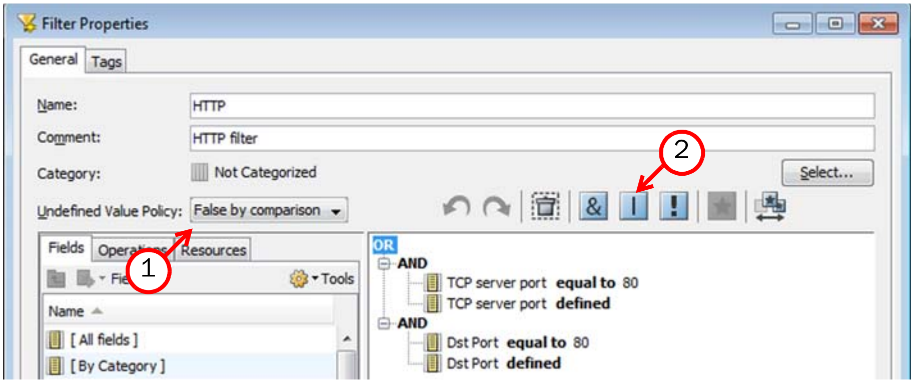

Name: HTTP

Comment: HTTP filter

Category: Not Categorized

Undefined Value Policy: False by comparison

```
OR
  AND
    TCP server port  equal to  80
    TCP server port  defined
  AND
    Dst Port  equal to  80
    Dst Port  defined
```

2. If there is no comparison operation (AND or OR) at the correct level, add one using the shortcut buttons above the editing panel. You cannot add a criterion without a comparison that defines the relationship between the criteria you add.
   - You can nest comparisons to create more complex filters, for example, you can create two AND sections under an OR condition to match either of two sets of criteria.
3. Drag and drop a log field from the **Fields** tab in the left panel to the expression. The Field also adds the default operation and value for that field.
   - The filters are organized in different categories. For example, **Access Control** contains fields relevant to Access rules and **Inspection** contains fields relevant to Inspection rules. Each field is included in several categories.
   - You must drop the field on the comparison operation (the word AND or OR) in which you want to include the field, unless the expression is empty (reads **<Add expression>**).
   - See Log Entry Fields (page 1212) for a table that describes the log fields.
4. (*Optional*) To change the field-specific operation, right-click the field that was added and select the operation from the **Change to** submenu.

- If the field contains one or more values (such as a specific IP address), you may first have to delete or remove the value(s) to be able to change the operation.
- The selection available depends on the selected field and whether the field already contains one or more values. The most common of these are (see the *Management Center Reference Guide* for more information):
  - **Equal to** allows you to match a single value (for example, an IP address).
  - **Between** allows you to match a range (for example, a range of TCP/UDP ports).
  - **Is Defined** allows you to match any non-empty value in the field.
  - **In** allows you to match any of several alternative values (for example, a whole network or both an IPv4 address and an IPv6 address).

5. To change the type of the included value or add additional values (for fields and operations that can accept changes/additions), right-click the criterion you are defining and select a type of value from the **Add** submenu.
   - This allows you, for example, to make a source IP address field accept a whole network instead of the default single IP address and to insert additional values into an In operation.
6. Depending on the operation and type of field, define the value or values you want the filter to match by typing the information, selecting an option from a list (such as a rule action) or dragging and dropping an element into the field.
   - The Is Defined operation does not include a specific value.
7. Continue adding fields in the same way until your desired criteria is defined. If you make a mistake, you can use the Undo shortcut above the editing panel or remove items as explained in Removing Filtering Criteria from Filters (page 192).
8. Click **OK**. The Filter you created is now available for use.

Removing Filtering Criteria from Filters

You can remove values or whole rows from Filter elements. If you want to replace a field value with a different value of the same type (for example, replace one IPv4 address with a different IPv4 address), you can edit the existing value without removing the old value.

▼ To remove filtering criteria from a Filter

1. Right-click the exact field/value/section you want to remove and select **Remove**.
   - If you right-click a field that has no specific value defined, the field is removed without further prompting.
   - If you right-click a specific value (or placeholder for a value), the value and its placeholder is removed without further prompting.
   - If you right-click a field that contains a specific value or that contains several values (or several empty placeholders for values), the Remove dialog opens.
2. If the Remove dialog opens, select the part of the criteria definition to remove and click **OK**.
   - **Remove whole sub-expression** deletes the whole section (the row you right-clicked and all its subordinate rows).
   - **Remove one value** deletes the value you select in the list (including the placeholder). The list includes all values from the row you right-clicked (including empty placeholders).
   - **Remove all values** deletes all values (including the placeholders) on the row that you right-clicked.
   - **Remove selected expression only** deletes the criteria (entire row) but leaves any subordinate rows intact. If any of the subordinate rows are not valid to be moved up a level, the operation fails.

Organizing Filter Elements

Prerequisites: Creating and Editing Filter Elements

Filters can be organized using Filter Tags. The Filter Tags you create are added as sub-branches to the **By Filter Tag** branch of the filter listing. Adding Filter Tags to Filter elements you create makes it easier to find Filters when you want to apply them to some operation. You can select several Filter Tags for each Filter that you create as well as for any of the default filters that are already in the system.

Creating New Filter Tags

▼ To create a Filter Tag

1. Select **Configuration→Configuration→Monitoring**. The Monitoring Configuration view opens.
2. Expand **Other Elements** in the tree.
3. Right-click **Filters** in the tree and select **New→Filter Tag**. The Filter Tag Properties dialog opens.
4. Type a **Name** for the Filter Tag.
5. (*Optional*) Enter a **Comment** for your own reference.
6. Click **OK**.

Changing the Tag of a Filter

Filters are arranged according to Filter Tag in the By Filter Tag branch. You can select any number of Filter Tags for Filters that you have created to organize the tree.

▼ To add or remove Filter Tags for Filters

1. Select **Configuration→Configuration→Monitoring**. The Monitoring Configuration view opens.
2. Expand **Other Elements→Filters** in the tree and select one or more Filter element(s) that you want to categorize differently.

Tip – **If you are removing the Filter Tag references in Filters so that you can delete a Filter Tag, select all filters in the All Filters branch. The selection is allowed to include filters that do not refer to the Filter Tag you want to remove.**

3. Right-click a selected Filter for a menu and do one of the following:
   - Select **Add Tag→Filter Tag** and select the Filter Tag you want to add in the dialog that opens. We recommend that you do not assign the default **System** Filter Tag to any Filter element yourself.
   - Select the Filter Tag you want to remove from the **Remove Tag** submenu. You cannot remove the **System** or **Correlations** Filter Tags from default system elements.

Alternatively, you can edit the Filter Tags of Filter elements created in your system on the Tags tab in each Filter element’s Properties dialog.

Chapter 12

Working With Diagrams

Diagrams allow you to visualize your network security environment.

The following sections are included:

Getting Started with Diagrams

Prerequisites: None

Diagrams allow you to generate a model of the elements you have already configured, or design a model of your network and configure Management Center elements at the same time. You can also use diagrams to view and monitor the status of the elements in your system.

What You Can Do with Diagrams

Diagrams allow you to:

- Maintain and understand the network structure.
- Monitor the status of your network graphically.
- Illustrate your network topology.
- Configure Management Center elements and other network devices while designing your network.
- Store and print network topologies.

What Should I Know before I Begin

- There are three types of diagrams: *Connectivity Diagrams* that show the status of the connections between elements that belong to the same configuration, *VPN Diagrams* that show the status of VPN connections, and *IP Diagrams* that are diagrams of an IP network.
- In addition to creating diagrams manually, you can create diagrams automatically from elements that are monitored in the System Status view. The System Status view also automatically displays a Connectivity Diagram or VPN Diagram to show the connectivity status of an engine you select.

Configuration Overview

1. Create a new diagram (see Creating Diagrams (page 197)).
2. Define the diagram background color and image (see Defining the Diagram Background (page 198)).
3. Insert elements to the diagrams either manually or automatically (see Adding Elements to Diagrams (page 199)).
4. Customize the diagram layout (see Arranging Elements in Diagrams (page 201)).
5. Make manual or automatic connections between the elements in the diagram (see Connecting Elements in Diagrams (page 201)).
6. Create relationships between diagrams (see Creating Links Between Diagrams (page 202)).

Related Tasks

- ▶ Viewing Diagrams (page 203).
- ▶ Printing Diagrams (page 204).
- ▶ Exporting Diagrams as Images (page 205).

Creating Diagrams

Prerequisites: None

Diagrams are created and modified in the Diagram Editor.

There are two methods for creating diagrams. You can create a diagram from previously configured elements or add new elements to the Management Center and draw the diagram simultaneously. When you create network elements as you draw the diagram, the elements are added to the system configuration with all the other elements. If you make changes to the configured elements (such as an IP address change), the changes are automatically updated in your diagram. This reduces the need to update your documentation.

Illustration 12.1 Diagram Editing Toolbar in Diagram Editor

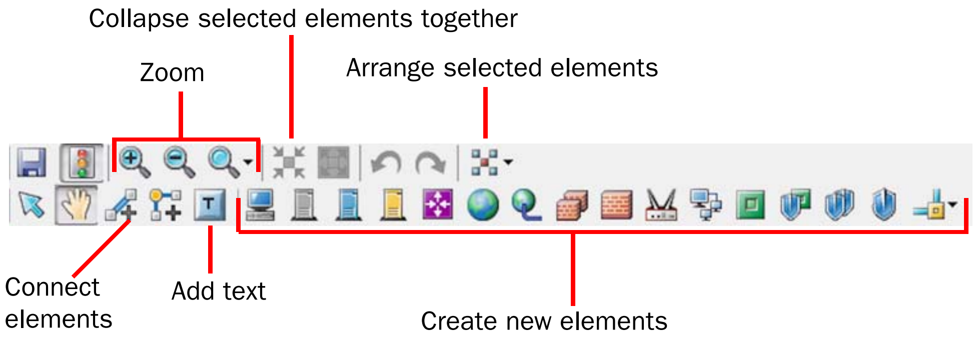
Collapse selected elements together
Zoom
Arrange selected elements
Connect elements
Add text
Create new elements

▼ To create a new diagram

1. Select **Configuration→Configuration→Monitoring** from the menu.
2. Right-click **Diagrams** and select **New→Connectivity Diagram**, **New→VPN Diagram**, or **New→IP Diagram** according to which diagram type you want to create. A blank diagram opens.

What's Next?

- If you want to change the background image or color, continue by Defining the Diagram Background.
- Otherwise, proceed to Adding Elements to Diagrams (page 199).

Defining the Diagram Background

Background images and colors enhance the appearance of your diagrams and make layout easier. For example, you can use a map as the background image and arrange the elements in the diagram according to their geographical location.

▼ To define the diagram background

1. Right-click the diagram and select **Edit**. The diagram opens for editing.
2. Right-click the diagram's background and select **Properties**. The Diagram Background Properties dialog opens.
3. Define the background properties as explained in the table below:

Option	Explanation
Color (*Optional*)	The color of the background.
Background	Select whether to use an image, a map, or an empty background. Click Select to select the image.
Watermark (*Optional*)	Select this option if you want the elements in the diagram to better stand out from the background image.
Keep Aspect Ratio (*Only with Image background*)	If this option is selected the height and width ratio of the original image is retained in the background image.
Width (*Only with Image background*)	The width of the background image as pixels/inch.
Height (*Only with Image background*)	The height the background image as pixels/inch.

4. Click **OK**.

What's Next?

▶ To continue defining a new diagram, proceed to Adding Elements to Diagrams (page 199).

Adding Elements to Diagrams

This section instructs you on how to add new elements to diagrams both manually and automatically.

What's Next?

- ▶ If you want to insert new elements to Diagrams manually, proceed to Inserting New Elements Manually.
- ▶ If you want to create Diagrams from Configured Elements manually or automatically, proceed to Creating Diagrams from Configured Elements.

Inserting New Elements Manually

▼ To insert a new element to a diagram manually

1. Select the type of element you want to create in the toolbar and click the diagram in the place where you want to add the element. The new element appears grayed-out because it has not been configured yet.
2. Double-click the new element. The Properties dialog for that element opens.
3. Define the element's properties. For more information, see Defining IP Addresses (page 719).
4. Click **OK**. The newly configured element is no longer grayed-out, and it is added to the appropriate category in the All Elements list.

What's Next?

- ▶ To change the diagram layout, see Arranging Elements in Diagrams (page 201).

Creating Diagrams from Configured Elements

In addition to creating IP Diagrams, VPN Diagrams, and Connectivity Diagrams manually from configured elements, you can also generate all types of diagrams automatically from the elements that are monitored in the Management Center. The monitored elements are displayed in the System Status view.

Automatically generated IP Diagrams and VPN Diagrams are based on routing information, so they show elements that belong to the same routing configuration. Automatically generated Connectivity Diagrams contain elements that belong to the same configuration as defined in server, Firewall, IPS engine, and Layer 2 Firewall properties.

▼ To create a diagram from configured elements

1. Right-click **Diagrams** and select **New→Connectivity Diagram**, **New→VPN Diagram**, or **New→IP Diagram**. A blank diagram opens.
2. Browse to the type of element you want to add to the diagram.
3. *(In automatic diagram creation)* Select an element that is monitored from the list in the Resources panel and drag and drop it in the diagram.

4. Add the elements to the diagram.
   - *(In automatic diagram creation)* Right-click the element's icon and select **Auto Generate→Add from Configuration** or **Auto Generate→Add from Routing**. The elements that belong to the same configuration or routing configuration are added to the diagram. The links between the elements are also added automatically.
   - *(In manual diagram creation)* Drag the existing elements from the list in the Resources panel and drop them in the diagram.
5. Click the **Save** button. The Network Diagram Properties dialog opens.
6. Name the diagram and add a free-form comment for your own reference *(optional)*.
7. (*Optional*) Select a parent diagram for the new diagram.
8. Click **OK**.

You can edit the automatically generated diagram in the same way as any diagram. For example, you can add more elements to the diagram by dragging and dropping configured elements from the list in the Resources panel. This may be useful for documenting your system although the elements may not belong to the same configuration.

What's Next?

▶ To change the diagram layout, see Arranging Elements in Diagrams (page 201).

Adding Text Comments to a Diagram

▼ To add text comments to a diagram

1. Click the **Text Tool** button in the toolbar.
2. Click the diagram in the place where you want to add the comment. The Text Properties dialog opens.
3. Select the desired font properties and type in the comment.
4. Click **OK**. The comment appears in the diagram.
5. If necessary, select the comment and drag it in the diagram to change its position.

What's Next?

▶ To change the diagram layout, see Arranging Elements in Diagrams (page 201).

Arranging Elements in Diagrams

Prerequisites: Creating Diagrams

To move the elements, click the **Hand Tool** button in the toolbar and drag and drop the elements to their places. In addition to moving individual elements by dragging and dropping, you can also use the Hand Tool to move all the elements in a diagram.

Alternatively, you can change the layout of your diagram by selecting one of the automatic layout options.

▼ To change the layout automatically

1. Click the **Select Components** tool (arrow button) in the toolbar.
2. Select the elements you want to arrange.
3. Select **Tools**→**Layout** and one of the options. The selected elements in the diagram are arranged in the layout you chose.

Connecting Elements in Diagrams

Prerequisites: Creating Diagrams

After you add elements to a diagram, you can connect the elements automatically or manually. The automatic connections are based on the IP addresses of the elements or on the Routing configuration.

To help you maintain the clarity of your diagrams, you can also use turning points. Turning points are useful, for example, in complex diagrams where connections would otherwise cross each other.

A turning point is created by clicking the **Add Turning Point** button in the toolbar, clicking in the middle of the connection between the two elements and dragging the turning point to the direction of your choice to create an angle.

Connecting Elements Automatically

▼ To connect elements automatically in an IP Diagram

1. Click the **Select Components** button in the toolbar.
2. Select the elements you want to connect.
3. Select **Tools**→**Auto Connect by Routing** or **Tools**→**Auto Connect by Addresses** from the menu**.** The connections between the elements appear.

Connecting Elements Manually

▼ To connect elements manually

1. Click the **Connection Tool** button in the toolbar.
2. Click the first element to be linked.
3. Drag the pointer to the second element, and release the left mouse button. The elements are connected unless the connection is not allowed. You may have to insert a Network element in between two elements before you can connect them.

If the connection has a red, crossed-out circle on it, the connection is not valid. This can happen, for example, if you connect an element to a network with an address range that does not include that element's IP address. Such connections are allowed so that you can plan a configuration change without actually editing the elements until you are finished.

Creating Links Between Diagrams

Prerequisites: Creating Diagrams

You can specify a parent diagram to organize the diagrams into a hierarchy. You can also create links between diagrams that share the same parent diagram.

What's Next?

- To define the diagram hierarchy, see Specifying a Parent Diagram.
- To create links between diagrams, see Creating Links from One Diagram to Another (page 203).

Specifying a Parent Diagram

Parent diagrams allow you to organize diagrams in a hierarchy. For each diagram, you can specify another diagram as its parent diagram. All diagrams that have the same parent diagram are shown as branches below the parent diagram in the tree.

▼ To specify a parent diagram

1. Right-click the diagram for which you want to specify a parent diagram and select **Properties**. The Properties dialog opens.
2. Click **Select**. The Select Element dialog opens.
3. Select a parent diagram from the diagram list and click **Select**.
4. Click **OK**.

Your diagram is now displayed as a branch under the selected parent diagram in the tree

What's Next?

- Creating Links from One Diagram to Another (page 203)

Creating Links from One Diagram to Another

You can attach links to any of the elements in the diagram. When you double-click a link, another diagram opens. When you have a large network, linking allows you to use just a network icon to represent individual network segments in a top-level diagram, and then create links from them to more detailed diagrams.

Before you can add a link, you must have at least two diagrams under the same parent diagram.

Note – You cannot create links between two diagrams that have different parent diagrams.

▼ To add a diagram link to an element

1. Right-click the element you want to function as a link, and select **Link To**. The Select Link Target dialog opens.
2. Select the diagram to link to and click **Select**.
3. Click the **Save** button in the toolbar to save the diagram.

Viewing Diagrams

▼ To open a diagram for viewing

↳ Right-click the diagram you want to open and select **Preview Connectivity Diagram**, **Preview VPN Diagram** or **Preview IP Diagram**. The selected diagram opens for viewing.

Adjusting the Element Details in Diagrams

The details of the elements (name, IP address, type, etc.) can be shown as text below the elements in the diagram and as tooltips that only show when you place the mouse pointer over an element in a diagram.

▼ To select details shown below the elements

1. Select **View→Diagram Attribute Level** from the menu.
2. Select one of the items to adjust which details (No Labels, Names Only, Main Details, All Details) are shown below each element in the diagram.

▼ To select details shown as tooltips

1. Select **View→Tooltip Attribute Level** from the menu.
2. Select one of the items to adjust which details (No Tooltips, Names Only, Main Details, All Details) are shown as tooltips for each element when you place the mouse pointer on them in the diagram.

Collapsing and Expanding Groups of Elements in Diagrams

You can collapse and expand clustered engine elements (a Firewall Cluster, an IPS Cluster, or a Layer 2 Firewall Cluster) or groups of elements (any group of elements you select) to hide the individual nodes or elements. This can make a complicated diagram much simpler, as the individual items can be viewed only when needed. This can be done while viewing as well as editing diagrams.

▼ To collapse elements

1. Click the **Select Components** button (arrow button) and select two or more elements, or at least one of the nodes belonging to an expanded cluster.
2. Click the **Collapse Selected Component(s)** button in the toolbar. If several elements were selected, they are now shown as a single cloud with the tooltip Collapsed Area. Clustered elements are collapsed into an appropriate cluster icon.

▼ To expand elements

1. Select a clustered element or a group of collapsed elements by clicking its icon.
2. Click the **Expand Selected Component(s)** button in the toolbar. The individual nodes or elements are displayed. The nodes are shown with a yellow sign containing a minus sign to illustrate their association to an expanded cluster.

Zooming and Navigating Diagrams

In addition to the various tools for zooming the diagram (Zoom In, Zoom Out, etc.), you can also use the Diagram Navigation tool which offers you a quick way to zoom into selected parts of the diagram.

▼ To use the Diagram Navigation tool

1. If the Diagram Navigation panel is not displayed, select **View**→ **Diagram Navigation** to activate the Diagram Navigation tool.
2. Click and drag in the area outside the black rectangle to zoom in and out.
3. Click inside the black rectangle to adjust which part of the diagram is shown.
4. Select **View**→**Zoom**→**Default Zoom** to see the default zoom view again.

Printing Diagrams

You can directly print the diagrams with your printer.

▼ To print a diagram

1. Open the diagram you want to print (see Viewing Diagrams (page 203)).
2. (*Optional*) Select **File**→**Print Preview**. A preview opens.
3. (*Optional*) Select the print area by dragging the white box to the part of the diagram you want to print.
4. *(Optional)* Adjust the print area size by dragging the border of the white box.
5. Click the **Print** button in the toolbar.

Exporting Diagrams as Images

You can export the diagrams that you have created. You can save diagrams in various graphics file formats or as a PDF document.

▼ To export a diagram

1. Open the diagram you want to export (see Viewing Diagrams (page 203)).
2. Select **File**→**Export**. The Export dialog opens.

Note – To export the diagram as an image, select the Export item at the main level of the File menu. Exporting through the Export sub-menu creates an XML export file.

3. Select the folder into which you want to save the diagram.
4. Enter a **File name**.
5. Select the format from the **Files of Type** list.
6. Click **Export**. The diagram is saved with the file name and format you specified.

Chapter 13

Incident Cases

When suspicious activity is detected, it is important to collect all the information about the incident and act quickly. The *Incident Case* element is a tool for investigating incidents of suspicious activity.

The following sections are included:

Getting Started with Incident Cases

Prerequisites: None

What Incident Cases Do

The Incident Case element allows you to gather together all the data, actions, system configuration information, and files related to a specific incident. This information is useful for effective incident investigation, and also helps to provide evidence in the case of legal action.

Limitations of Incident Cases

When you write a Memo in an incident case, text does not automatically wrap in the Incident Memo Properties dialog. You must press Enter at the end of each line if you want your text to break to a new line.

Memos are permanently deleted when you remove them from an incident case. Other types of attachments can be added again if you remove them by mistake. See Attaching Data to Incident Cases (page 210).

Configuration Overview

Illustration 13.1 Elements in the Configuration of Incident Cases

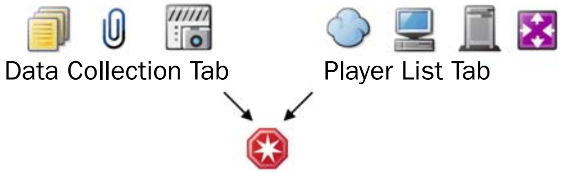

1. Create an incident case. See Creating a New Incident Case (page 209).
2. Attach relevant Logs, Policy Snapshots, Memos and files to the incident case. See Attaching Data to Incident Cases (page 210).
3. Add the network elements involved in the incident to the incident case. See Adding Players to Incident Cases (page 213).
4. Add journal entries to the incident case. See Adding Journal Entries to Incident Cases (page 213).

What's Next?

▶ To begin the configuration, proceed to Creating a New Incident Case (page 209).

Creating a New Incident Case

Prerequisites: None

▼ To create a new incident case

1. Select **Configuration→Configuration→Monitoring**.
2. Right-click **Incident Cases** and select **New Incident Case**. The Incident Case Properties dialog opens.

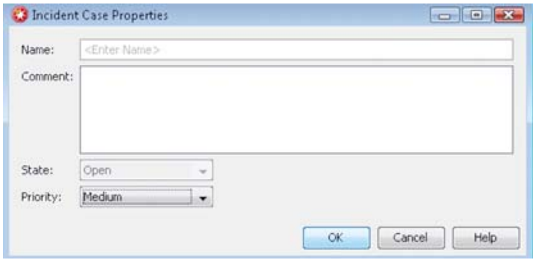

State: Open

Priority: Medium

3. Enter a unique **Name**, (*optional*) **Comment**, and *(optional)* **Priority**.
   - Priority information is for your own reference, and is not used by the system.
   - The **State** cannot be changed during the creation of the new incident case. To change the state, see Changing the State of an Incident Case (page 214).
4. Click **OK**. The new incident case opens for editing.

What's Next?

- If you want to set a context for handling the Incident Case in the Management Client, continue by Setting an Incident Context.
- Attach Logs, Policy Snapshots, Memos, and Files to the Incident Case as instructed in Attaching Data to Incident Cases (page 210).
- Add Players to the Incident Case as instructed in Adding Players to Incident Cases (page 213).
- Write Journal entries as instructed in Adding Journal Entries to Incident Cases (page 213).

Setting an Incident Context

Prerequisites: None

You can set a context for solving a particular Incident Case in the Management Client. The system then generates a list of your actions in the Management Client and attaches the selected Incident Case in the audited actions. You can view the actions on the History tab of the Incident Case.

▼ To set the Incident Context

1. Select **View→Layout→Incident Context Toolbar**. The Incident Context selection appears in the toolbar.
2. Select the Incident Case to which you want to add information. The status bar switches to orange when an Incident Case is selected.

What's Next?

- Attach Logs, Policy Snapshots, Memos, and files to the Incident Case as instructed in Attaching Data to Incident Cases.
- Add Players to the Incident Case as instructed in Adding Players to Incident Cases (page 213).
- Write Journal entries as instructed in Adding Journal Entries to Incident Cases (page 213).

Attaching Data to Incident Cases

Prerequisites: Creating a New Incident Case /Opening an Incident Case for Editing

You can use the Data Collection tab to attach information that is needed to provide context for investigating the incident.

The types of data that can be attached to an incident case are Logs, Policy Snapshots, Memos, and Files.

What's Next?

- To attach Logs, see Attaching Logs and Audit Entries to Incident Cases.
- To attach Policy Snapshots, see Attaching Policy Snapshots to Incident Cases (page 211).
- To write and attach Memos, see Attaching Memos to Incident Cases (page 212).
- To attach other types of files, see Attaching Files to Incident Cases (page 212).

Attaching Logs and Audit Entries to Incident Cases

You can attach selected single logs or audit entries, selected groups of logs or audit entries, or all logs or audit entries that match your query criteria to an incident case. It is also possible to attach log and audit data to incident cases directly from the Logs view.

▼ To attach logs or audit entries

1. Select **File→Attach→Logs** in the Data Collection tab of the incident case. The Logs view opens.
2. (*Optional*) To restrict the data that is displayed, specify your search criteria in the Query panel. See Filtering Logs in the Logs View (page 156) for details.
3. To attach a single log or a selected group of logs, select the log(s) you want to attach, right-click the selection and select **Export→Attach Log Events to Incident**. To attach all logs that match your filtering criteria, right-click any log entry and select **Export→Attach Log Events to Incident**. The Attach Logs to Incident Case dialog opens.

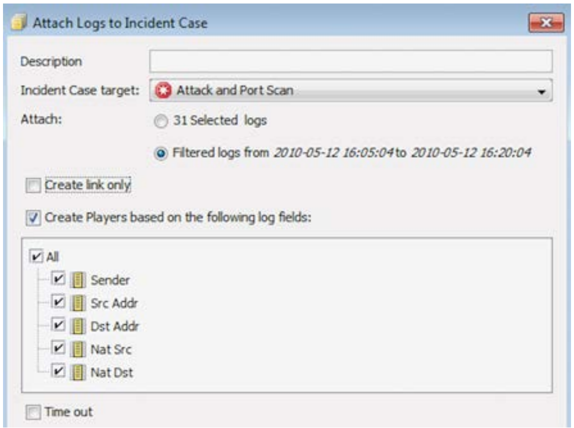

Attach Logs to Incident Case

Description

Incident Case target: Attack and Port Scan

Attach: ○ 31 Selected logs

● Filtered logs from *2010-05-12 16:05:04* to *2010-05-12 16:20:04*

□ Create link only

☑ Create Players based on the following log fields:

- ☑ All
  - ☑ Sender
  - ☑ Src Addr
  - ☑ Dst Addr
  - ☑ Nat Src
  - ☑ Nat Dst

□ Time out

4. Enter a unique **Description** for the log(s). This is used as the name of the data item in the Incident Case.
5. Select **Other** from the **Incident Case Target** pull-down list. The Select dialog opens.
6. Select the Incident Case that you want to attach the logs to and click **Select**.
7. Specify which logs you want to **Attach**:
   - Select **Selected logs** to add only the specific log entries you selected.
   - Select **Filtered logs from** to add all the log entries that match your query criteria.
8. (*Optional*) Select **Create Link Only** if you want to create a reference to the logs without duplicating the log files.
9. (*Optional*) Select **Create also Players based on the following log fields** and select the fields if you want to automatically fetch the player information related to these logs.
10. (*Optional*) Select **Time out** to automatically stop the export if it takes too long.
11. Click **OK**. The selected log(s) are attached to the incident case and appear in the Data Collection tab.

What's Next?

▶ If you are finished working with data attachments, proceed to Adding Players to Incident Cases (page 213) or Adding Journal Entries to Incident Cases (page 213).

Attaching Policy Snapshots to Incident Cases

Policy Snapshots help to establish which policies were in place at the time of the incident.

▼ To attach a policy snapshot

1. Select **File→Attach→Policy Snapshot** in the Data Collection tab. The Select Policy Snapshot dialog opens.
2. Select the policy snapshot you want to attach and click **Select**. The policy snapshot is attached to the incident case and appears in the Data Collection tab.

What's Next?

- To write and attach Memos, see Attaching Memos to Incident Cases.
- To attach other types of files, see Attaching Files to Incident Cases (page 212).
- Otherwise, if you are finished working with data attachments, add players as instructed in Adding Players to Incident Cases (page 213) or add journal entries as instructed in Adding Journal Entries to Incident Cases (page 213).

Attaching Memos to Incident Cases

Memos allow you to take notes about the incident. You can copy and paste from the clipboard into a memo. Memos can be edited and deleted after they are added to the incident case. If you want to add permanent, read-only comments and notes about the incident, see Adding Journal Entries to Incident Cases (page 213).

▼ To attach a memo

1. Select **File→Attach→Memo** in the Data Collection tab. The Incident Memo Properties dialog opens.
2. Give the memo a descriptive **Name**. This is used as the name of the data item in the incident case.
3. Type or paste the text of the memo in the **Text** field.
4. Click **OK**. The memo is added to the Data Collection list.

What's Next?

- To attach other types of files, see Attaching Files to Incident Cases (page 212).
- Otherwise, if you are finished working with data attachments, add players as instructed in Adding Players to Incident Cases (page 213) or add journal entries as instructed in Adding Journal Entries to Incident Cases (page 213).

Attaching Files to Incident Cases

You can attach any type of files, such as saved reports, text files, saved e-mail messages, packet capture files, or screen captures to incident cases. If you want to attach a newer version of a file, you must first remove the existing attachment from the incident case.

▼ To attach a file

1. Select **File→Attach→File** in the Data Collection tab. The Open dialog opens.
2. Browse to the file you want to attach. Select the file and click **Open**. The file is added to the Data Collection list.

What's Next?

- If you are finished working with data attachments, add players as instructed in Adding Players to Incident Cases (page 213) or add journal entries as instructed in Adding Journal Entries to Incident Cases (page 213).

Adding Players to Incident Cases

Prerequisites: Creating a New Incident Case /Opening an Incident Case for Editing

A player is any element that was involved in the incident. The related players can be fetched automatically when attaching logs or audit entries (see Attaching Logs and Audit Entries to Incident Cases (page 210)), or you can add players manually. Alternatively, you can copy and paste elements into the Player List tab of an incident case.

▼ To add a player

1. Select **File→New Player** in the Player List tab. The Player Properties dialog opens.
2. Give the player a unique **Name**. This name is used in the Player List.
3. Enter the **IPv4 Address** of the player or enter the **DNS Name** of the player and click **Resolve** to resolve the DNS name to an IP address. **Elements** shows a list of player elements.
4. *(Optional)* Add a **Comment**.
5. Click **OK**. The player is added to the Player list.

What's Next?

▶ If you are finished working with the player list, attach data as instructed in Attaching Data to Incident Cases (page 210) or add journal entries as instructed in Adding Journal Entries to Incident Cases.

Adding Journal Entries to Incident Cases

Prerequisites: Creating a New Incident Case

The journal allows you to record your observations and comments about administrator actions during the incident investigation. This is especially useful when more than one administrator is investigating the same incident simultaneously. Journal entries are automatically marked with a timestamp, and once a journal entry is added, it cannot be removed.

▼ To add a journal entry

➥ Type the journal entry in the **Additional Comment** field and click **Commit**. The journal entry is added to the list.

Working With Existing Incident Cases

Prerequisites: Creating a New Incident Case

After you have created an incident case, you can edit the contents and properties of the incident case.

Opening an Incident Case for Editing

▼ To open an incident case for editing

1. Select **Configuration→Configuration→Monitoring**.
2. Select **Incident Cases** in the tree view.
3. Right-click the incident case and select **Edit Incident Case**. The Incident Case view opens.

4. Edit the Incident Case as needed:
   - Attach data (see Attaching Data to Incident Cases (page 210)).
   - Add players (see Adding Players to Incident Cases (page 213)).
   - Write journal entries (see Adding Journal Entries to Incident Cases (page 213)).

Changing the Priority of an Incident Case

The default Priority of an incident case is **Low** when an incident case is created. The Priority information is only for your own reference, and it is not used by the system. As the investigation of the incident case progresses, you can change its priority as necessary.

▼ To change the priority of an Incident Case

1. Right-click the Incident Case you want to change the priority for and select **Properties**. The Properties dialog opens.
2. Select the new Priority and click **OK**.

Changing the State of an Incident Case

The default State of an incident case is **Open** when an incident case is created. The State information is only for your own reference, and it is not used by the system. As the investigation of the incident case progresses, you can change its state accordingly.

▼ To change the state of an Incident Case

1. Right-click the Incident Case you want to change the state for and select **Properties**. The Properties dialog opens.
2. Select the new **State**:
   - **Open**: The incident case has been created, but investigation has not begun.
   - **Under Investigation**: The incident case is actively being investigated.
   - **False Positive**: Legitimate activity was incorrectly interpreted as suspicious. There is actually no incident.
   - **Closed**: The investigation is finished.
3. Click **OK**. The state of the incident case is changed.

Checking Incident History

The Incident History tab shows all the logs and audit entries that track actions performed in this Incident Case view.

▼ To check the incident history

➥ Select **View→Incident History**. The Incident History tab opens.

CONTROLLING ENGINES

In this section:

CHAPTER 14

CONTROLLING ENGINE OPERATION

You can command and set options for Firewall engines, Layer 2 Firewall engines, IPS engines, Master Engines, and Virtual Firewalls through the Management Client.

The following sections are included:

Commanding Engines Remotely

Prerequisites: The engines must have a valid working configuration

You can control Firewall engines, Layer 2 Firewall engines, IPS engines, Master Engines, and Virtual Firewalls remotely through the Management Client through each engine element's right-click menu. The commands available depend on the type of component. In a cluster, the commands that affect the operating state of the engines can only be given to the individual nodes, not to the whole cluster.

You can also give commands and set options for more than one engine at a time by Shift-selecting or Ctrl-selecting the elements.

This subject is covered in the following topics:

- Turning Engines Online
- Turning Engines Offline (page 219)
- Setting Nodes to Standby (page 219)
- Rebooting Nodes (page 219)
- Refreshing the Currently Installed Policy (page 220)

Turning Engines Online

Engines in the offline state (status icon is blue and status text reads offline) can be turned online through the right-click menu if there are no configuration issues that would prevent the node/cluster from operating normally. Typical issues that may prevent a node from going online include policy issues, automatic tests failing, or heartbeat connection problems between nodes in clusters (see Node Does not Go or Stay Online (page 1102) for more information).

Note – You may also be able to give commands to nodes in the unknown state (gray icon), but you will not see a change of status. Since the actual operating status is not available, the node may already be online, in which case you will receive an error if you try to command the node online.

▼ To command a Security Engine to go Online

1. Select **Monitoring→System Status**.
2. Expand the tree until you see the individual engine nodes.
3. Right-click the offline engine node and select either of the following:
   - **Commands→Go Online**: the selected node goes to the online state and starts processing traffic according to the installed policy.
   - **Commands→Lock Online**: The node goes online and stays online even if a process tries to change its state. Serious error conditions can still force the node to go offline.
4. Click **OK**. The status is updated shortly.

Note – If the cluster is set to standby mode, only one node at a time can be online. Commanding a standby node online switches the current online node to standby. See Adjusting Firewall Clustering Options (page 556).

Turning Engines Offline

In the offline state, engines stop processing traffic, but remain otherwise operational and ready to be turned online again either automatically or by an administrator's command, depending on the configuration.

Caution – When you turn a node offline, it stops processing traffic. On Firewalls, traffic is stopped unless there are other cluster nodes that can take over.

▼ To command a Security Engine to go offline

1. Select **Monitoring→System Status**.
2. Expand the tree until you see the individual engine nodes.
3. Right-click the online or standby node and select either of the following:
   - **Commands→Go Offline**: the selected node goes offline and stops processing traffic.
   - **Commands→Lock Offline**: the node goes offline and stays offline even if a process tries to change its state.
4. Click **OK**. The status is updated shortly.

Setting Nodes to Standby

When a cluster runs in standby mode, only one node at a time processes traffic. The other running nodes are on standby. Standby nodes keep track of the traffic so that they can take over if the active node fails. Only one node at a time is in the online state, and the rest are either in the standby or offline state.

When you command an online node to standby, a standby node in the cluster (if there is one) automatically goes online to take over the traffic.

▼ To command a Security Engine to Standby

1. Select **Monitoring→System Status**.
2. Expand the tree until you see the individual engine nodes.
3. Right-click an Offline or Online node and select **Commands→Standby**. A confirmation dialog opens.
4. Click **OK**. The status is updated shortly.

Rebooting Nodes

Caution – If you are rebooting a cluster, reboot the nodes one by one to avoid breaks in service. If you command all nodes to reboot, all of the nodes reboot at the same time.

▼ To reboot a node

1. Select **Monitoring→System Status**.
2. Expand the tree until you see the individual engine nodes.
3. Right-click the node and select **Commands→Reboot**. A confirmation dialog opens.
4. Click **OK**. You can monitor the rebooting process by following the changes in the status of the element.

Refreshing the Currently Installed Policy

You can reinstall the currently installed policy of one or more components to transfer any configuration changes you have made in the system or in the policy itself since the last policy installation. Each type of Security Engine has its own type of policy. Inspection Policies can be used by all types of Security Engines.

Note – In clusters, all nodes must be either operational or explicitly disabled for the policy installation to succeed (see Disabling Nodes of a Cluster Temporarily (page 224)).

▼ Refreshing the currently installed policy

1. Select **Monitoring→System Status**.
2. Expand the tree until you see the cluster-level engine elements.
3. Right-click the single engine or engine cluster and select **Current Policy→Refresh**. The Task properties dialog opens.
4. Click **OK**.

Related Tasks

- ▶ Creating and Managing Policy Elements (page 645).

Backing up and Restoring Dynamic Routing Configurations

Dynamic routing is configured on the Firewall node's command line. See Configuring Dynamic Routing (page 236). You can back up and restore the engine's current dynamic routing configuration through the Management Client.

▼ To back up or restore a dynamic routing configuration

1. Select **Monitoring→System Status**.
2. Expand the tree until you see the individual engine nodes.
3. Right-click the node and select one of the following:
   - **Configuration→Dynamic Routing→Backup**: saves the node's current dynamic routing configuration.
   - **Configuration→Dynamic Routing→Restore**: restores a saved dynamic routing configuration on the node.

   A progress summary opens in a new tab.
4. Click **Close** when the backup/restoration is complete.

Removing Virtual Security Engines from a Master Engine

When you remove a Virtual Security Engine from a Master Engine, the Virtual Security Engine goes offline and stops processing traffic.

▼ To remove a Virtual Security Engine from a Master Engine

1. Select **Monitoring→System Status**.
2. Expand the tree until you see the cluster-level engine elements.

3. Right-click the Virtual Firewall and select **Commands→Remove Virtual Firewall from Master Engine**. A new tab opens to show the progress of the operation.
   - The Virtual Firewall is removed from the Master Engine. The Virtual Firewall element is kept in the system.
   - You can associate the Virtual Firewall with a different Virtual Resource to activate the Virtual Firewall on a different Master Engine. See Editing Virtual Security Engine Properties (page 410).

Moving Virtual Security Engines to Another Master Engine

You can move an active Virtual Security Engine from one Master Engine node to a different node in the same Master Engine. The Virtual Security Engine remains online and continues to process traffic during the move.

▼ To move a Virtual Security Engine to Another Master Engine

1. Select **Monitoring→System Status**.
2. Expand the tree until you see the cluster-level engine elements.
3. Right-click the Virtual Firewall and select **Commands→Move Virtual Firewall**. A Select Element dialog opens.
4. Select the Master Engine node to which you want to move the Virtual Firewall and click **Select**. The Virtual Firewall is moved to the selected Master Engine node.

Commanding Engines Locally

Prerequisites: None

Under normal circumstances, you should control the engines remotely through the Management Client. For abnormal situations, there are limited tools for giving some basic commands (such as go online/offline) through the engine’s command line interface:

- The available tools are listed in Engine Commands (page 1157).
- For information on how to access the engine command line, see Getting Started with the Engine Command Line (page 232).

Setting Engine Options

Prerequisites: None

Enabling/Disabling Engine Status Monitoring

By default, monitoring is automatically activated for all engines, but can be turned off as necessary.

▼ To disable/re-enable Security Engine monitoring

1. Select **Monitoring→System Status**.
2. Expand the tree until you see the top-level engine elements.
3. Right-click a single engine or engine cluster and select **Options**.
4. Deselect/select **Monitored**. Shortly, the status changes to Not Monitored and the icons associated with the element turn white.

Enabling/Disabling Firewall/VPN Diagnostics

Diagnostics mode provides more detailed log data for troubleshooting purposes. Diagnostics are only available for Firewalls, excluding Layer 2 Firewalls and SOHO Firewalls.

Note – Disable the diagnostics after troubleshooting to avoid overloading the Log Server with log data.

▼ To enable/disable diagnostics

1. Select **Monitoring→System Status**.
2. Expand the tree until you see the top-level engine elements.
3. Right-click the Firewall/Firewall Cluster element on which you want to apply diagnostics and select **Options→Diagnostics**. The Diagnostics dialog opens.
   - Click the top-level cluster icon, not the individual node(s) under the main icon.
4. Select/deselect the **Diagnostic** option at the top.
5. Select the diagnostic modes you want to apply to the selected Firewall/Firewall Cluster when diagnostics are enabled.
6. Click **OK**. The changes are applied immediately.

Disabling/Enabling User Database Replication

Firewall engines have a local replica of the Management Server’s internal LDAP database, which can store accounts for end-users for authentication purposes. By default, all changes are immediately replicated from the Management Server’s database to the local replicas on Firewall engines.

Master Engines have one local replica of the Management Server’s internal LDAP database for each Domain in which a Virtual Security Engine has users in the internal LDAP database. Changes are replicated from the Management Server’s database to the local replicas if Virtual Security Engines are associated with the Master Engine. The information that is replicated to Master Engines depends on what information is accessible to the Virtual Security Engines

You can disable the replication of the Management Server’s internal LDAP database to a Firewall engine or a Master Engine.

Note – Disabling or enabling replication of the Management Server’s internal LDAP database for a Master Engine also disables or enables replication of the Management Server’s internal LDAP database for all Virtual Security Engines hosted by the Master Engine.

▼ To disable/re-enable Firewall user database replication

1. Select **Monitoring→System Status**.
2. Expand the tree until you see the top-level engine elements.
3. Right-click a Single Firewall, Firewall Cluster, or Master Engine and select **Options**.
4. Deselect/select **User DB Replication**.

Enabling/Disabling Status Surveillance

By default, there is no warning to administrators if the status of the engines changes to an unknown state. You can optionally activate the status surveillance feature, which generates an alert if a single engine or none of the engines in a cluster send a status update for 15 minutes.

▼ To enable/disable engine status surveillance

1. Select **Monitoring→System Status**.
2. Expand the tree until you see the top-level engine elements.
3. Right-click a single engine or engine cluster and select **Options**.
4. Select/deselect **Status Surveillance**.

What's Next?

▶ Make sure System Alerts are escalated so that notification is sent out if status surveillance detects a failure. See Getting Started with Alert Escalation (page 266).

Enabling/Disabling SSH Access to the Engine

Secure remote access to the engines is provided by the SSH daemon process. This process can be started and stopped remotely. For maximum security, we recommend disabling SSH access whenever it is not used.

Alternatively, you can enable and disable SSH access when logged in to the node as explained in Reconfiguring Basic Engine Settings (page 233).

SSH uses TCP port 22. Make sure the connections are allowed in the policies of any Firewalls involved in the communications (including the Firewall that you are trying to contact).

▼ To enable/disable SSH access

1. Select **Monitoring→System Status**.
2. Expand the tree until you see the individual engine nodes.
3. Right-click a node and select **Commands→Enable SSH** or **Commands→Disable SSH**. The SSH process is started or stopped on the engine.

Changing the Engine Password

The password for access to the engine command line can be changed remotely through the Management Client as explained in this section. The user account for command line access is always `root`. Alternatively, if you remember the old password, you can change the password when logged in on the node as explained in Reconfiguring Basic Engine Settings (page 233).

▼ To change the command line password remotely

1. Select **Monitoring→System Status**.
2. Expand the tree until you see the individual engine nodes.
3. Right-click a node and select **Commands→Change Password**. The Change Password dialog opens.
4. Type in the new password in both fields and click **OK**. The new password is effective immediately.

Changing NetLink State Manually

Prerequisites: Creating NetLinks, Creating an Outbound Multi-Link Element

To change a NetLink's state, it must be operational. NetLinks with gray (unknown) status cannot be commanded.

▼ To change NetLink state manually

1. Select **Monitoring→System Status**. The System Status view opens.
2. Expand the **Multi-Link** branch and browse to the correct Firewall.
3. Right-click the NetLink whose state you want to change and select:
   - **Commands→Force NetLink Enable** to set the NetLink to a permanently active state regardless of probing results.
   - **Commands→Force NetLink Disable** to set the NetLink to a permanently disabled state regardless of probing results.
   - **Commands→Reset to Auto** to return the NetLink to automatic operation in which the NetLink is used based on probing results.

Disabling/Enabling Cluster Nodes

Prerequisites: None

Disabling Nodes of a Cluster Temporarily

Disabling a node is required to allow continued management of the other cluster members if one node goes out of operation. When you disable a node, you can physically remove it from the cluster without removing its definition from the system.

Disabling the node tells the other nodes and the Management Server that it is not necessary to try to contact it, preventing unnecessary communications attempts, alerts, test failures, etc. Disabling a node also allows policy installations on the other nodes when one node is shut down or malfunctions. No commands can be sent to a disabled node and no monitoring information is available for it.

▼ To disable a node in a cluster

1. Select **Monitoring→System Status**.
2. Expand the tree until you see the individual engine nodes.
3. If the node is online, right-click the node that you want to disable and select **Commands→Lock Offline**.
4. Double-click the cluster element. The properties dialog for the cluster opens.
5. Switch to the **Cluster** tab
6. Select the **Disabled** option in the Nodes table for the node(s) you want to disable and click **OK**.
7. Refresh the policy of the cluster.

When you re-enable the node, follow the procedure explained in Re-Enabling Disabled Cluster Nodes (page 225).

Re-Enabling Disabled Cluster Nodes

When a cluster node has been disabled, its configuration is typically made obsolete by policy installations done on the other cluster nodes. This prevents the node from operating normally and may in some cases disturb the operation of the whole cluster. Follow the procedure below to make the node operational.

▼ To enable a node in a cluster

1. (*Recommended*) Before connecting the disabled node (the same physical device or a replacement), set the node to the initial configuration state using the `sg-reconfigure` command on the engine command line. See Reconfiguring Basic Engine Settings (page 233).

Caution – If you re-introduce a disabled node that has a working configuration, the node must receive the heartbeat traffic from other nodes and accept it (based on certificates), or the node will consider itself the only available cluster member and go online. This may prevent the whole cluster from processing traffic.

2. Double-click the cluster element. The properties dialog for the cluster opens.
3. Switch to the **Cluster** tab.
4. Deselect the **Disabled** option in the Nodes table for the node(s) you want to re-enable and click **OK**.
5. Refresh the security policy to ensure that all nodes have the same configuration.
   - If the policy installation is unsuccessful, return the previously disabled node to the initial configuration state, see Reconfiguring Basic Engine Settings (page 233).
6. *(Optional)* Right-click the node and select **Commands→Go Online** or **Commands→Standby** to return the node to operation.

Editing Engine Configurations

Prerequisites: None

The engines can be configured in the following ways:

- The network card drivers, mapping of physical interfaces on the network cards to Interface IDs, and speed/duplex settings are defined using the Configuration Wizard on the engine command line. See Reconfiguring Basic Engine Settings (page 233).
- Other engine-specific settings are defined in the engine elements' properties in the Management Client. See Modifying Existing Engine Elements (page 384).
- The routing information is defined mostly in the Routing view (see Configuring Routing (page 587)). However, static ARP entries, multicast routing, and policy routing are defined in the engine element's properties. See Modifying Existing Engine Elements (page 384).

Chapter 15

Stopping Traffic Manually

The Access rules mainly determine which traffic is stopped, but you can also terminate or blacklist connections manually when you want the action to be temporary.

The following sections are included:

Terminating Connections Manually

Prerequisites: None

On Firewalls, Layer 2 Firewalls, and IPS engines with inline interfaces, you can terminate any current connection. For example, you can remove an inactive connection that has not been properly closed. Terminating an open connection alone does not prevent any new connection from opening again. This action is not available for IPS engines with capture interfaces.

▼ To terminate a connection manually

1. Select **Monitoring→Connections**.
2. Select a Security Engine from the list. The Connections view for the selected engine opens.
3. Select one or more connections in the table (use Shift-select or Ctrl-select to select more than one).
4. Right-click a selected row and select **Terminate**. For each selected connection, a confirmation message appears with information about the connection. When you click **Yes**, the engine terminates the connection.

Blacklisting Connections Manually

Prerequisites: Enabling Blacklist Enforcement

You can blacklist traffic manually on Firewalls, IPS engines, and Layer 2 Firewalls. For example, you can temporarily block a suspicious or disruptive source of communications while you conduct further investigations. You can create blacklist entries in the Blacklist view, Connections view, Monitoring view, and Logs view. The blacklist is not necessarily applied to all traffic; the Access rules determine how the blacklist is used.

Note – If a connection is allowed by a rule placed above the blacklist rule in the Access rules, the connection is allowed regardless of the blacklist entries. Check the logs to see which connections are discarded based on blacklisting.

▼ To create manual blacklist entry

1. Create a new blacklist entry in one of the following ways:
   - In the Blacklist view, Connections view, or Logs view: Right-click a row in the table and select **New Blacklist Entry** or **New Entry**.
   - To create a blacklist entry for a specific Security Engine: right-click the Security Engine element in the Connections view, Monitoring view, or Logs view, and select **New Blacklist Entry** or **Blacklist→New Entry**.

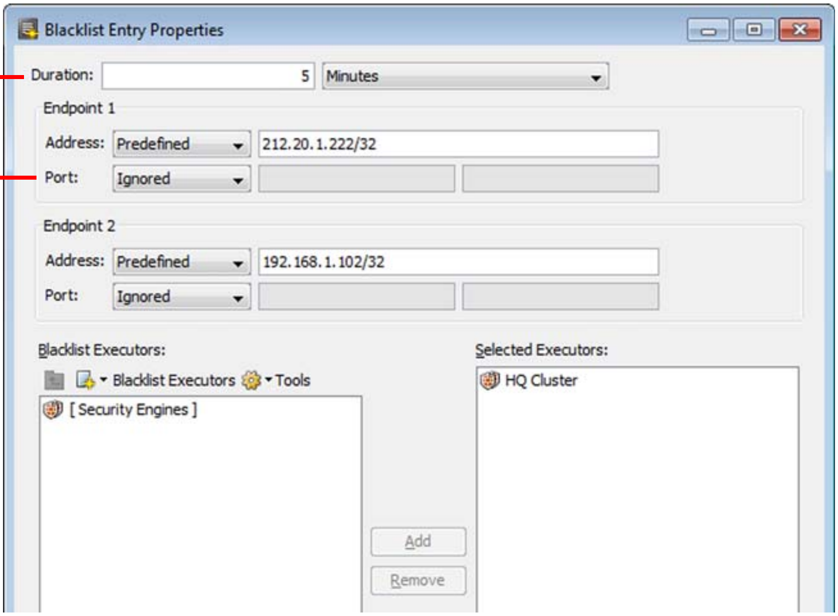

Blacklist Entry Properties

Duration: 5 Minutes

Endpoint 1

Address: Predefined 212.20.1.222/32

Port: Ignored

Endpoint 2

Address: Predefined 192.168.1.102/32

Port: Ignored

Blacklist Executors: [Security Engines]

Selected Executors: HQ Cluster

Leaving **Duration** as 0 cuts only the current connections.

Port ranges are available only if the protocol is TCP or UDP.

2. Select the **Duration** for how long this entry will be kept.
   - If you leave the value as 0, the entry only cuts the current connections. Otherwise, the entry is enforced for the specified period of time.
3. Select the **Address** to blacklist for Endpoint 1 and Endpoint 2.
   - **Any**: Matches any IP address.
   - **Predefined**: Matches the specific IP address and prefix you enter in the field to the right of the Address type list.
     - For example, the /24 prefix blacklists all the addresses in the same C-class network.
     - The default /32 prefix blacklists only the specific IP address you enter.
4. Select the **Port** to blacklist for Endpoint 1 and Endpoint 2.
   - **Ignored**: Matches any port.
   - **IPredefined TCP** and **Predefined UDP**: Matches the specific source and destination ports that you enter in the fields to the right of the Port type list.
5. Select the **Blacklist Executors** that enforce the blacklist entry.
6. Click **OK**. The blacklist entry is sent to the executor and the traffic is blocked.

Related Tasks

- Enabling Blacklist Enforcement (page 817).
- Configuring Automatic Blacklisting (page 818).
- Monitoring Connections, Blacklists, VPN SAs, Users, and Routing (page 111).
- Editing Access Rules (page 672).

CHAPTER 16

WORKING ON THE ENGINE COMMAND LINE

Although the engines are managed remotely, some operations on the Linux command line on the engines are useful for troubleshooting and local maintenance operations.

The following sections are included:

Getting Started with the Engine Command Line

Prerequisites: None

Nearly all aspects of engine configuration are done through the Management Client, but some engine settings and options must be defined and configured on the command line. Command line tools allow you to define and configure these settings and options.

What You Can Do on the Engine Command Line

- Reconfigure the engine's keyboard layout, time zone, network card settings, and network card to Interface ID mapping.
- Create scripts that run when the engine changes its state.
- (Re-)establish contact between the engine and the Management Server.
- Manually revert to the previous configuration.
- Run various general and Stonesoft-specific tools that assist in problem solving.

Limitations of the Engine Command Line

Changes made on the engine command line apply only to the node on which they were made. If you want to apply the changes to other engines, such as all the nodes in a cluster, you must make the same changes separately on the command line of each engine.

Some engine configuration options, such as network interface settings, cannot be changed through an SSH console. To be able to change these settings, you must connect using a serial cable or connect a display and keyboard directly to the engine machine.

The Management Server contact settings that are displayed in the Engine Configuration Wizard (sg-reconfigure) do not show the engine's actual working configuration (transferred whenever the engine's policy is installed/refreshed), but instead display the values that were set when the node was initialized.

What do I Need to Know Before I Begin?

All command line tools that are available for Single Firewalls are also available for Virtual Firewalls. However, there is no direct access to the command line of Virtual Firewalls. Commands to Virtual Firewalls must be sent from the command line of the Master Engine.

Accessing the Engine Command Line

Prerequisites: None

▼ To access the command line on the engines

1. Connect to the engine in one of the following ways:
   - Physically through a serial console using a null-modem cable (9600bps, 8 databits, 1 stopbit, no parity).
   - Physically using a display and keyboard connected directly to the engine machine.
   - Remotely using an SSH client. SSH access to the engine can be enabled and disabled through the Management Client as explained in Enabling/Disabling SSH Access to the Engine (page 223).
2. Log in as **`root`** and enter the engine password.
   - If you forget the password, you can change it in the Management Client as explained in Changing the Engine Password (page 223).

Reconfiguring Basic Engine Settings

Prerequisites: Accessing the Engine Command Line (Saving an Initial Configuration for Security Engines)

During the installation of the engine, the engine's keyboard layout (for command line use), time zone (for command line use), network card settings, and network card to Interface ID mapping were defined. The `sg-reconfigure` command allows you to change these settings. The procedure also allows you to re-establish a trust relationship between the engine and the Management Server if the trust is lost. On engines that are fully configured, each operation can be done individually without changing the other options.

▼ To reconfigure engine settings

1. Launch the Configuration Wizard using one of the following commands:
   - **`sg-reconfigure --no-shutdown`**: The Configuration Wizard starts without shutting the engine down. Network interface settings cannot be changed in this mode.
   - **`sg-reconfigure`**: The engine shuts down and the Configuration Wizard starts. All options are available if you have a local connection. If you have a remote SSH connection, you cannot change network interface settings (engine uses the "no-shutdown" mode).
2. Use the Configuration Wizard to change the settings. Detailed descriptions of the options can be found in the *Firewall/VPN Installation Guide* and the *IPS and Layer 2 Firewall Installation Guide*.
   - The Management Server contact details are not used by the engine after a policy has been installed from the Management Server. They are shown for your reference only.
   - If you select the **Switch to Initial Configuration** option, all configuration and policy information that has been transferred to the engine is cleared. You must install a policy on the engine before it can be operational again.

Caution – Do not select Switch to Initial Configuration on the Prepare for Management Contact screen unless you want to reset the engine to the post-installation state (including a policy that allows communication only between the engine and the Management Server).

   - If you select **Contact Management Server** and enter a new one-time password, the Management Server and the engine re-establish their trust relationship. Select this option when you want to replace a missing or expired certificate, or if the trust relationship with the Management Server is lost for any other reason, such as changing the Management Server's IP address.

Note – If there is a Firewall between a remote engine and the Management Server, you must allow the connection in the Firewall's IPv4 Access rules. If there is a NAT device between a remote engine and the Management Server, you must also configure NAT rules for the connection. Otherwise, the remote engine cannot contact the Management Server.

Related Tasks

- ▶ Engine Commands (page 1157).

Enabling 256-bit Security Strength for Engines

Prerequisites: Creating a New Internal ECDSA Certificate Authority, Saving an Initial Configuration for Security Engines, Accessing the Engine Command Line

You can optionally use 256-bit encryption for the connection between Security Engines and the Management Server. This requires both the engines and the Management Server to be version 5.5 or higher. You must also use an Internal ECDSA Certificate Authority to sign certificates for system communication. When you enable 256-bit security strength, the Management Server and the engine re-establish their trust relationship.

▼ To enable 256-bit Security Strength for engines

1. Launch the Configuration Wizard using one of the following commands:
   - **`sg-reconfigure --no-shutdown`**: The Configuration Wizard starts without shutting the engine down. Network interface settings cannot be changed in this mode.
   - **`sg-reconfigure`**: The engine shuts down and the Configuration Wizard starts. All options are available if you have a local connection. If you have a remote SSH connection, you cannot change network interface settings (engine uses the “no-shutdown” mode).
2. Select **Next** on each page until the Prepare for Management Contact page opens.

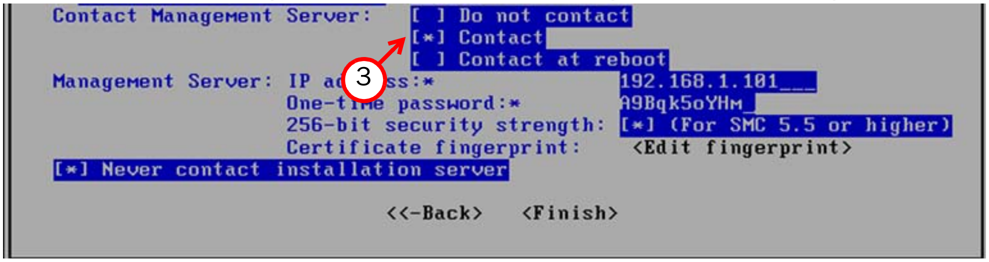

```
Contact Management Server:  [ ] Do not contact
                            [*] Contact
                            [ ] Contact at reboot
Management Server: IP address:*               192.168.1.101___
                   One-time password:*        A9Bqk5oYHm_
                   256-bit security strength: [*] (For SMC 5.5 or higher)
                   Certificate fingerprint:    <Edit fingerprint>
[*] Never contact installation server

                        <<-Back>   <Finish>
```

3. Select **Contact** or **Contact at Reboot** and press the spacebar.
4. Enter the Management Server IP address and the one-time password.

Note – The one-time password is engine-specific and can be used only for one initial connection to the Management Server. Once initial contact has been made, the engine receives a certificate from the Management Center for identification. If the certificate is deleted or expires, you need to repeat the initial contact using a new one-time password.

5. Select **256-bit Security Strength** and press the spacebar to use 256-bit encryption for the connection to the Management Server. 256-bit encryption must also be enabled for the Management Server. See Creating a New Internal ECDSA Certificate Authority (page 348).
6. (*Optional*) Highlight **Edit Fingerprint** and press **`Enter`**. Fill in the Management Server’s certificate fingerprint (also shown when you saved the initial configuration). Filling in the certificate fingerprint increases the security of the communications.
7. Highlight **Finish** and press **`Enter`**. The engine now tries to make initial Management Server contact. The progress is displayed on the command line.

Related Tasks

- ▶ Engine Commands (page 1157).

Creating Engine Scripts

Prerequisites: Accessing the Engine Command Line

Engine scripts run when the engine changes its state. The following scripts can be created:

Table 16.1 Possible Scripts on the Engines

Triggering Event	Script Location and Name
Engine boots up	`/data/run-at-boot`
Administrator refreshes or installs the policy	`/data/run-at-policy-apply`
Engine enters the Online state	`/data/run-at-online`
Administrator issued the 'Lock Online' command	`/data/run-at-locked-online`
Engine enters the Offline state	`/data/run-at-offline`
Administrator issued the 'Lock Offline' command	`/data/run-at-locked-offline`
Engine enters the Standby state	`/data/run-at-standby`

The script names and locations cannot be changed. If the scripts are not found, engine operation continues as normal. If a script is found, it is executed and a log entry is created. To stop this from happening, you must delete or move the script.

Note – If you want to use a script in a cluster, remember to create the script on all the nodes in the cluster or copy the script to all of them, so that all the nodes function in the same way when their state changes.

▼ To create a script

1. Create a text file with the commands you want the engine to execute (the first line of the script must be `#!/bin/sh`) in one of the following ways:
   - Create and edit the script on the engine's command line using the `vi` text editor.
   - Create and edit the script on a different host and transfer the file to the engine, for example, using SSH.
2. Save the script in the correct folder on the engine (see Table 16.1).
3. Make the file executable by typing the following command:
   `chmod a+x /data/<script name>`

The script is executed whenever the engine encounters the triggering event for running the script.

Related Tasks

- ▶ Engine Commands (page 1157).

Restoring a Previous Configuration Manually

Prerequisites: Accessing the Engine Command Line

If the engine loses management connectivity due to a faulty configuration, the previous configuration can be restored manually through the command line.

Note – You must restore the configurations separately on each node of a cluster. All nodes in a cluster must have the exact same configuration (as indicated by an identification code that is generated at each policy installation).

▼ To manually revert to the previous configuration

➥ Enter the following command:

```
/usr/lib/stonegate/bin/sgcfg -a -d /data/config/policy/previous apply
```

Related Tasks

▶ Engine Commands (page 1157).

Configuring Dynamic Routing

Prerequisites: Accessing the Engine Command Line

Dynamic routing with the Quagga suite is configured on the Firewall node's command line using the vtysh shell for Quagga routing software.

Note – Configuring dynamic routing for Virtual Firewalls requires SSH access to the command line of the Master Engine. We recommend that you disable SSH access whenever it is not needed and that you make sure your Access rules allow SSH access to the engines from the administrators' IP addresses only.

Detailed instructions for configuring the Quagga suite can be found on the Quagga Software Routing Suite web site at http://www.quagga.net/docs.php. For more information about other commands related to dynamic routing configuration, see Engine Commands (page 1157).

▼ To configure dynamic routing

1. Enter the following command to create an empty protocol-specific configuration file: `touch /data/config/quagga/ospfd.conf`.
2. Enter the following command to start the ospf routing daemon: `sg-dynamic-routing restart|force-reload`.

Caution – The protocol-specific configuration file (for example, ospfd.conf) must exist in the `/data/config/quagga` directory, and the routing daemon must be running before you configure dynamic routing using vtysh. Otherwise, vtysh cannot save protocol-specific configurations.

3. Enter the following command to start the Quagga command shell: `vtysh`.

Sending Commands to Virtual Security Engines

Prerequisites: Accessing the Engine Command Line

All command line tools that are available for Single Firewalls are also available for Virtual Firewalls. Commands to Virtual Firewalls are sent from the command line of the Master Engine. For more information about command line tools, see Engine Commands (page 1157).

▼ To send commands to Virtual Security Engines

1. Connect to the command line on the Master Engine as explained in Accessing the Engine Command Line (page 232).
2. Enter commands in the following format: **`se-virtual-engine [`*`options`*`]`**

Table 16.2 Options for `se-virtual-engine` Command

Option	Description
**`-h	--help`**	Shows the help message for the se-virtual-engine command.
**`-l	--list`**	Lists the active Virtual Security Engines.

`-v <`*`virtual engine ID`***`> |`**
`--virtual-engine=<`*`virtual engine ID`***`>`**
	Specifies the ID of the Virtual Security Engine on which to execute the command.

**`-e	--enter`**	Enters the command shell for the Virtual Security Engine specified with the `-v` or `--virtual-engine` option. To exit the command shell, type `exit`. Using the command shell is recommended if you need to send multiple commands to the Virtual Security Engine.

`-E "<`*`command [options]`***`>" |`**
`--execute="<`*`command [options]`***`>"`**
	Executes the specified command on the Virtual Security Engine specified with the `-v` or `--virtual-engine` option. Executing individual commands is recommended if you only need to send a few commands to the Virtual Security Engine. or if you need to send the same command to multiple Virtual Security Engines.

MANAGEMENT CENTER CONFIGURATION

In this section:

Chapter 17

Configuring Automatic Software Updates

You can configure the Management Server to automatically download and install dynamic update packages, remote upgrades for engines, and licenses.

The following sections are included:

Getting Started with Automatic Updates and Engine Upgrades

Prerequisites: None

Automatic updates are available for dynamic update packages, remote upgrades for engines, and licenses. On Firewalls, anti-virus signatures are always updated automatically when anti-virus is active.

What Automatic Updates and Engine Upgrades Do

The Management Server can automatically perform the following tasks:

- Check for new dynamic update packages and automatically download and install them according to your selection.
- Check for new engine upgrade packages. Engine upgrades can also be automatically downloaded, but they must always be installed manually.
- Upgrade the licenses.
- When automatic updates and engine upgrades are active, you can also view information regarding the maintenance contract and support level of your licenses in the Management Client (see Checking Maintenance Contract Information (page 129)).

Limitations

- *(Multiple Management Servers only)* Downloaded dynamic update packages are visible only on the Management Server to which they are imported. When the dynamic update packages are activated, the changes are made to the Management Database and then replicated between the Management Servers.
- There are no automatic updates for the Management Center software.
- New engine software versions may require an upgraded version of the Management Center. Check the Release Notes for compatibility information before upgrading the engines.
- Upgrades and updates (both automatic and manual) require an active maintenance/support contract with Stonesoft.
- If you select the **Notify when updates become available** setting, you must manually download the updates and engine upgrades.

Configuring Automatic Updates and Engine Upgrades

Prerequisites: None

The Management Server can periodically check the Stonesoft web site for new dynamic update packages, engine upgrades, and licenses. This feature is active by default. In an environment with multiple Management Servers, automatic updates must be enabled on the Management Server that controls the Shared Domain.

There are several options for how new updates and upgrades are handled. The automatic updates and engine upgrades require the Management Server to be able to connect to the Stonesoft servers at https://update.stonesoft.com and https://smc.stonesoft.com either using HTTPS on port 443 or through an HTTP proxy. Additionally, you must have a valid maintenance/support contract with Stonesoft.

▼ To configure automatic updates and engine upgrades

1. Right-click the Management Server element and select **Properties**. The Management Server Properties dialog opens.

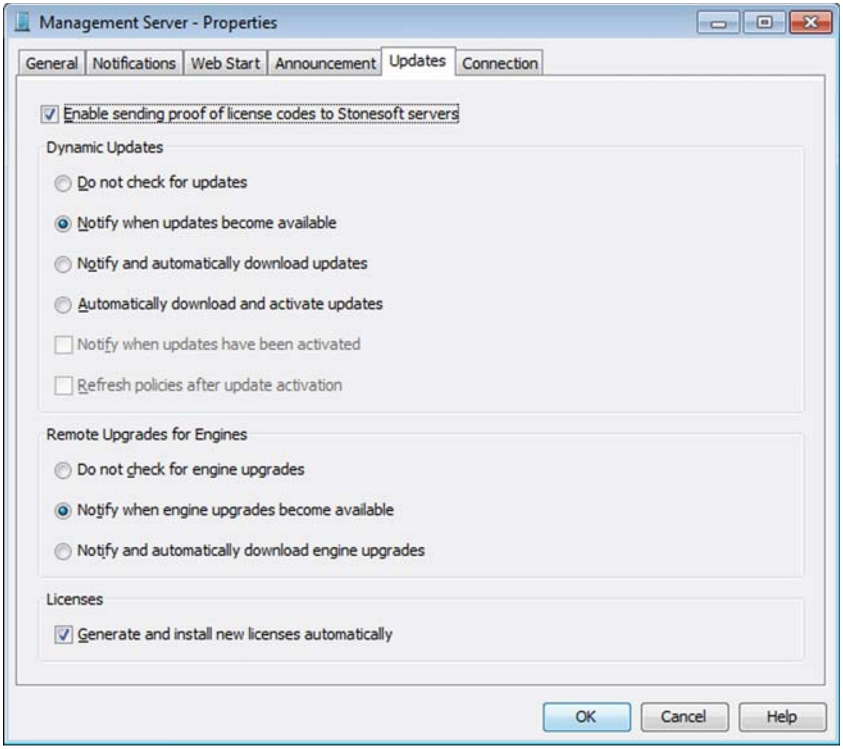

Management Server - Properties

☑ Enable sending proof of license codes to Stonesoft servers

Dynamic Updates

○ Do not check for updates
● Notify when updates become available
○ Notify and automatically download updates
○ Automatically download and activate updates
□ Notify when updates have been activated
□ Refresh policies after update activation

Remote Upgrades for Engines

○ Do not check for engine upgrades
● Notify when engine upgrades become available
○ Notify and automatically download engine upgrades

Licenses

☑ Generate and install new licenses automatically

2. Switch to the **Updates** tab.
3. Select **Enable sending proof of license codes to Stonesoft servers**. This allows you to select settings for dynamic updates and for engine and license upgrades.

4. Configure the **Dynamic Updates** settings:

Setting	Description
Do not check for updates	You are not notified of new dynamic updates.
Notify when updates become available	You receive an alert when a new dynamic update becomes available. You must manually download and activate the update.
Notify and automatically download updates	You receive an alert when a new dynamic update becomes available. The Management Center also automatically downloads the update. You must manually activate the update.
Automatically download and activate updates	The Management Center automatically downloads and activates the new dynamic updates. *(Multiple Management Servers)* This setting must be enabled on the Management Server that controls the Shared Domain.
Notify when updates have been activated	*(Optional)* You receive an alert when the dynamic updates have been activated. This option becomes available when you select **Automatically download and activate updates**. You must refresh the policies before the updates take effect. If **Refresh policies after update activation** is selected, the policies are refreshed automatically. Otherwise, you must refresh the policies manually.
Refresh policies after update activation	(*Optional*) The Management Center automatically refreshes the policies after activating the dynamic updates. This option becomes available when you select **Automatically download and activate updates**.

Note – Because update packages may change system elements, the policies may require editing after update activation.

5. Select one of the **Remote Upgrades for Engines** settings:

Setting	Description
Do not check for engine upgrades	You are not notified of new engine upgrades.
Notify when engine upgrades become available	You receive an alert when a new engine upgrade becomes available. You must manually download and install the update.
Notify and automatically download engine upgrades	You receive an alert when a new engine upgrade becomes available. The Management Center automatically downloads the new engine upgrade. You must manually install the update. *(Multiple Management Servers)* This setting must be enabled on the Management Server that controls the Shared Domain.

6. (*Optional*) Select **Generate and install new licenses automatically** to automatically regenerate and install the licenses required for upgrading Stonesoft system components to a major new release.

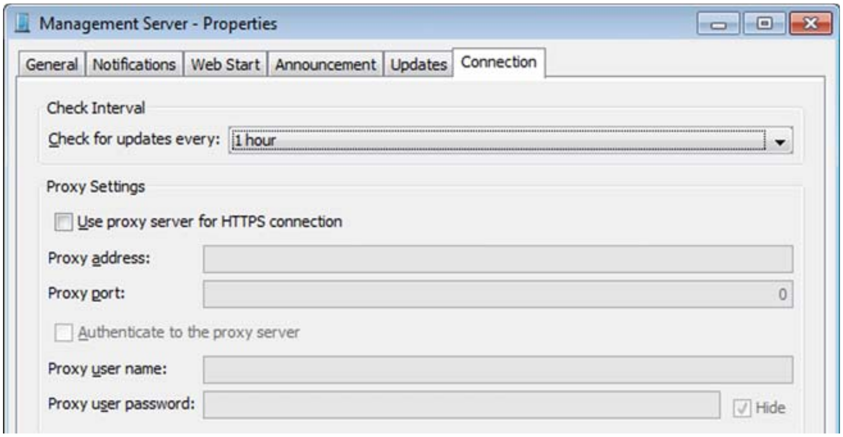

Check Interval

Check for updates every: 1 hour

Proxy Settings

□ Use proxy server for HTTPS connection

Proxy address:

Proxy port: 0

□ Authenticate to the proxy server

Proxy user name:

Proxy user password: ☑ Hide

7. (*Optional*) Switch to the **Connection** tab and select the **Check Interval** to define how often the system checks for new updates.

8. (*Optional*) If the connection from the Management Server to the Stonesoft servers requires a proxy server, select **Use proxy server for HTTPS connection** and enter the **Proxy Address** and **Proxy Port**.

   - If the proxy server requires user authentication, select **Authenticate to the proxy server** and enter the **Proxy User Name** and **Proxy Password**.

9. Click **OK**.

Related Tasks

- Checking Maintenance Contract Information (page 129).
- Scheduling Tasks (page 1040).
- Getting Started with Licenses (page 1044).
- Getting Started with Upgrading the SMC (page 1056).
- Getting Started with Upgrading Engines (page 1062).
- Manual Dynamic Updates (page 1069).

CHAPTER 18

ADMINISTRATOR ACCOUNTS

The rights and privileges of SMC administrators are defined with administrator accounts.

The following sections are included:

Getting Started with Administrator Accounts

Prerequisites: None

An administrator account specifies the actions for which the administrator has permissions (create new elements, browse logs, etc.).

How Administrator Accounts Can Be Configured

- Sets of administrator privileges are defined as reusable lists.
- Each list of privileges is applied to a specific group of elements.
- Several different pairs of privileges and elements can be applied to a single administrator account. These privileges can include, for example, viewing access to some elements and editing access to other elements.
- You can also create unrestricted accounts for “superusers” that can perform any action on any element. Some maintenance tasks require an unrestricted account.

What Do I Need to Know Before I Begin

Accounts that are used to log in to the Management Client can also be used to log in the Web Portal, but Web Portal User accounts are created separately. See Defining Web Portal User Accounts (page 298).

Configuration Overview

1. (*Optional*) Define customized reusable lists of allowed tasks for accounts with restricted permissions. See Defining Administrator Roles (page 249).
2. (*Optional*) Define customized reusable lists of elements for defining access rights for restricted accounts. See Defining Access Control Lists (page 252).
3. Create an administrator account for each administrator. See Defining Administrator Accounts (page 253).
4. (*Optional*) Configure strength requirements and expiration intervals for administrator passwords. See Defining Password and Login Settings for Administrators (page 259).

Caution – Do not use shared accounts. Using shared accounts makes auditing difficult and may make it difficult to discover security breaches.

Defining Administrator Roles and Access Control Lists

Prerequisites: None

You can use Administrator Roles and Access Control Lists in accounts that define restricted administrator permissions. You can either use the predefined Administrator Roles and Access Control Lists or create new ones.

What's Next?

- To customize the administrator permissions, start by Defining Administrator Roles.
- To define the elements the administrator is allowed to access, proceed to Defining Access Control Lists (page 252)

Defining Administrator Roles

Administrator Roles specify a restricted set of permissions that include, for example, the right to create, modify, and delete elements.

Caution – Any changes you make to an Administrator Role are applied immediately to every administrator account that uses the role (possibly including the account you are currently logged in with). Make sure that the permissions are correct before you apply changes to existing Administrator Roles.

To define an Administrator Role

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Right-click **Access Rights** and select **New→Administrator Role** or right-click an existing Administrator Role to edit and select **Properties**. The Administrator Role Properties dialog opens.

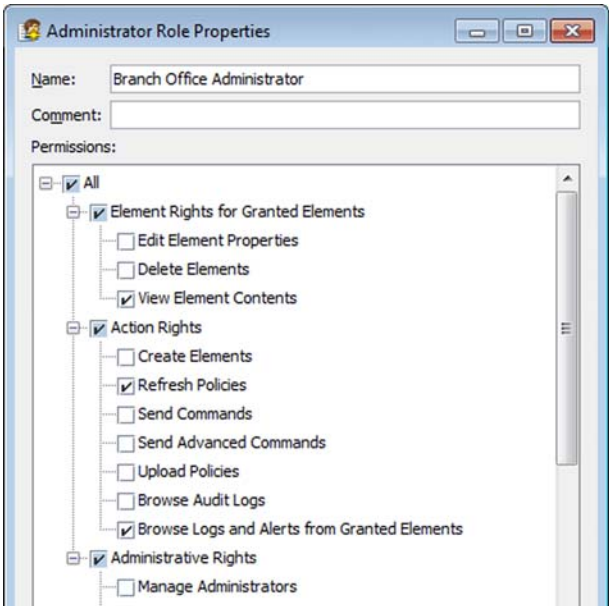

Administrator Role Properties

Name: Branch Office Administrator

Comment:

Permissions:

- ☑ All
  - ☑ Element Rights for Granted Elements
    - □ Edit Element Properties
    - □ Delete Elements
    - ☑ View Element Contents
  - ☑ Action Rights
    - □ Create Elements
    - ☑ Refresh Policies
    - □ Send Commands
    - □ Send Advanced Commands
    - □ Upload Policies
    - □ Browse Audit Logs
    - ☑ Browse Logs and Alerts from Granted Elements
  - ☑ Administrative Rights
    - □ Manage Administrators

3. *(New Administrator Role only)* Enter a unique **Name**.

4. Select the permissions that are applied to the elements selected for the role:

Table 18.1 Administrator Role Permissions

Permission	Description
Edit Element Properties	Allows editing the elements' properties.
Delete Elements	Allows deleting the elements.
View Element Contents	Allows viewing the contents of the elements.
Create Elements	Allows creating new elements.
Refresh Policies	Allows refreshing policies on engines if both the engines and the policies are selected for the role and the policies are allowed policies for the engines. It also allows creating and running Policy Refresh Tasks.
Send Commands	Allows sending basic commands to the selected engines. Basic commands allow an administrator to turn engines online and offline, reboot engines, and create and run sgInfo tasks and Remote Upgrade Tasks. If the role includes this permission and the permission to Browse Logs and Alerts from Granted Elements, the administrator can also terminate connections from the selected elements and browse and create blacklist entries for them.
Send Advanced Commands	Allows sending advanced commands to the selected engines. Advanced commands allow an administrator to enable and disable SSH access to the engine command line, change an engine password, and take Traffic Captures (see Taking a Traffic Capture (page 126)).
Upload Policies	Allows uploading policies on engines if both the engines and the policies are selected for the role and the policies are allowed policies for the engines. It also allows creating and running Policy Upload Tasks.
Browse Audit Logs	Allows browsing logs about administrator actions and events.
Browse Logs and Alerts from Granted Elements	Allows browsing logs from the selected elements and acknowledging alerts about them. If the role includes this permission and the permission to Send Commands, the administrator can terminate connections from the selected elements and browse and create blacklist entries for them.
Manage Administrators	Allows viewing and managing Administrator, Web Portal User, Administrator Role, and Access Control List elements.
Manage Users	Allows adding, removing, and modifying users in the Authentication Server, internal user database, and external directory servers.
Manage Alerts	Allows viewing and managing Alert, Alert Chain, and Alert Policy elements, and installing Alert Policies.
Manage Backups	Allows viewing and managing backups, and creating and running Backup Tasks.
Manage Licenses	Allows viewing, installing, binding, unbinding, and removing licenses.

Table 18.1 **Administrator Role Permissions (Continued)**

Permission	Description
Manage Logs	Allows creating and running log data tasks (Export Log Tasks, Archive Log Tasks, and Delete Log Tasks).
Manage Log Pruning	Allows pruning logs with Immediate Discard and Discard Before Storing filters.
Manage Reports	Allows viewing and managing reports.
Manage Updates and Upgrades	Allows downloading, importing, activating, and removing dynamic update packages. It also allows downloading, importing, and removing engine upgrades.
Manage VPNs	Allows viewing and managing elements related to VPNs.
View System Alerts	Allows browsing and acknowledging System Alerts, which are alerts about the internal operation of the system.

5. Click **OK**.

What's Next?

- ▶ If you want to define reusable groups for elements for which administrators have rights, proceed to Defining Access Control Lists (page 252).
- ▶ Otherwise, the new Administrator Role is now ready to be used when defining administrator accounts. See Defining Administrator Accounts (page 253).

Defining Access Control Lists

An Access Control List defines a group of elements for which an administrator has rights. There are several predefined Access Control Lists in **Administration→Access Rights→Access Control Lists**.

▼ To create an Access Control List

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Right-click **Access Rights** and select **New→Access Control List** or right-click an existing Access Control List and select **Properties**. The Access Control List Properties dialog opens.

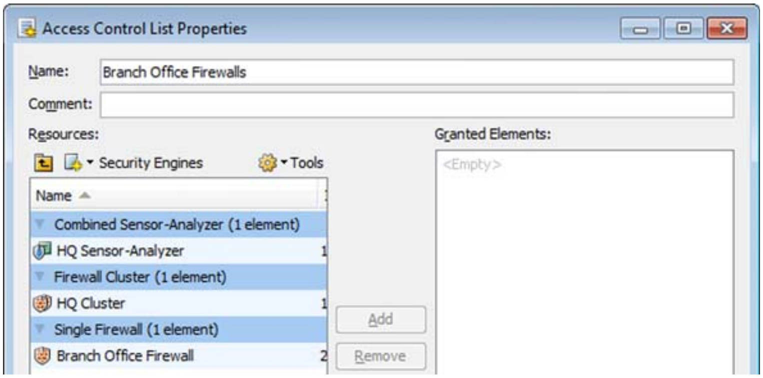

Name: Branch Office Firewalls

Name	
Combined Sensor-Analyzer (1 element)	
HQ Sensor-Analyzer	1
Firewall Cluster (1 element)	
HQ Cluster	1
Single Firewall (1 element)	
Branch Office Firewall	2

3. *(New Access Control List only)* Give the Access Control List a unique **Name** and enter an optional free-form **Comment** for your own reference.
4. Select the element(s) you want to add to the Access Control List from **Resources** and click **Add**. The selected element(s) are added to the **Granted Elements** list in the right panel.
5. Click **OK**.

What's Next?

- ▶ To define new administator accounts that use the new Access Control List, proceed to Defining Administrator Accounts (page 253).
- ▶ To assign the Access Control List to an existing restricted account, proceed to Defining Rights for Restricted Administrator Accounts (page 256).

Related Tasks

- ▶ Defining Administrator Permissions on Engines (page 530)

Defining Administrator Accounts

Prerequisites: (*Optional*) Defining Administrator Roles and Access Control Lists

An account with unrestricted permissions is automatically created during installation to ensure there is a “superuser” account available in the system. With this first account, you can create the necessary administrator accounts for daily management tasks.

The administrator accounts for the users of the optional Web Portal are defined with Web Portal User elements. See Defining Web Portal User Accounts (page 298). All other administrator accounts are defined with Administrator elements.

You can authenticate administrators using a password stored in the SMC’s internal database or you can use a RADIUS-based Authentication method provided by an external RADIUS server or by the Authentication Server. See Authenticating Administrators Using RADIUS Methods (page 262).

Creating a New Administrator Element

▼ To create an Administrator element

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Right-click **Access Rights** and select **New→Administrator**. The Administrator Properties dialog opens.

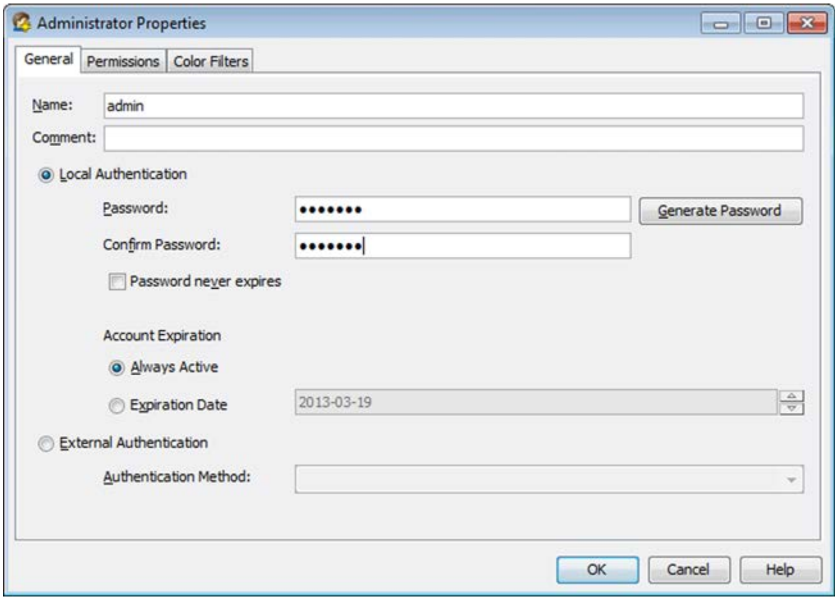

Name: admin

Comment:

● Local Authentication

Password: •••••••

Confirm Password: •••••••

□ Password never expires

Account Expiration

● Always Active

○ Expiration Date 2013-03-19

○ External Authentication

Authentication Method:

3. Give the administrator a unique **Name**. This is the username that the administrator uses to log in to the Management Client.
4. Select the type of authentication: **Local Authentication** (authentication by the Management Server) or **External Authentication** (RADIUS-based authentication by an external RADIUS server or by the Authentication Server).

5. Configure the authentication options according to the type of authentication:

Type of Authentication	Setting	Configuration
Local Authentication	Password	Enter and confirm the password, or click **Generate Password** to generate a random 7-digit alphanumeric password. **Note** - a generated password is a one-time password. The Administrator is prompted to enter a new password at the first login.
	Password never expires	(*Optional*) Select this option if you want the password to always be valid. Selecting this option overrides the password expiration settings in the Administrator Password Policy. See Defining Password and Login Settings for Administrators (page 259) for more information.
	Account Expiration	Select whether the account is Always Active or enter an Expiration Date.
External Authentication	Authentication Method	Select an Authentication Method provided by an external RADIUS authentication server or the Authentication Server component. See Authenticating Administrators Using RADIUS Methods (page 262) for more information.

Note – We recommend that passwords be at least eight characters long and contain a combination of numbers, letters, and special characters. Secure passwords are never based on personal information such as names, birthdays, social ID numbers, phone numbers, street names, or registration plate numbers.

What's Next?

▶ Continue by Defining Administrator Permissions (page 255).

Related Tasks

▶ Defining Password and Login Settings for Administrators (page 259)

Defining Administrator Permissions

▼ To define an administrator's permissions

1. Switch to the **Permissions** tab in the Administrator element's properties.

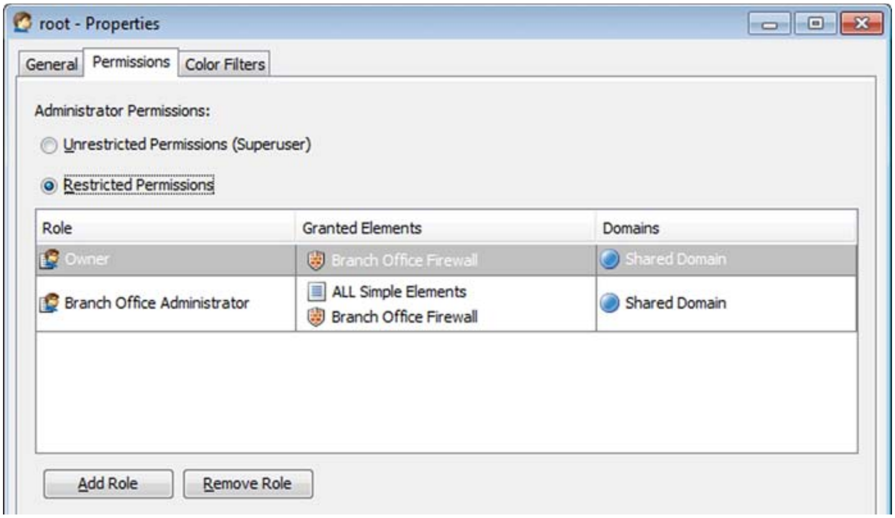

Administrator Permissions:

○ Unrestricted Permissions (Superuser)

● Restricted Permissions

Role	Granted Elements	Domains
Owner	Branch Office Firewall	Shared Domain
Branch Office Administrator	ALL Simple Elements Branch Office Firewall	Shared Domain

2. Select one of the following options:
   - **Unrestricted Permissions (Superuser)**: The administrator can manage all elements and perform all actions without any restrictions.
   - **Restricted Permissions**: The administrator has a limited set of rights that apply only to the elements granted to the administrator.

Caution – Select only the minimum necessary permissions for each Administrator account.

What's Next?

- If you selected **Restricted** permissions for this administrator, proceed to Defining Rights for Restricted Administrator Accounts (page 256).
- If you selected **Unrestricted Permissions** for this administrator and you want to limit the log entries viewable by this account, proceed to Restricting the Logs an Administrator Can View (page 257).
- If you selected **Unrestricted Permissions** for this administrator and you want to adjust the display colors of log entries for this account, proceed to Customizing Log Colors (page 258).
- Otherwise, click **OK**. The administrator account is ready for use.

Related Tasks

- Defining Administrator Permissions on Engines (page 530)

Defining Rights for Restricted Administrator Accounts

To define the permissions in detail as explained below, the Administrator element must have **Restricted Permissions** selected as the administrator permissions level.

▼ To define the administrator rights for a restricted account

1. Switch to the **Permissions** tab in the Administrator Properties dialog.

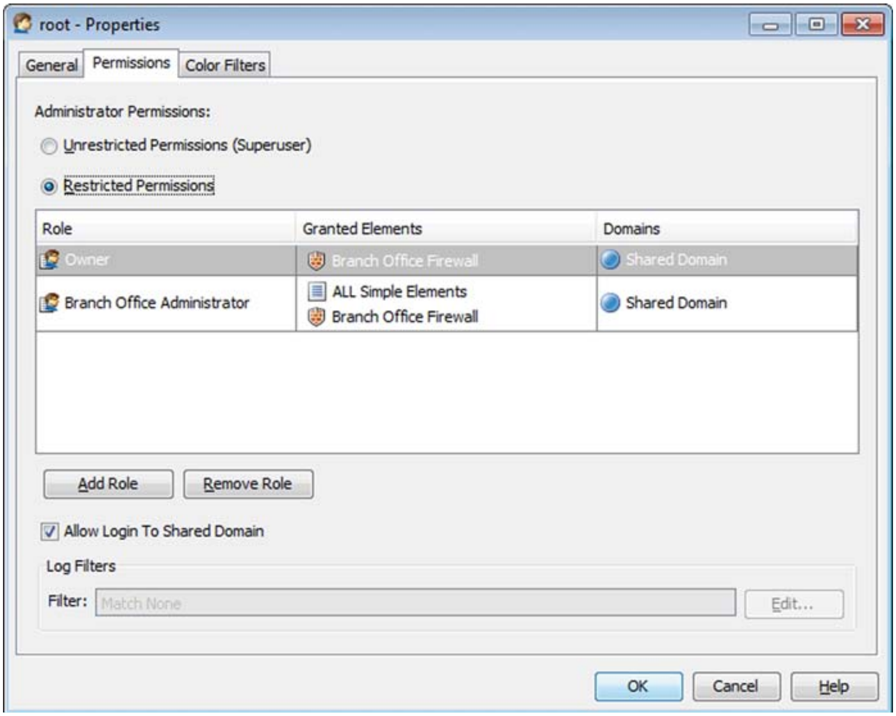

Administrator Permissions:

○ Unrestricted Permissions (Superuser)

● Restricted Permissions

Role	Granted Elements	Domains
Owner	Branch Office Firewall	Shared Domain
Branch Office Administrator	ALL Simple Elements Branch Office Firewall	Shared Domain

☑ Allow Login To Shared Domain

Log Filters

Filter: Match None

2. Click **Add Role**. A new Administrator Role appears in the list above.
3. Click the **Role** cell and select the Administrator Role that defines the rights you want to set. In addition to any customized roles, there are four predefined Administrator Roles:
   - **Operator:** Can view the properties of selected elements. Can send commands to selected engines, refresh and upload policies, and browse logs and alerts from selected elements.
   - **Editor**: Can create, edit, and delete selected elements. Can send commands to engines, refresh and upload policies, and browse logs and alerts from selected elements.
   - **Owner**: Can view the properties of selected elements, and edit and delete the selected elements.
   - **Viewer**: Can view the properties of selected elements.
4. Double-click the **Granted Elements** cell and select the elements to which the rights granted by the Administrator Role apply.
   - The **Set to ALL** action also depends on the type of elements. For example, if you browse to Firewalls and click **Set to ALL**, the item “All Firewalls” is added.
   - You can also select one or more predefined or user-created Access Control Lists. “Simple elements” includes all elements except elements that have a dedicated system Access Control List (for example, there are dedicated Access Control Lists for different types of security engines and their policies). See the *Management Center Reference Guide* for a description of all the pre-defined Access Control lists.

5. (*Optional*) If Domain elements have been configured, click the **Domain** cell to select the Domain(s) in which the rights granted by the Administrator Role and the selected elements apply.
   - You can leave the default **Shared Domain** selected in the Domains cell. All the elements automatically belong to the predefined Shared Domain if Domain elements have not been configured. See Domains (page 283) for more information on Domains.
   - You can also select the **ALL Domains** Access Control List to grant permissions for all the Domains that exist in the system.
6. (*Optional*) Repeat Steps 2 to 5 to define additional administrator permissions.
7. (*Optional*) If Domain elements have been configured, leave **Allow Login to Shared Domain** selected if you want to grant the administrator permission to log in to the Shared Domain. Otherwise, the administrator is only allowed to log in to the specified Domain(s).

In addition to the privileges you explicitly set, administrators who are allowed to create and delete elements automatically have privileges to view, edit, and delete elements they create themselves, even if they are not allowed to view, edit, or delete any other elements.

What's Next?

▶ If the administrator has the right to view logs, you can further restrict the logs the administrator can view as explained in Restricting the Logs an Administrator Can View.

▶ If you want to adjust the display colors of log entries for this account, proceed to Customizing Log Colors (page 258).

▶ Otherwise, click **OK**. The administrator account is ready for use.

Related Tasks

▶ Defining Administrator Permissions on Engines (page 530)

Restricting the Logs an Administrator Can View

If an administrator is allowed to view logs and alerts, you can create local filters that are applied to the log data before it is displayed to the administrator. The filters that you create here are specific only to the Administrator element in question, unless you save them as permanent Filter elements. See Saving Local Filters (page 189).

This section explains how you can select log filters in Administrator elements. See Selecting Log Browsing Permissions for a Web Portal User (page 301) for information on selecting log browsing permissions and filters for Web Portal users.

▼ To select log filters in an Administrator element

1. Switch to the **Permissions** tab in the Administrator element properties.
2. Under Log Filters, click **Edit**. The Local Filter Properties dialog opens. For information on how to create local filters, see Creating and Editing Local Filters (page 187).

What's Next?

▶ If you want to adjust the display colors of log entries individually for this account, proceed to Customizing Log Colors (page 258).

▶ Otherwise, click **OK**. The administrator account is ready for use.

Customizing Log Colors

Prerequisites: None

By default, certain logs are shown with a colored background in the Logs view. The colors are determined by Log Color Filters. You can customize the default log colors used by default in all administrator accounts or define administrator-specific log colors in the Administrator element's properties. To use customized colors for logs, you must also create the filter(s) that match those logs. Only administrators with the right to manage administrator accounts can customize log colors.

▼ To customize default log color filters

1. In the Administration Configuration view, right-click the Administrator whose log colors you want to change and select **Properties**. The Administrator Properties dialog opens.
2. Switch to the **Color Filters** tab.

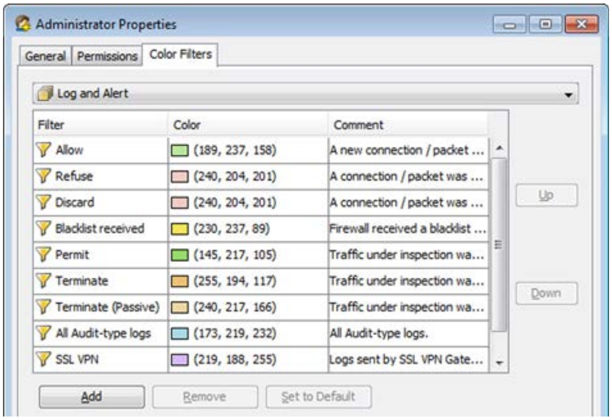

Filter	Color	Comment
Allow	🟩 (189, 237, 158)	A new connection / packet ...
Refuse	🟥 (240, 204, 201)	A connection / packet was ...
Discard	🟥 (240, 204, 201)	A connection / packet was ...
Blacklist received	🟨 (230, 237, 89)	Firewall received a blacklist ...
Permit	🟩 (145, 217, 105)	Traffic under inspection wa...
Terminate	🟧 (255, 194, 117)	Traffic under inspection wa...
Terminate (Passive)	🟨 (240, 217, 166)	Traffic under inspection wa...
All Audit-type logs	🟦 (173, 219, 232)	All Audit-type logs.
SSL VPN	🟪 (219, 188, 255)	Logs sent by SSL VPN Gate...

3. Select the log type for which you want to modify color filter(s):
   - **Log and Alert**: colors for logs and alerts displayed in the Logs view.
   - **Connections**: colors for currently open connections displayed in the Connections view.
   - **Blacklist**: colors for blacklist entries in the Blacklist view.
   - **VPN SAs**: colors for Internet Exchange Keys (IKE) and IPsec protocols displayed in the VPN SAs view.
   - **Users**: colors for different users in the Users view.
   - **Routing**: colors for routing entries displayed in the Routing Monitoring view.
4. (*Optional*) To define a new filter for a log type, click **Add** and double-click the **Empty Filter** cell in the new color filter row. The Local Filter Properties dialog opens. For information on how to create local filters, see Creating and Editing Local Filters (page 187).
5. Double-click the **Color** cell of the filter for which you want to modify the color and select a color from the palette or click **More Colors** to select a custom color. The selected colors are assigned to the filters and they are used whenever logs match the filter.
6. Click **OK**.

Defining Password and Login Settings for Administrators

Prerequisites: None

You can define settings for password strength, password expiration, failed logins, and actions related to temporary and long-term inactivity in the Administrator Password Policy. The Administrator Password Policy is applied to all administrator accounts defined with Administrator and Web Portal User elements. The Administrator Password Policy is not enforced by default.

> **Note – The Password Never Expires option in the Administrator element properties overrides the password expiration settings. If the option is selected, the password expiration settings do not apply to the account.**

Enabling Enforcement of Password Settings

▼ To enable enforcement of password settings

1. Select **File→System Tools→Password→Enforce Password Settings** from the main menu. A confirmation dialog opens showing the current password policy settings.

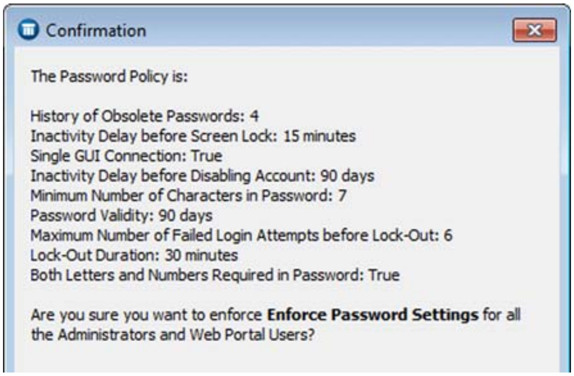

The Password Policy is:

History of Obsolete Passwords: 4
Inactivity Delay before Screen Lock: 15 minutes
Single GUI Connection: True
Inactivity Delay before Disabling Account: 90 days
Minimum Number of Characters in Password: 7
Password Validity: 90 days
Maximum Number of Failed Login Attempts before Lock-Out: 6
Lock-Out Duration: 30 minutes
Both Letters and Numbers Required in Password: True

Are you sure you want to enforce **Enforce Password Settings** for all the Administrators and Web Portal Users?

Table 18.2 Password Policy Settings

Setting	Description
History of Obsolete Passwords	The number of previous passwords that an administrator cannot reuse.
Inactivity Delay before Screen Lock	The number of minutes after which an administrator who is idle is automatically logged out.
Single GUI Connection	Defines whether an administrator can open only a single session at a time to the Management Client or to the Web Portal.
Inactivity Delay before Disabling Account	The maximum number of days an administrator account can be inactive before it is disabled.
Minimum Number of Characters in Password	The minimum number of characters an administrator password must contain.
Password Validity	The number of days after which administrator passwords expire and must be changed.

Table 18.2 Password Policy Settings (Continued)

Setting	Description
Maximum Number of Failed Login Attempts before Lock-Out	The maximum number of failed login attempts before an administrator account is locked.
Lock-Out Duration	The duration (in minutes) for which the administrator account is locked when the maximum number of failed login attempts is reached.
Both Letters and Numbers Required in Password	Defines whether administrator passwords must contain both letters and numbers.

2. Click **Yes**. The password policy settings are now applied to new administrator passwords.

What's Next?

- If you want to modify the current password policy settings, continue by Defining Password Policy Settings (page 260).
- Otherwise, the password policy configuration is complete.

Defining Password Policy Settings

The administrator password policy is defined in the `SGConfiguration.txt` file that is stored on the Management Server. You can either use the default values for each setting or modify the settings as necessary.

▼ To define administrator password policy

1. Exit the Management Client.
2. Stop the Management Server service through the operating system's service management feature or using the **`sgStopMgtSrv`** script.
3. Browse to the `<`*`installation directory`*`>/data` directory on the Management Server.

Note – If you installed the Management Server in the `C:\Program Files\Stonesoft\Management Center` **directory in Windows, some program data may be stored in the** `C:\ProgramData\Stonesoft\Management Center\data` **directory.**

4. Edit `SGConfiguration.txt` and configure the following parameters as needed:

Parameter Name	Description
`OBSOLETE_PASSWORD_QUEUE_LENGTH`	Enter a number to prevent the reuse of the specified number of previous passwords. The default number is 4.
`GUI_INACTIVITY_MAX_IN_MINUTES`	Enter the number of minutes after which administrators who are idle are logged out of the Management Client. The default is 15 minutes. Setting the parameter to zero minutes disables the screen lock for administrators.
`SINGLE_GUI_SESSION`	Define (as **True** or **False**) whether a single login per administrator account can be used to log in to the Management Client at a time. The default value is **True**. When the value is **True**, only a single login per account is allowed.
`DAYS_OF_ACCOUNT_INACTIVITY_MAX`	Enter the maximum number of days an administrator account can be inactive before it is disabled. The default is 90 days.
`PASSWORD_CHARACTER_NUMBER_MINIMUM`	Enter the minimum number of characters administrator passwords must contain. The default is 7 characters.
`PASSWORD_EXPIRATION`	Enter the number of days after which the administrator password must be changed. The default is 90 days.
`FAILED_AUTH_ATTEMPT_MAX`	Enter the maximum number of failed login attempts after which the administrator account is locked. The default is 6 attempts.
`TIME_LOGIN_LOCKED`	Enter the number of minutes for which the administrator account is locked when the maximum number of failed logins is reached. The default is 30 minutes. You must define a value for `FAILED_AUTH_ATTEMPT_MAX` to use this.
`PASSWORD_BOTH_LETTER_AND_NUM_REQUIRED`	Define (as **True** or **False**) if administrator passwords must contain both letters and numbers. The default is **True**.

5. Save the changes you have made to the file.
6. Restart the Management Server service.
7. Launch the Management Client.

Changing Administrator Passwords

Prerequisites: None

If you have not configured administrator passwords to automatically expire, we recommend that you change administrator passwords regularly. An administrator who has the right to manage administrator accounts can change any other administrator's password. All administrators can change their own passwords in the Management Client or the Web Portal.

▼ To change another administrator's password

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Select **Access Rights→Administrators** or **Access Rights→Web Portal Users**.
3. Right-click the Administrator or Web Portal User element and select **Properties**.
4. Enter and confirm the **Password**.
5. Click **OK**.

▼ To change your own administrator password

1. Select **File→System Tools→Password→Change Password**. The Change Password dialog opens.
2. Enter your current password in the **Old Password** field.
3. Enter a new password in the **New Password** field.
4. Confirm the new password in the **Confirm New Password** field.
5. Click **OK**.

Authenticating Administrators Using RADIUS Methods

Prerequisites: You must have an external authentication server that supports the RADIUS protocol

You can authenticate administrators and Web Portal users using RADIUS-based authentication methods provided by an external authentication server, Active Directory server, or the optional Authentication Server component. You can enable the use of RADIUS authentication methods separately for each administrator or Web Portal User account.

The Management Server's internal user database does not allow external authentication servers to query the administrator account information, so you must maintain accounts separately both in the SMC and on an external directory server that the external authentication server can access. The administrator name must be the same in both user databases.

If you are using the Authentication Server component and have linked Administrator or Web Portal User accounts on an external directory server to user accounts in the Authentication Server, only the PAP RADIUS method can be used. If you are using the Authentication Server component and have not linked user accounts, any RADIUS method can be used.

▼ To set up RADIUS-based authentication for administrators and Web Portal users

1. Define one of the following types of server elements to represent the external server:
   - Define a RADIUS Authentication Server element (see Defining RADIUS or TACACS+ Authentication Servers (page 858)).
   - Define an Active Directory Server element and add a RADIUS-based Authentication Method (see Defining Active Directory Server Elements (page 830)).
   - Install the optional Authentication Server component (see Integrating Authentication Server Services (page 861)).

Note – The shared secret used in the communications is defined in the RADIUS Authentication Server or Active Directory Server element, or in the RADIUS Client settings of the Authentication Server element.

2. Create a rule in your Policy allowing traffic from your Management Server to the external RADIUS server, or from your Management Server to the Authentication Server.
3. Right-click the Management Server and select **Properties**.
4. Select the **RADIUS Method** for the authentication method you want to use for authenticating the Management Server's communications with the external RADIUS server
   - **PAP** (Password Authentication Protocol): This is the only supported method for linked Authentication Server users.
   - **CHAP** (Challenge-Handshake Authentication Protocol).
   - **MSCHAP, MSCHAP 2**: Microsoft versions of the CHAP protocol. We recommend using MSCHAP 2 if it is supported by the RADIUS server.
   - **EAP-MD5**: Extensible Authentication Protocol with an MD5 Hash.

Caution – The security of the system requires that these communications remain confidential. Whichever security method you choose, we recommend transferring the communications between the Management Server and the RADIUS server over secure networks only.

5. (*RADIUS Authentication Servers only*) Set up the external server for use with the Management Server.
   - Define the Management Server as a RADIUS client on your server.
   - Define the same authentication method on your server as you selected in the Management Server properties in the previous step.
6. Select the **Authentication Method** for external authentication in the properties of each Administrator or Web Portal User account (as explained in Defining Administrator Accounts (page 253)). The Administrator element's name must be the same as in the user database that the external server uses.

Disabling Administrator Accounts

Prerequisites: Creating a New Administrator Element

If an administrator account is no longer needed, you can disable the administrator account to remove access for the administrator.

▼ To disable an administrator account

1. Select **Administration→Access Rights→Administrators**.
2. Right-click the Administrator element and select **Disable Administrator**. A confirmation dialog opens.
3. Click **Yes** to confirm that you want to disable the Administrator.
4. The Administrator is marked as **Disabled** and all scheduled Tasks created by the Administrator are ignored.

Deleting Administrator Accounts

When you delete an administrator account that has been marked as obsolete, all references to the administrator are removed, for example, from the element modification history. Audit entries that reference the administrator are preserved.

Note – There must be at least one account with unrestricted privileges in the system. It is not possible to delete the last remaining unrestricted account.

See Deleting Obsolete Elements (page 89) for general information on deleting elements that have been marked as obsolete.

CHAPTER 19

ALERT ESCALATION

The alerting system can be configured to escalate the alerts generated so that notifications are sent to the administrators through multiple channels.

The following sections are included:

Getting Started with Alert Escalation

Prerequisites: None

Alerts notify you in case something unexpected or suspicious happens. It is vital to the health of the system that alerts do not go unnoticed.

What Alert Escalation Does

The alert escalation process starts when an alert is triggered by a system warning or error, or because traffic matches a rule that is configured to trigger an alert. The Log Server can send out different types of notifications to any number of administrators. Alert escalation stops when one of the administrators acknowledges the alert or when all configured alert notifications have been sent.

Limitations

- Alert escalation for the SSL VPN is configured in the SSL VPN Administrator interface.
- Only one e-mail recipient can be configured for each notification. To send an e-mail to several people at the same time, you must configure an e-mail group on the mail server or configure several notifications consecutively without delays.
- Test alerts cannot be sent with Firewall, Layer 2 Firewall, or IPS engines as the sender, and they always have default Severity and Situation information.

Configuration Overview

Illustration 19.1 Elements in the Configuration

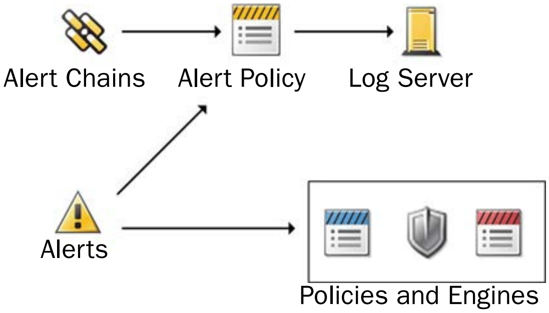

1. (*Optional*) Create new Custom Alert elements to allow more fine-grained matching in the Alert Policy. See Defining Custom Alerts (page 267).
2. Make sure that alerts are generated for all necessary events. See Defining What Triggers an Alert (page 268).
3. Create lists of actions to take when an alert is escalated. See Defining Alert Chains (page 268).
4. Define which alerts are matched against which Alert Chain. See Defining Alert Policies (page 274).

Creating Alerts

Prerequisites: None

There are three predefined Alert elements in the system:

- the *System alert*, which is reserved for alerts about the system operation.
- the *Default* alert, which is a ready-made element that defines the alert that is triggered if no specific alert is triggered.
- the *Test alert*, which is used in Alert Policies for testing Situations.

Additionally, you can define custom Alert elements, which are useful if you want to configure different alert notifications for different types of events.

You can create new Alert elements for traffic matches and test failures, and create specialized handling rules for them in your Alert Policy. System events are always associated with the System Alert element.

What's Next?

- ▶ To create or modify a custom Alert element, proceed to Defining Custom Alerts.
- ▶ To configure which events generate alerts, proceed to Defining What Triggers an Alert (page 268).
- ▶ To send a test alert, see Testing Alerts (page 281).

Defining Custom Alerts

▼ To define a custom alert

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Right-click **Alert Configurations** and select **New→Custom Alert**. The Custom Alert Properties dialog opens.

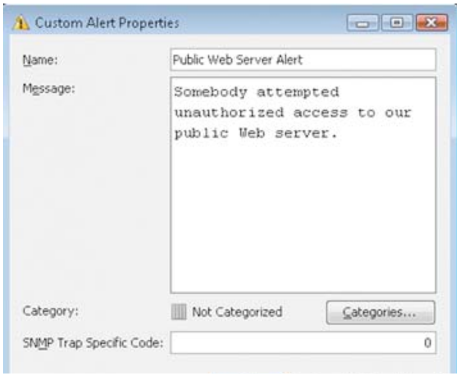

Name: Public Web Server Alert

Message: Somebody attempted unauthorized access to our public Web server.

Category: Not Categorized

SNMP Trap Specific Code: 0

3. Give the alert definition a unique **Name**. This is the name displayed in the Logs view.
4. *(Optional)* Type a Message for the alert. This is also shown in logs.
5. *(Optional)* If you are using SNMP traps for delivering alerts, define the SNMP trap specific code in the **SNMP Trap Specific Code** field. This is included in the SNMP messages sent. The value depends on how the receiving system is set up and the purpose you have planned for this alert.

6. Click **OK**. The new custom Alert element appears in the Alerts folder.

If you edited a custom Alert that is already in use in some Alert Policy, the changes you made take effect only after you refresh the Alert Policy on the Log Server(s).

What's Next?

▶ Use the custom Alert by Defining What Triggers an Alert.

Defining What Triggers an Alert

System Alerts or custom Alerts are always triggered by a situation matching in the system. The Situation can be a warning or error in the operation of the Stonesoft Management Center, a test failure, a rule match, or a match to a pattern defined in a Situation element. In addition to the System alerts triggered by internal events in the SMC, you can configure alerts on the following:

- You can configure a rule in your Firewall, Layer 2 Firewall, or IPS Policy to trigger an alert. See Defining Access Rule Logging Options (page 684).
- You can activate Status Surveillance on engines to trigger an alert when the Management Server does not receive the status updates for a while. See Enabling/Disabling Engine Status Monitoring (page 221).
- You can configure the tester to issue an alert whenever a test fails (for example, when a network link goes down). Some tests are run on the engines by default. See Getting Started with the Engine Tester (page 520).
- Server Pool Monitoring Agents can trigger alerts when they detect problems with the servers. See Getting Started with Inbound Traffic Management (page 618).
- You can set thresholds for monitored items in Overviews to trigger alerts when the threshold is reached. See Setting Thresholds for Monitored Items (page 109).

What's Next?

▶ To define how administrators are notified of new alerts in the system, proceed to Defining Alert Chains (page 268).

Defining Alert Chains

Prerequisites: None

Alert Chains define which notification channels are used to send alert notifications to administrators. Alert Chains are used in Alert Policies. The Alert Policy defines which Alert Chains are triggered by which Alerts.

What's Next?

▶ If you want administrators to be notified of alerts outside of the Management Client and you have not defined Alert channels, start by Defining Alert Channels (page 269).

▶ Otherwise, start by Creating New Alert Chains (page 271) or Modifying Existing Alert Chains (page 271).

Defining Alert Channels

Alert Channels are ways to send notifications to administrators. If you want to send alerts over e-mail, as SMS text messages, through a custom script, or as SNMP traps, you must configure the alert channels you want to use for each Log Server that handles alert escalation.

▼ To define the alert channels

1. Select **Monitoring→System Status**. The System Status view opens.
2. Expand **Servers**, right-click the Log Server and select **Properties**. The Log Server Properties dialog opens.
3. Switch to the **Alert Channels** tab.

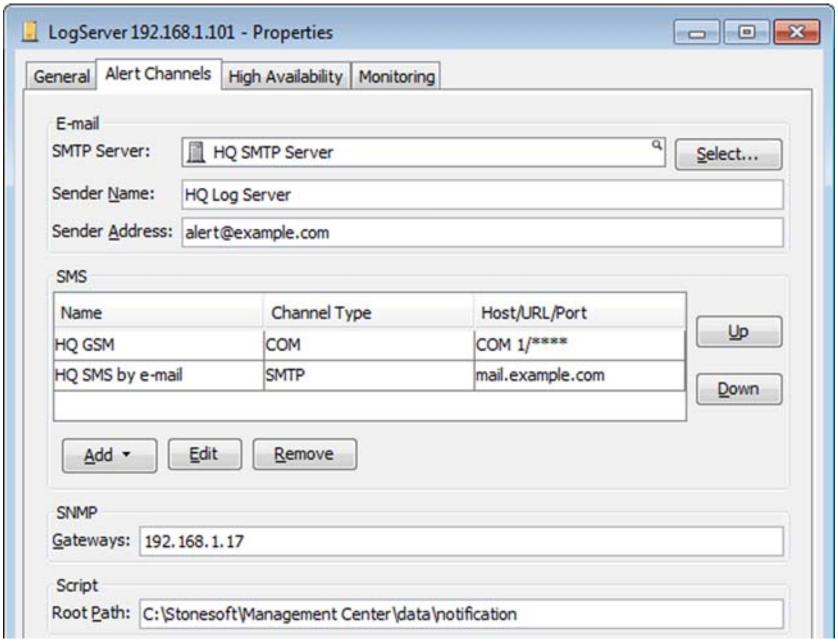

SMTP Server: HQ SMTP Server

Sender Name: HQ Log Server

Sender Address: alert@example.com

Name	Channel Type	Host/URL/Port
HQ GSM	COM	COM 1/****
HQ SMS by e-mail	SMTP	mail.example.com

Gateways: 192.168.1.17

Root Path: C:\Stonesoft\Management Center\data\notification

Tip – You can configure some or all of the alert channels in any combination for each Log Server.

4. (*E-mail*) Configure the following settings:

Setting	Configuration
SMTP Server	Select or create the SMTP Server that is used to send the alert notifications as e-mail. See Creating SMTP Server Elements (page 280) to create a new SMTP Server element.
Sender Name (*Optional*)	Enter the name to be used in the From field of the e-mail. If this setting is left blank, the Default Sender Name defined in the SMTP Server Properties is used.
Sender Address (*Optional*)	Enter the e-mail address to be used in the From field of the e-mail. If this setting is left blank, the Default Sender Address defined in the SMTP Server Properties is used.

5. (*SMS*) Click **Add** and select the Channel Type:
   - **SMTP**: SMS messages are sent using an SMTP server.

Note – You can add multiple SMS Channel Types. The SMS channels are used in the order in which they are listed. Use the Up and Down buttons to change the order of the channels if necessary.

6. (*SMS*) Configure the settings according to the Channel Type:

Channel Type	Setting	Configuration
SMTP	Name	Enter a unique name.
	SMTP Server	Select or create the SMTP Server that is used to send the SMS notifications. See Creating SMTP Server Elements (page 280) to create a new SMTP Server element.
	Account (*If required by SMTP server*)	Enter the user name for connecting to the SMTP server.
	Password (*If required by SMTP server*)	Enter the password for connecting to the SMTP server.
	Local Host name	Enter the DNS host name of the Log Server.
	E-mail Sender Address (*Optional*)	Enter the e-mail address to be used in the From field of the e-mail. If this setting is left blank, the Default Sender Address defined in the SMTP Server Properties is used.
	Recipient	Enter the domain of the SMTP server to which the mail is sent for processing.
	Subject	Enter a subject for the message.

7. (*SNMP*) Enter the host name or IP address of the SNMP **Gateways** to which the alert notifications are sent as SNMP traps.
   - You can specify a list of gateways separated by semi-colons.
   - If your SNMP gateway port is something other than the default port 162, specify the port number by typing a colon and the port after the host name (for example, snmp-gw:4390).
8. (*Script*) Enter the **Root Path** where alert handling scripts are stored on this Log Server.
   - The default location is <installation directory>/data/notification.
   - Do not define the script name here; you add it in the alert chain at each place where you want to call a particular script. You can use multiple scripts. For information on creating the script, see Using Custom Scripts for Alert Escalation (page 278).
9. Click **OK**.

What's Next?

▶ Repeat these steps for all Log Servers you plan to use for Alert Escalation, then continue by Creating New Alert Chains (page 271).

Creating New Alert Chains

You can create several Alert Chains. Alert Chains are used in the Alert Policy, which can escalate different alerts to different Alert Chains according to rules set in the Alert Policy.

▼ To create a new alert chain

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Right-click **Alert Configurations** and select **New→Alert Chain**. The Alert Chain Properties dialog opens.

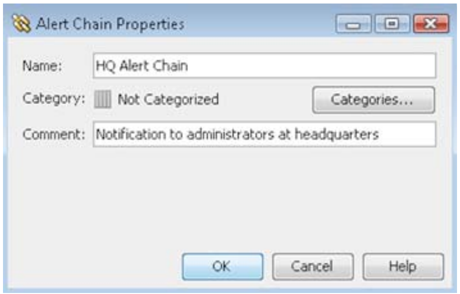

Name: HQ Alert Chain

Category: Not Categorized

Comment: Notification to administrators at headquarters

3. Give the Alert Chain a unique **Name**.
4. Click **OK**. The new alert chain is created and the Alert Chain opens for editing.

What's Next?

▶ Continue by defining the alert escalation chain as explained in Editing Alert Chains (page 272).

Modifying Existing Alert Chains

▼ To modify an existing alert chain

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Browse to **Alert Configurations→Alert Chains**.
3. Right-click the alert chain you want to modify and select one of the following:
   - **Properties**: modify the properties information (such as the name).
   - **Edit Alert Chain:** modify the rules for escalating alerts and modify the Alert Chain as explained in Editing Alert Chains (page 272).

Editing Alert Chains

An Alert Chain is defined in rows that are read from top to bottom. Each row represents a different recipient or a different alert channel to a given recipient.

Tip – It is not mandatory to add any rows to an Alert Chain. For example, you can create a chain that just uses the Final Action to automatically acknowledge or stop the escalation of alert entries that the Alert Policy directs to the chain.

▼ To edit an alert chain

1. Add a row:
   - In an empty Alert Chain, right-click the Final Action row and select **Rule→Add Rule**.
   - In an Alert Chain with existing rules, right-click a rule and select **Rule→Add Rule Before** or **Rule→Add Rule After**.
2. Select the Alert **Channel**. The options are:
   - **Custom Script**: alerts are sent for processing to a script you create. See Using Custom Scripts for Alert Escalation (page 278).
   - **Delay**: processing is paused for the specified time before the next rule is applied.
   - **SMS**: an SMS text message is sent.
   - **SMTP**: an e-mail message is sent.
   - **SNMP**: an SNMP trap is sent.
   - **User Notification**: a blinking icon appears at the bottom right corner of the selected administrators' Management Client. The icon works as a shortcut to a Logs view pre-filtered to display the currently active alerts.
3. Specify the **Destination** of the alert notification. The destination information varies according to the selected alert channel:
   - **Custom script:** Enter the name (or full path) of the script file (note that the root path of the script is defined in the Log Server properties).
   - **SMS:** Enter the recipient's mobile phone number
   - **SMTP:** Enter the recipient's e-mail address. Only one address is allowed. Use a mail group address or create additional rows without delays to send the e-mail to multiple recipients at the same time.
   - **SNMP:** Not editable, because the SNMP server is defined in the Log Server's properties (the actual SNMP trap that is sent depends on the alert event).
4. (*Optional*) To set a limit to how many alerts the designated recipient is sent, double-click the **Threshold to Block** cell and use the dialog that opens to set the following options:
   - **Pass on Max**: Enter the maximum number of alerts that the recipient is to receive. After this threshold is reached, any rows with this recipient are ignored.
   - **During**: Enter the time period in hours (**h**) and minutes (**min**) for counting the number of alerts to the recipient.
   - **Notify first blocking**: select this option if you want the system to notify the alert recipient once alert blocking has started.
   - **No moderation**: select this option if you do not want to set any threshold for blocking.

Note – If you leave the Threshold to Block cell empty, there is no blocking threshold for the alerts sent to the recipient.

5. (*Mandatory for the Delay channel, optional for other channels*) Enter the **Delay**, define a pause (in minutes) before the next row of the alert chain is processed.
   - The purpose of the delay is to give to the recipient of the notification time to acknowledge the alert before the next notification is sent.
   - If sending the notification through the selected channel fails, the delay entered here is ignored. If you want to add delays that are always valid, add a row with **Delay** as the alert channel and set the delay on that row.
6. Continue adding further rows until the Alert chain is complete.

What's Next?

- ▶ If you have not defined what happens when the end of the Alert Chain is reached, proceed to Defining the Final Action of an Alert Chain (page 273).
- ▶ If you have defined what happens when the end of the Alert Chain is reached, save the finished Alert Chain. If it is not included in your Alert Policies yet, proceed as explained in Defining Alert Policies (page 274).
- ▶ Otherwise, reload the Alert Policies that use this Alert Chain (see Installing Alert Policies (page 276)).

Defining the Final Action of an Alert Chain

The Final Action row in an Alert Chain determines what happens when all Alert Channels in the Alert Chain have been tried, but none of the Administrators have acknowledged the alert.

▼ To define the Final Action

1. Select the **Final Action** that the system takes if the last row of the alert chain is reached:
   - **None**: The alert escalation ends. No further notifications are sent, but the alert stays in the list of active alerts.
   - **Acknowledge**: The alert escalation ends and the system automatically acknowledges the alert (removes it from the list of active alerts).
   - **Redirect**: The alert escalation continues in the Alert Chain you select in the other selection box.
   - **Return**: Returns the processing to the Alert Policy. The Alert Policy matching continues from the next row. If there is another matching rule in the Alert Policy, the alert escalation continues. If no further matches are found, the escalation ends and the alert stays in the list of active alerts.
2. Click the Save icon in the toolbar.

What's Next?

- ▶ If the Alert Chain is not included in your Alert Policies yet, proceed as explained in Defining Alert Policies (page 274).
- ▶ Otherwise, reload the Alert Policies that use this Alert Chain (see Installing Alert Policies (page 276)).

Defining Alert Policies

Prerequisites: Defining Alert Chains

Alert Policies determine the criteria for selecting which alerts generated by various sources are escalated to which Alert Chains.

Firewalls, Layer 2 Firewalls, IPS engines, and SMC servers are possible sources for alerts. If Domain elements have been configured in the system, you can select a Domain as a Sender in an Alert Policy in the Shared Domain. Additionally, alert events can be directed to different Alert Chains depending on the time of the day or the Alert element used in the triggering configuration.

What's Next?

▶ Proceed to Creating New Alert Policies or Modifying Existing Alert Policies.

Creating New Alert Policies

▼ To create a new alert policy

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Right-click **Alert Configurations** and select **New→Alert Policy**. The Alert Policy Properties dialog opens.

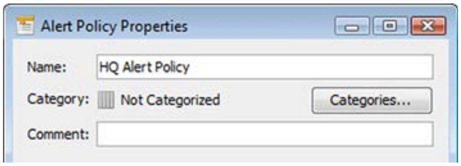

Name: HQ Alert Policy

Category: Not Categorized

Comment:

3. Give the alert policy a unique **Name**.
4. Click **OK**. The new alert policy opens for editing.

What's Next?

▶ Continue by defining the alert policy rules as explained in Editing Alert Policy Rules (page 275).

Modifying Existing Alert Policies

▼ To edit an alert policy

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Browse to **Alert Configurations→Alert Policies** in the element tree. The existing alert policies are displayed.
3. Right-click the appropriate alert policy and select one of the following:
   - Select **Properties** to change the properties information (such as the name).
   - Select **Edit Alert Policy** to modify the actual alert policy rules and proceed as explained in Editing Alert Policy Rules (page 275).

Editing Alert Policy Rules

▼ To edit alert policy rules

1. Add a new rule in the Alert Policy:
   - In an empty Alert Policy, right-click the rule table and select **Rule→Add Rule**.
   - In an Alert Policy with existing rules, right-click a rule and select **Rule→Add Rule After** or **Rule→Add Rule Before**.

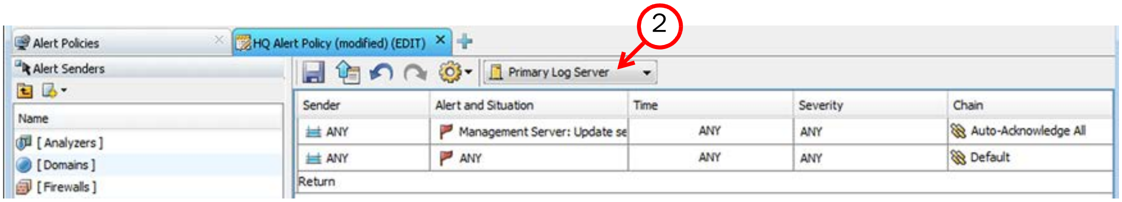

Sender	Alert and Situation	Time	Severity	Chain
ANY	Management Server: Update se	ANY	ANY	Auto-Acknowledge All
ANY	ANY	ANY	ANY	Default
Return				

2. Select the Log Server to which this Alert Policy applies.
3. (*Optional*) Specify the **Sender** that this alert rule matches:
   - You can add Firewalls, Layer 2 Firewalls, IPS engines, Groups, or Servers, or right-click and select **Set to ANY**.
4. (*Optional*) Specify the **Alert and Situation** that this rule matches.
5. (*Optional*) Double-click the **Time** cell and select the days and times when you want the rule to be applied. Click **OK**.
   - If you do not specify any validity time, the alert rule is always applicable.
   - The time is entered as UTC (GMT) time. You must calculate yourself the effects of the time difference to your local time. Also, take into account that UTC does not adjust to daylight saving time (summer time).
6. (*Optional*) Double-click the **Severity** cell and specify the **Severity** value or the range of **Severity** values that this rule matches.
   - To define a single Severity value, select **Severity** and one of the Severity options.
   - If you want the rule to match a range of Severities, select **Severity Range** and define the range in the From and To lists.
7. Select which Alert **Chain** is processed when an alert event matches this rule.
8. Repeat these steps to add any further rules, then click the Save icon in the toolbar.

What's Next?

▶ Activate the Alert Policy on your Log Server by Installing Alert Policies (page 276).

Installing Alert Policies

Prerequisites: None

Changes made to alert policies take effect when you install the Alert Policy on a Log Server.

▼ To install an Alert Policy

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Browse to **Alert Configurations→Alert Policies**.
3. Right-click the Alert Policy you want to install and select **Install Policy**. The Task Properties window opens.
4. Select the Log Server(s) on which you want to install the Alert Policy and click **Add**.
5. Click **OK**.

After you install a policy on a Log Server, you can simply right-click the Log Server element and select **Current Alert Policy→Refresh** to transmit any changes you have made in the Alert Policy to the Log Server.

Acknowledging Alerts

Prerequisites: There must be at least one active alert in the system

Once your system starts generating alert entries, they are sent to the Log Server. Alert entries are displayed together with other types of log entries in the Logs view. Alert entries are marked with an icon on their left side to help them stand out from the other log entries. Alerts that have not yet been acknowledged are marked with an *Unacknowledged Alert* icon (a yellow triangle with an exclamation point and an asterisk).

Alert entries can be acknowledged in the Logs view and in the Web Portal. See the Web Portal *Online Help* for more information about acknowledging alerts through the Web Portal. Once an alert entry is acknowledged by an administrator, it is removed from the list of active alerts and all Alert Chain processing for that alert entry is stopped. Alert entries can be acknowledged individually, or you can acknowledge all the active alert entries at once.

Alert entries can be managed like other log entries. For instructions on using the log management tools, see Getting Started with Log Data Management (page 1016).

What's Next?

▶ Acknowledge alerts either by Acknowledging Individual Alerts (page 277) or by Acknowledging All Active Alerts (page 277).

Acknowledging Individual Alerts

Note – When you acknowledge an alert entry, alert escalation stops for that alert entry and no new notifications are sent out from the Log Server to administrators.

▼ To acknowledge one or a few selected alert entries

1. Select **Monitoring→Logs**. The Logs view opens.
   - Alternatively, you can select **Monitoring→System Status** in the menu and click the number indicating the total number of active alerts to open all of the active alerts in the Logs view.
2. Select **Alert** from the list at the top of the Query panel.
3. (*Optional*) Define additional filtering criteria to find the desired alerts.
4. Click **Apply** in the Query panel to filter the view.
5. Select one or more alert entries in the log entry table.
6. Click the Tools icon and select **Acknowledge**. A confirmation dialog opens.
7. Select **Yes** to acknowledge the alert(s).

Acknowledging All Active Alerts

Note – When you acknowledge an alert entry, alert escalation stops for that alert entry and no new notifications are sent out from the Log Server to administrators.

▼ To acknowledge all displayed alerts

1. Select **Monitoring→Logs**. The Logs view opens.
   - Alternatively, you can select **Monitoring→System Status** in the menu and click the number indicating the total number of active alerts to open all of the active alerts in the Logs view.
2. Click the Tools icon and select **Acknowledge All**. A Confirmation dialog opens.
3. Select **Yes** to acknowledge the alerts.

Using Custom Scripts for Alert Escalation

Prerequisites: None

All custom scripts must be stored in the same root path, which is defined for each Log Server in the Log Server element's properties. See Defining Alert Channels (page 269). If the Log Server is installed in the default installation directory, the path is:

Table 19.1 Default Custom Script Path

Operating System	New Installation of Version 5.4 or Higher	Upgrade to Version 5.4 or Higher
Windows	C:\Program Files\Stonesoft\data\notification\	C:\Stonesoft\StoneGate\data\notification\
Linux	/usr/local/stonesoft/notification/	/usr/local/stonegate/data/notification/

There is an example notification script in this directory called `notify.bat` in Windows and `notify.sh` in Linux that can be modified for your own use. On Linux, the sgadmin user needs read, write, and execute permissions in the script's directory.

The alert information is given to the script as command arguments as described in the table below.

Table 19.2 Arguments Passed to the Custom Scripts

Argument Number	Content	Description
1	Alert ID	Identifies the alert uniquely in the system.
2	Alert Name	The name defined in the alert properties.
3	Alert Originator	The IP address of the component that generated this alert.
4	Alert Date	The date when the alert was originally generated.
5	Alert Message	A short alert description.
6	Alert Severity	The Severity value of the alert from 1 to 10 where 1 is the least severe and 10 is the most severe. The numerical Severity value corresponds to the following Severity value in the generated alert: 1= Info, 2-4=Low, 5-7=High, and 8-10=Critical.
7	Alert Short Description	The contents of the Comment field in the alert properties.
8	Event ID	IPS only: reference to the event ID that triggered the alert.
9	Situation Description	Long description of the Situation that triggered the alert.

When the alert script is executed, the output (stdout) is appended to the `notify.out` file in the script's directory. The error output (stderr) is appended to the `notify.err` file in the script's directory. The Linux script in the illustration below is an example on how to create an operating system log entry using the custom script alert notification.

Illustration 19.2 Example Custom Alert Script

```
#!/bin/sh
# This script uses the ‘logger’ utility to create an operating system
# log entry of the Log Server alert notification.
PATH=/bin:/usr/bin

# Create log entry: “Stonesoft Alert (<ALERT_ID>): Severity <SEVERITY>
#    : <ALERT_NAME> : <ALERT_DESCRIPTION>”

/usr/bin/logger “Stonesoft Alert ($1): Severity $6 : $2 : $5”

exit 0
```

Alert scripts stored in the correct directory (as defined in the Log Server element's properties) can be called from Alert Chains by their name. See Editing Alert Chains (page 272).

Setting up a Dedicated Alert Server

Prerequisites: You must have at least two Log Servers

The Alert Server features are part of every Log Server installation. However, you can also forward all alerts from one Log Server to another. You can, for example, use one Log Server just for handling alerts and have all the other Log Servers send the alerts they generate and receive to the Log Server dedicated to handling alerts.

As a result, Alert Chains and Alert Policies do not need to be managed for each server, which may help in simplifying alert management. On the other hand, you must consider that a loss of that one Log Server then can mean that no alerts are received from anywhere in the system.

▼ To set up a dedicated Alert Server

1. Right-click a Log Server that you do not want to handle alerts anymore, and select **Properties**. The Properties dialog for that Log Server opens.
2. Select the **Forward all Alerts to** option.
3. Select the Log Server you want to use as a dedicated Alert Server.
4. Click **OK** to close the dialog. The Log Server starts forwarding the alerts immediately.

Repeat as necessary for other Log Servers.

Creating SMTP Server Elements

Prerequisites: None

SMTP Server elements define the properties of SMTP servers that are used to send e-mail and SMS notifications. SMTP Server elements can be used in the properties of the following elements:

Table 19.3 Elements Where SMTP Servers can be Used

Element	Purpose
Authentication Server	Notification of changes to user accounts. See Defining Authentication Server Notification Channels (page 868) for more information.
Log Server	Alert notification. See Defining Alert Channels (page 269) for more information.
Management Server	E-mailing Reports. See E-Mailing Reports (page 182) for more information.

▼ To create an SMTP Server element

1. Select **Configuration→Configuration→Security Engine**. The Security Engine Configuration view opens.
2. Browse to **Network Elements→Servers**.
3. Right-click Servers and select **New→SMTP Server**. The SMTP Server Properties dialog opens.

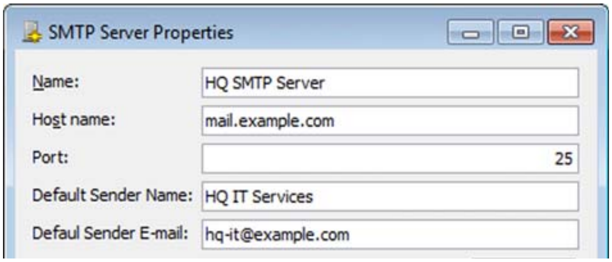

SMTP Server Properties

Name: HQ SMTP Server

Host name: mail.example.com

Port: 25

Default Sender Name: HQ IT Services

Defaul Sender E-mail: hq-it@example.com

4. Configure the settings as described below:

Setting	Configuration
Name	Enter a unique name.
Host Name	Enter the DNS host name (*recommended*) or the IP address of the SMTP Server.
Port	Enter the port number if the SMTP Server uses a port other than the default port 25.
Default Sender Name	Enter the name to be used in the From field of the e-mail. This default value can be overridden in the properties of the element where the SMTP Server is used.
Default Sender E-mail	Enter the e-mail address to be used in the From field of the e-mail. This default value can be overridden in the properties of the element where the SMTP Server is used.

5. Click **OK**.

Testing Alerts

Prerequisites: None

You can test the correct functioning of alerting by sending a test alert. It is treated in the same way as any other alert. The test alerts have the lowest possible severity value “Low” and they have their own specific Situation. Test alerts cannot be sent with Firewall, Layer 2 Firewall, or IPS engines as the sender. The sender is always an SMC server, so it is not possible to test how alerts from other components are handled using the test alert.

Note – A test alert is escalated only if the Alert Policy rules match test alerts.

▼ To test an alert

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Browse to **Alert Configurations→Alerts**.
3. Right-click the alert to be tested and select **Test Alert**. The Select Alert Server dialog opens.
4. Select the server that will send the alert and click **Select**.
5. Check that the alert entry appears in the Logs view.

What’s Next?

▶ Acknowledge the alert as explained in Acknowledging Alerts (page 276).

CHAPTER 20

DOMAINS

Domain elements allow you to restrict which elements are displayed to the administrators in the Management Client and in the optional Web Portal. They also allow you to define in which administrative Domain(s) an administrator has permissions. Configuring Domains requires a special license.

The following sections are included:

Getting Started with Domains

Prerequisites: None

Domain elements help you in managing large networks and in defining administrator permissions.

How Domains Can Be Configured

- Domain elements allow you to group together elements that belong to specific configurations (for example, elements that belong to a specific customer or site).
- You can use Domains to divide responsibilities between administrators, so that administrators only have access to elements in specific Domains.

What Do I Need to Know Before I Begin

- You must have a special license to be able to configure Domain elements. The number of Domains that you can create depends on the license.
- The predefined *Shared Domain* is meant for all the elements that do not belong to a particular customer or site. All predefined system elements belong to the Shared Domain. If there is no Domain license in the Management Center or no Domains have yet been configured, all the elements belong to the Shared Domain.
- The elements in the Shared Domain are displayed to all administrators when they are logged in to any Domain in the Management Client.
- When a Domain is active it can be viewed and edited. When there is more than one Management Server, individual Domains can be activated on different Management Servers. A Domain is active on the Management Server that controls the Domain. A single Domain is active on only one Management Server at a time.
- Domains, Management Servers, Log Pruning Filters, and Administrator accounts with unrestricted permissions are elements that automatically belong to the Shared Domain. You can only create these elements in the Shared Domain, and you cannot move them to any other Domain.
- Licenses and update packages always belong to the Shared Domain.
- If you have Master Engine and Virtual Security Engine elements, the Master Engine must either belong to the Shared Domain or to the same Domain as the Virtual Security Engines. See Getting Started with Virtual Engines (page 404) for more information.

Configuration Overview

1. Generate and install your purchased license for Domains, see Generating New Licenses (page 1047).
2. Define the Domain element(s) in the Shared Domain. See Creating Domains (page 285).
   - If you have already configured Domain elements, you must log in to the Shared Domain to create more Domains. See Logging in to a Domain (page 286).
3. Log in to each Domain and define the elements that belong to the Domain. See Logging in to a Domain (page 286).
   - If you want to move existing elements from one Domain to another, you must first log in to the Domain where the elements are stored and then move the elements to the other Domain. See Moving Elements Between Domains (page 288).

Caution – The elements in the Shared Domain are displayed to all the administrators when they log in to any Domain in the Management Client. Make sure that each Domain-specific element is assigned to the right Domain.

Creating Domains

Prerequisites: You must have installed a license for Domains and you must be logged in to the Shared Domain

Your Domain license defines how many Domains you can create. Only an administrator who has unrestricted permissions can create Domains.

▼ To create a Domain

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Expand the **Other Elements** branch.
3. Right-click **Domains** and select New Domain. The Domain Properties dialog opens.

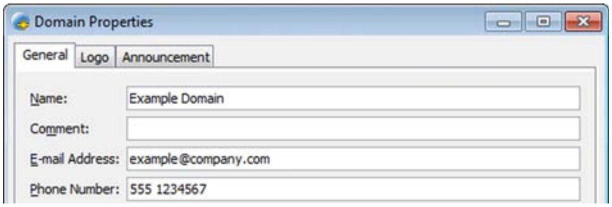

Name: Example Domain

Comment:

E-mail Address: example@company.com

Phone Number: 555 1234567

4. Give the Domain a unique **Name**.
5. (*Optional*) Fill in the **E-mail Address** and **Phone Number** fields with information that you want to be displayed in the Domain Overview. You can enter, for example, the information for the contact person at your company or the administrator responsible for the Domain.

What's Next?

- ▶ If you want to define an announcement that is shown to Web Portal users, continue by Writing Announcements to Web Portal Users (page 305).
- ▶ If you want to define a logo for the Domain, continue by Defining a Domain Logo.
- ▶ Otherwise, click **OK**. The Domain is ready for use. If you want to move existing elements to the Domain, see Moving Elements Between Domains (page 288). To create new elements in the Domain, continue by Logging in to a Domain (page 286).

Defining a Domain Logo

You can define a logo for each Domain. The Domain logo is displayed in the upper right corner of the Management Client window and next to the Domain's name in the Domain Overview. The Domain logo is also displayed in the Web Portal.

▼ To define a Domain logo

1. Switch to the **Logo** tab in the Domain element's properties dialog.

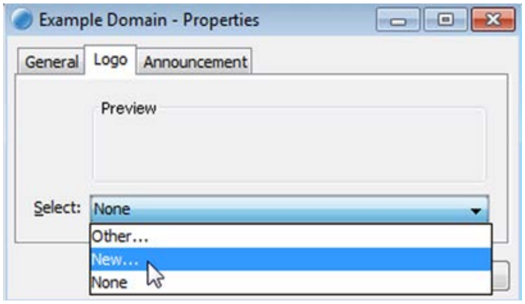

Select: None

2. Select the logo option from the list:
   - Select **None** to clear a previously selected logo so that the Domain has no logo.
   - Select a logo from the list or select **Other** and select a logo in the dialog that opens.
   - Select **New** to import a new logo. See Using the Domain Overview (page 290).
3. Click **OK**. The **Logo** tab shows a preview of the selected Logo.

What's Next?

▶ If you want to define an announcement that is shown to Web Portal users, continue by Writing Announcements to Web Portal Users (page 305).

▶ Otherwise, click **OK**. The Domain is ready to use. To move existing elements to the Domain, see Moving Elements Between Domains (page 288). To create new elements in the Domain, continue by Logging in to a Domain.

Related Tasks

▶ Using the Domain Overview (page 290)

Importing a New Domain Logo

Prerequisites:

▼ To import a new Domain Logo

1. In the Logo Properties dialog, enter a unique **Name** for the logo.
2. Click **Browse** and select the **Image** (a .jpg, .gif, .png, .tif, .bmp, or .pnm file) in the dialog.
3. Click **OK**. The Logo tab of Domain Properties shows a preview of the selected logo.
4. Click **OK** to close the Domain Properties.

Logging in to a Domain

Prerequisites: Creating Domains

If you only have permissions in a single Domain, you automatically log in to the Domain when you log in to the Management Client. If you have permissions in more than one Domain, you must log in to the correct Domain before managing elements that belong to the Domain. You can be logged in to more than one Domain at the same time.

When there is more than one Management Server, it is possible to have different Domains on different Management Servers. A Domain is active on the Management Server that controls it, and each Domain is active on only one Management Server at a time. When you are logged in to a Domain on the Management Server where the Domain is active, the Domain can be viewed and edited. When you log in to a Domain on a Management Server where it is not active, you can only view the elements in the Domain. Logging in to a Domain on a Management Server where it is not active may be necessary if the Management Server where the Domain is active cannot be accessed (for example, due to problems with network connectivity).

You can log in to a Domain in the following ways:

- in the Management Client's login dialog
- in the Domain Overview that opens automatically after Management Client login when you have permissions in more than one Domain

- in the Administration Configuration view.

▼ To log in to a Domain in the Management Client's login dialog

1. To log in to a Domain in the Management Client's login dialog, enter the **User Name** and **Password** as usual.
2. In the Server Address field, select or enter the server's IP address and enter the name of the Domain separated by a slash (/) or a backslash (\) after the IP address.

Example `192.168.200.31/Example Domain or 192.168.200.31\Example Domain`.

- If **Remember Server Address** is selected, the Domain information is saved in the Server Address list for further logins.

3. Click **Login**.

▼ To log in to a Domain in the Domain Overview

1. If the Domain Overview is not open, select **Monitoring→Domain Overview**.
2. Right-click the Domain and select one of the following options:
   - **Login**: Log in to the Domain on the Management Server to which you are currently connected.
   - **Login to Controller** (*multiple Management Servers only*): Log in to the Domain on the Management Server where it is active.

▼ To log in to a Domain in the Administration Configuration view

1. In the Administration Configuration view, browse to **Other Elements→Domains** in the Administration tree.
2. Right-click a Domain and select one of the following options:
   - **Login**: Log in to the Domain on the Management Server to which you are currently connected.
   - **Login to Controller** (*multiple Management Servers only*): Log in to the Domain on the Management Server where it is active.

Logging out of all Domains

Prerequisites: None

If you have permissions in more than one Domain, you can log out of all Domains through the Domain Overview.

▼ To log out of all Domains in the Domain Overview

1. If the Domain Overview is not open, select **Monitoring→Domain Overview**.
2. Select **File→Logout from All**. You are logged out of all Domains.

Moving Elements Between Domains

Prerequisites: Logging in to a Domain

Most elements can be moved to a different Domain. You must be logged in to the Domain in which the elements are located. If there are multiple Management Servers, both Domains must be active on the same Management Server. However, the following elements always belong to the Shared Domain: predefined system elements, Domains, Management Servers, Licenses, update packages, Log Pruning Filters, and Administrator accounts with unrestricted permissions. You cannot move these elements from the Shared Domain.

▼ To move elements between Domains

1. Log in to the active Domain from which you want to move element(s). See Logging in to a Domain (page 286).
   - If there is more than one Management Server, you must log in to the Domain on the Management Server where the Domain is active.
2. Select the element(s) that you want to move to another Domain (Shift- or Ctrl-click to select multiple elements).

Tip – To hide elements from the Shared Domain, click the Tools icon and select Show Only Current Domain.

3. Right-click one of the selected elements and select **Tools→Move to Domain**. The Move to Domain dialog opens.

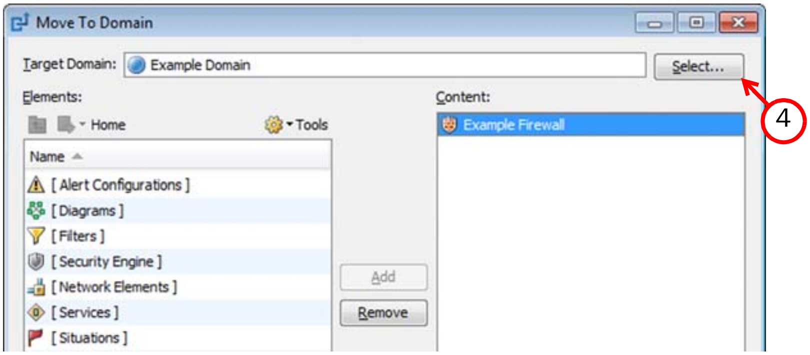

4. Click **Select** for the **Target Domain**.

Tip – You can select additional elements to move from the Elements list.

5. Click **Move**. The Management Server automatically checks if there are references from/to the selected elements in the current Domain. A new view opens, showing the reference status.
   - If the element you are moving references another element, you must either remove the references or move the referenced element as well.

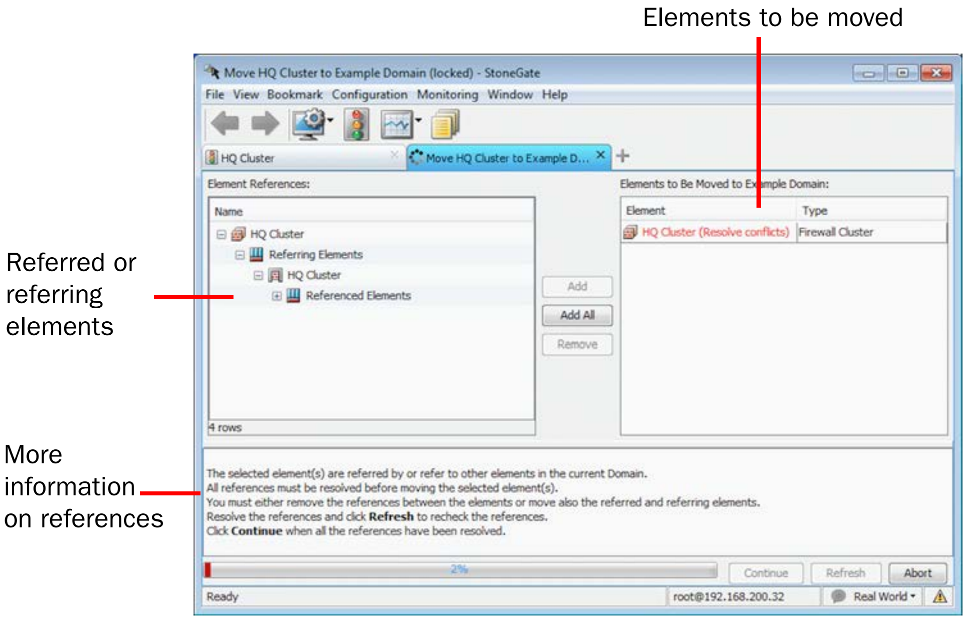
Elements to be moved
Referred or referring elements
More information on references
Move HQ Cluster to Example Domain (locked) - StoneGate
File View Bookmark Configuration Monitoring Window Help
HQ Cluster
Move HQ Cluster to Example D...
Element References:
Name
HQ Cluster
Referring Elements
HQ Cluster
Referenced Elements
4 rows
Add
Add All
Remove
Elements to Be Moved to Example Domain:
Element
Type
HQ Cluster (Resolve conflicts)
Firewall Cluster
2%
Continue
Refresh
Abort
Ready
root@192.168.200.32
Real World

The selected element(s) are referred by or refer to other elements in the current Domain.
All references must be resolved before moving the selected element(s).
You must either remove the references between the elements or move also the referred and referring elements.
Resolve the references and click **Refresh** to recheck the references.
Click **Continue** when all the references have been resolved.

6. If the **Element References** panel shows **Referring** or **Referenced Elements**, expand the branches to see the referring/referenced elements and resolve the element references in one of the following ways:
   - Select the referring/referenced element and click **Add** to also move that element.
   - Right-click the referring/referenced element and modify it to remove the references to the element that you want to move. For example, remove the element that you want to move from a policy that references the element.
7. If you have resolved a reference, click **Refresh** to view the current status of element references.
   - You must resolve all element references before moving the selected elements.
   - If the **Element References** panel is empty, there are no element references.
8. Click **Continue** to move the element(s) to the selected Domain.

What's Next?

▶ If you want to view or modify the moved elements, log in to the Domain to which you moved the elements. If there are multiple Management Servers, log in to the Domain on the Management Server where the Domain is active. See Logging in to a Domain (page 286).

Using the Domain Overview

Prerequisites: Creating Domains

The Domain Overview allows you to see at a glance the status of the Domains and their elements, so that you do not need to log in to each Domain to check its status.

The Domain Overview is available only to administrators who have permissions in more than one Domain. The Domain Overview only shows information from the Domains in which the administrator has permissions. The information in the Domain Overview depends on the administrator's rights. The Domain Overview shows the statuses of the elements and the number of alerts in each Domain as well as any other information (for example, e-mail addresses and telephone numbers) defined in the Domain properties.

If Domain elements have been configured, the Domain Overview also automatically opens when an administrator who has permissions in more than one Domain logs in to the Management Client without specifying the Domain in the login dialog.

▼ To open the Domain Overview

↪ Select **Monitoring→Domain Overview**.

Illustration 20.1 Domain Overview

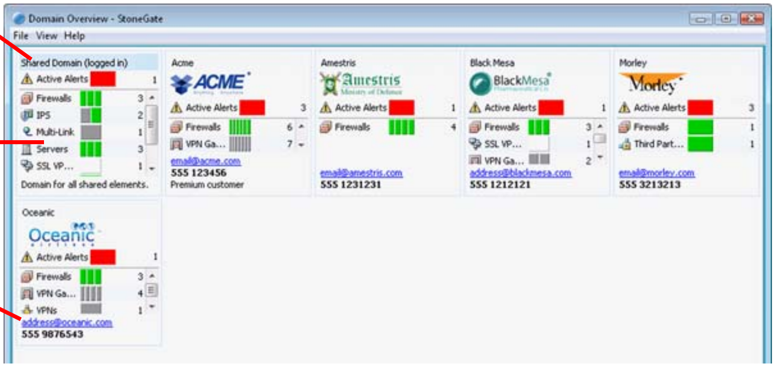

Shared Domain (logged in)
Active Alerts 🟥 1
Firewalls 3
IPS 2
Multi-Link 1
Servers 3
SSL VP... 1
Domain for all shared elements.

Acme
Active Alerts 🟥 3
Firewalls 6
VPN Ga... 7
email@acme.com
555 123456
Premium customer

Amestris
Active Alerts 🟥 1
Firewalls 4
email@amestris.com
555 1231231

Black Mesa
Active Alerts 🟥 1
Firewalls 3
SSL VP... 1
VPN Ga... 2
address@blackmesa.com
555 1212121

Morley
Active Alerts 🟥 3
Firewalls 🟩 1
Third Part... 🟩 1
email@morley.com
555 3213213

Oceanic
Active Alerts 🟥 1
Firewalls 3
VPN Ga... 4
VPNs 1
address@oceanic.com
555 9876543

Login information

To log in to a Domain, double-click the Domain information.

You can send e-mail through the shortcut.

Related Tasks

- ▶ Logging in to a Domain (page 286)
- ▶ Logging out of all Domains (page 287)

Deactivating and Activating Domains

Prerequisites: None

Deactivating Domains

To delete a Domain or to move it between Management Servers, you must first deactivate the Domain. Only administrators who have unrestricted permissions can deactivate a Domain. The Shared Domain cannot be deleted, but it can be moved to a different Management Server.

▼ To deactivate a Domain

1. View all Domains in one of the following ways:
   - Select **File→System Tools→Control Management Servers** to open the Control Management Servers dialog and switch to the **Domains** tab.
   - Select **Monitoring→Domain Overview** to open the Domain Overview.
2. Right-click the Domain and select **Deactivate Domain**.
3. A confirmation dialog opens. Click **Yes** to deactivate the Domain.

Related Tasks

▶ Deleting Domains (page 292)

Activating Domains

You can move a Domain from one Management Server to another by activating the Domain on a different Management Server. Only administrators who have unrestricted permissions can activate a Domain.

▼ To activate a Domain

1. Select **File→System Tools→Control Management Servers**. The Control Management Servers dialog opens.
2. Switch to the **Domains** tab.
3. Right-click the deactivated Domain and click **Select Controller**. The Select Element dialog opens.
4. Select the Management Server on which you want to activate the Domain and click **Select**. A confirmation dialog opens.
5. Click **Yes**. The Domain is activated on the selected Management Server.

Deleting Domains

Prerequisites: Deactivating and Activating Domains

Only administrators who have unrestricted permissions can edit and delete Domains. To delete a Domain, you must first deactivate it. See Deactivating and Activating Domains (page 291).

If you delete a Domain, all the elements that belong to the Domain are also deleted. If there are elements that you do not want to delete, you must move them to another Domain before deleting the Domain. See Moving Elements Between Domains (page 288). You cannot delete the predefined Shared Domain.

▼ To mark a deactivated Domain as obsolete

1. Select **File→System Tools→Control Management Servers**. The Control Management Servers dialog opens.
2. Switch to the **Domains** tab.
3. Right-click the deactivated Domain and select **Mark as Obsolete**.
4. Delete the Domain as instructed in Deleting Obsolete Elements (page 89).

Related Tasks

▶ Restoring Obsolete Elements (page 89).

CHAPTER 21

SETTING UP THE WEB PORTAL

The Web Portal provides browser-based access to logs, reports, and policy snapshots for specific authorized users. The Web Portal is provided by the Web Portal Server, which is an optional component that you can purchase for your Management Center.

The following sections are included:

Getting Started with Web Portal Access

Prerequisites: You must have licenses for running a Web Portal Server and for creating Web Portal Users

What the Web Portal Does

The Web Portal provides restricted clientless access to logs, reports, and policy snapshots. It is particularly useful for managed service providers for providing customers information about their system. There is no software for end-users to install; they can access the information using a web browser.

Limitations of the Web Portal

If Domains are configured in your system, each Web Portal User account is always restricted to working within a single Domain. Administrators with full unrestricted accounts can select between domains after logging in to the portal if the Web Portal Server element is in the Shared Domain.

In addition to the licensed limit for the number of user accounts you can create, each Web Portal Server also has a fixed limit for the maximum number of concurrent users.

Configuration Overview

1. Define a Web Portal Server element. See Defining Web Portal Server Settings (page 295).
2. If providing the Web Portal over HTTPS, generate a certificate for the server. See Activating HTTPS on the Web Portal Server (page 296).
3. Allow the necessary connections in the Firewall Access rules. See Allowing Web Portal Connections (page 297).
4. Install the Web Portal Server. See the *Installation Guide* for instructions.
   - We recommend placing the Web Portal Server in a DMZ network if you offer access to external users.
   - You must generate and install a License for the Web Portal Server.
5. Create Web Portal User accounts for the end-users. See Defining Web Portal User Accounts (page 298).
   - The number of Web Portal users you can configure is limited by license. You must generate and install a separate License file for the Web Portal Users.
   - Management Client administrator accounts are also valid in the Web Portal.
6. (*Optional*) Make installation-specific changes to your portal. See Customizing the Web Portal (page 303).

What's Next?

- ▶ To begin the configuration, proceed to Defining Web Portal Server Settings (page 295).

Related Tasks

- ▶ Writing Announcements to Web Portal Users (page 305)

Defining Web Portal Server Settings

Prerequisites: None

The Web Portal Server is a web server that offers end-users access either through plain HTTP or through the secure HTTPS protocol. The Web Portal Server retrieves data from other Management Center servers, filters it, and presents the resulting data to the users.

Caution – Always use HTTPS unless the connections are otherwise secured.

▼ To define a new Web Portal Server element

1. Select **Monitoring→System Status**.
2. Right-click **Servers** and select **New→Web Portal Server**. The Web Portal Server Properties dialog opens.

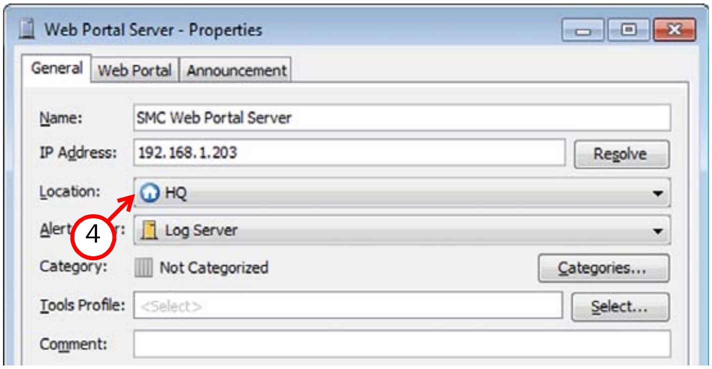

Name: SMC Web Portal Server

IP Address: 192.168.1.203

Location: HQ

Alert Server: Log Server

Category: Not Categorized

Tools Profile: <Select>

Comment:

3. Enter a unique **Name** and the **IP Address** for the Web Portal Server.
4. Select the correct **Location** if necessary in your environment. See Getting Started with System Communications (page 64).
5. Select the Log Server that is used as the **Alert Server** for sending the alerts generated by this Web Portal Server.
6. Switch to the **Web Portal** tab and select **Enable** to activate the Web Portal.
7. (*Optional*) Enter the **Host Name** that the Web Portal uses.
8. (*Optional*) Change the (TCP) **Port Number** that the service listens to. By default, port 8080 is used.

Note – Make sure the listening port is not already in use on the server. For ports reserved for Stonesoft system services, see Default Communication Ports (page 1167).

9. (*Optional*) If the Web Portal Server has several addresses and you want to restrict access to one address, specify this address in the **Listen Only on Address**.
10. (*Optional*) Select **Generate Server Logs** if you want log files to be generated on the file system when the Web Portal is accessed.

What's Next?

- ▶ (*Recommended*) To encrypt the connections between the clients and the server, continue by Activating HTTPS on the Web Portal Server.
- ▶ To use plain HTTP that sends the information in cleartext, click **OK** and continue by Allowing Web Portal Connections (page 297).

Activating HTTPS on the Web Portal Server

Prerequisites: Defining Web Portal Server Settings

To protect the transported information from eavesdropping, you can encrypt the communications by activating HTTPS on the Web Portal Server. If you secure the Web Portal connections using HTTPS, the Web Portal Server requires a certificate. You can either self-sign the certificate directly in the dialog or use an externally signed certificate:

- If you self-sign the certificate directly, web browsers display a warning to the users and require them to accept the certificate. The certificate is valid for one year. Renewing is done by recreating the certificate in the same way as a new certificate is created (explained below).
- Alternatively, you can sign the certificate using an external certificate authority that the clients already trust (such as one of the large commercial certificate authorities or a company-internal certificate authority that all clients are configured to trust).

Certificates have a fixed validity time (from a certain date and time to a certain date and time). Make sure the date, time, and timezone settings are correct on both the Management Server and the Web Portal Server computers, as mismatches may prevent using the Web Portal. Clients also check the certificate validity, but incorrect time settings on the client computers typically do not prevent the Web Portal from being used. Instead, browsers typically display a warning that users can dismiss.

▼ To activate HTTPS on the Web Portal Server

1. Select **Enable** on the Web Portal tab of the Web Portal Server properties dialog.
2. Click **Select** next to the Server Certificate field. The Select Element dialog opens.
3. Select the Server Protection Credentials element or create a new Server Protection Credentials element through the Tools menu at the top of the dialog. See Configuring Server Protection (page 796).
4. Click **OK**.

What's Next?

- ▶ Modify the Access rules to allow Web Portal connections as instructed in Allowing Web Portal Connections (page 297).

Allowing Web Portal Connections

Prerequisites: Defining Web Portal Server Settings

If the connections are routed through a firewall, you must modify the IPv4 Access rules in your security policy to allow the connections from the Web Portal Server to the Management Server and the Log Server, and the end-user connections to the Web Portal Server.

▼ To configure Access rules for the Web Portal

➥ Allow following connections in the IPv4 Access rules as needed:

Table 21.1 Access Rules for Web Portal Server

Source	Destination	Port/Service	Action
Web Portal users' networks (or hosts)	Web Portal Server	Port defined in the Web Portal Server Properties dialog.	Allow
Web Portal Server	Management Server	*SG Control* (8902-8913/TCP)	Allow
Web Portal Server	Log Server(s)	*SG Data Browsing (Web Portal Server)* (8916-8917/TCP)	Allow

Note – Remember to adjust the NAT rules as well if it is necessary in your network setup.

What's Next?

- Install the Web Portal Server's license and the Web Portal Server as instructed in the *Stonesoft Management Center Installation Guide*.
- Create Web Portal User accounts for the end-users. See Defining Web Portal User Accounts (page 298).
- To change the appearance of the portal, see Customizing the Web Portal (page 303).

Defining Web Portal User Accounts

Prerequisites: The number of Web Portal user accounts in the system must not exceed your licensed limit

End-users are defined using special Web Portal User accounts. Management Client administrator accounts are also valid in the Web Portal to allow testing and other internal use.

If Domain elements have been configured, each Web Portal user account always belongs to a specific Domain and only grants permissions to the Domain in which the Web Portal User element is created.

▼ To create a Web Portal user account

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Right-click **Access Rights** and select **New→Web Portal User**. The Web Portal User Properties dialog opens.

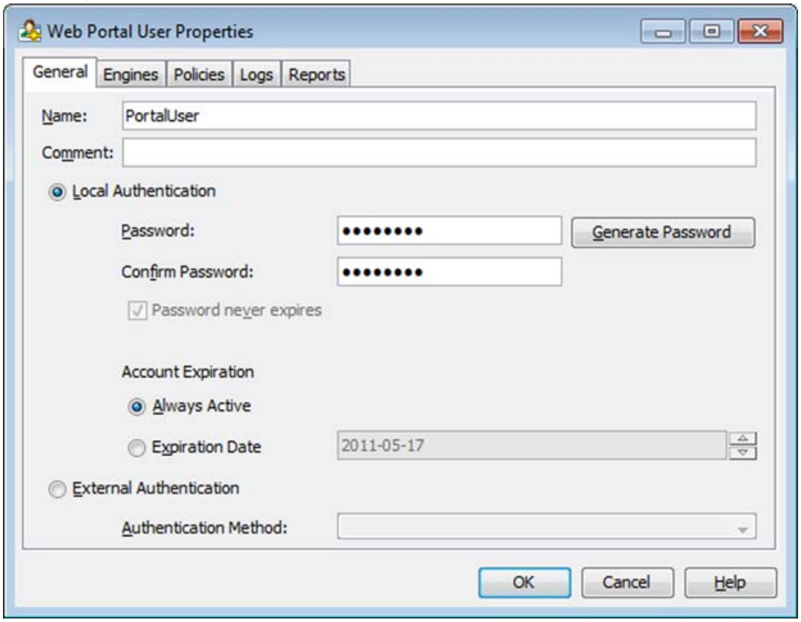

Name: PortalUser

Comment:

● Local Authentication

Password: •••••••••

Confirm Password: •••••••••

☑ Password never expires

Account Expiration

● Always Active

○ Expiration Date 2011-05-17

○ External Authentication

Authentication Method:

3. Enter a unique Web Portal user **Name** that the Web Portal user uses to log in to the Web Portal.
4. For **Internal Authentication**, configure the settings as follows:
   - Type in a **Password** and confirm it in the field below, or click **Generate Password** to generate a random 7-digit alphanumeric password.
   - Select whether the account is **Always Active** or enter an **Expiration Date**.
5. For **External Authentication**, select the Authentication Method from the list.
   - See Authenticating Administrators Using RADIUS Methods (page 262) for more information.

Caution – We recommend that passwords be at least eight characters long and contain a combination of numbers, letters, and special characters. Secure passwords are never based on personal information such as names, birthdays, social ID numbers, phone numbers, street names, or registration plate numbers.

What's Next?

- Continue by Granting Engines to a Web Portal User.

Granting Engines to a Web Portal User

To define a Web Portal user's permissions, you must first select the engines.

To select granted engines for a Web Portal User

1. Switch to the **Engines** tab in the Web Portal User element's properties.

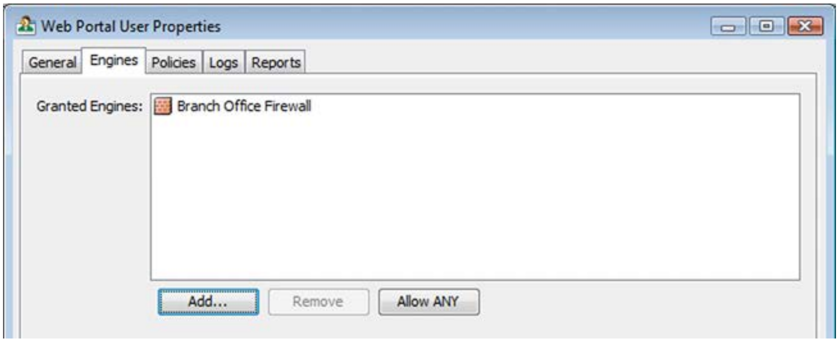

Granted Engines: Branch Office Firewall

2. Select the granted engines:
   - Click **Allow ANY** if you want to grant all the engines to the Web Portal user. If Domain elements have been configured, only the engines in the current Domain are granted to the Web Portal user.
   - Or click **Add** to select the engines. The Select Engine dialog opens.
3. Select the engine(s) from which the administrator is allowed to view logs and/or policy snapshots and click **Select**.

What's Next?

- Continue by Selecting Policy Permissions for a Web Portal User (page 300).

Related Tasks

- Selecting Report Data Permissions for a Web Portal User (page 302)
- Selecting Log Browsing Permissions for a Web Portal User (page 301)

Selecting Policy Permissions for a Web Portal User

You can select one or more policies from which the Web Portal user is allowed to view policy snapshots. You can only select policies that are installed on engines granted to the Web Portal user. See Granting Engines to a Web Portal User (page 299). You can define in detail which parts of the policies are shown in the policy snapshots.

▼ To select policy permissions for a Web Portal User

1. Switch to the **Policies** tab in the Web Portal User element's properties.

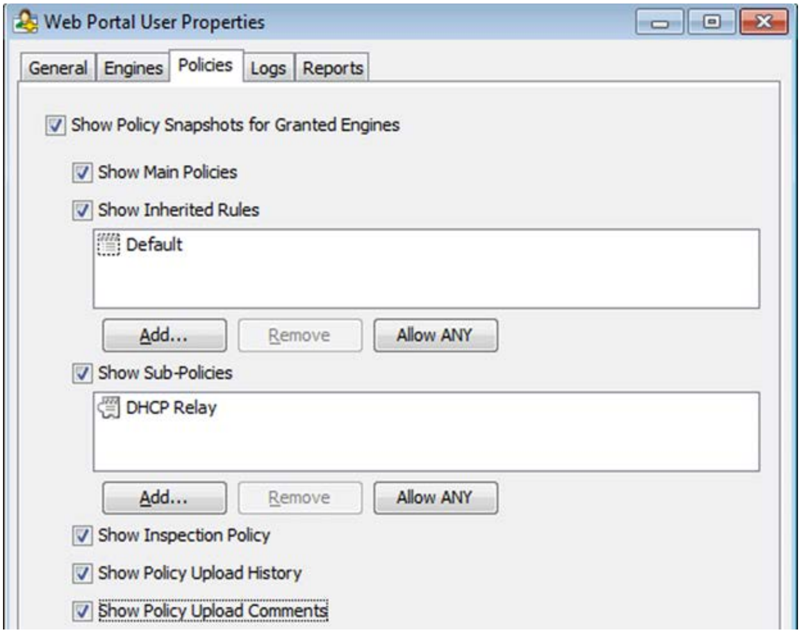

☑ Show Policy Snapshots for Granted Engines

☑ Show Main Policies

☑ Show Inherited Rules

Default

☑ Show Sub-Policies

DHCP Relay

☑ Show Inspection Policy

☑ Show Policy Upload History

☑ Show Policy Upload Comments

2. Define if the Web Portal user has access to policy snapshots:
   - Deselect **Show Policy Snapshots from Granted Engines** if you want to deny the Web Portal user access to policy snapshots. Proceed to Selecting Log Browsing Permissions for a Web Portal User (page 301) to define the log permissions for the Web Portal user.
   - Otherwise, make sure that the option is selected.
3. (*Optional*) Deselect **Show Main Policies** if the Web Portal user is not allowed to view the rules in the upper-level policies.
4. (*Optional*) Select **Show Inherited Rules** if the Web Portal user is allowed to view rules inherited from policy templates.
   - If you want to select from which templates inherited rules are shown, click **Add** and select the template(s) in the dialog that opens.
5. (*Optional*) Deselect **Show Sub-Policies** if the Web Portal user is not allowed to view information from sub-policies.
   - If you want to select the sub-policies from which rules are shown, click **Add** and select one or more sub-policies in the dialog that opens.
6. (*Optional*) Deselect **Show Policy Upload History** if the Web Portal user is not allowed to view and compare all policy snapshots from the granted engines.
   - If this option is selected, the Web Portal user can view and compare policy snapshots of any policies that have been installed on the granted engines (not only policy snapshots of policies granted to the Web Portal user).

7. (*Optional*) Deselect **Show Policy Upload Comments** if the Web Portal user is not allowed to view the comments that administrators have added at policy upload.

What's Next?

- Continue by Selecting Log Browsing Permissions for a Web Portal User.

Selecting Log Browsing Permissions for a Web Portal User

▼ To select log browsing permissions for a Web Portal User

1. Switch to the **Logs** tab in the Web Portal User element's properties.

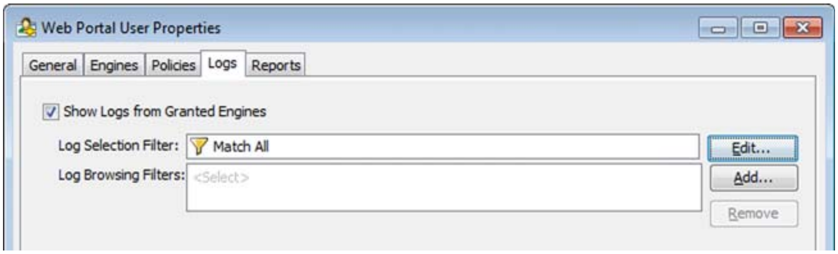

Web Portal User Properties

General | Engines | Policies | Logs | Reports

☑ Show Logs from Granted Engines

Log Selection Filter: Match All

Log Browsing Filters: <Select>

2. Define if the Web Portal user is allowed to view logs in the Web Portal:
   - Deselect **Show Logs from Granted Engines** to deny the Web Portal user access to logs. Proceed to Selecting Report Data Permissions for a Web Portal User (page 302).
   - Otherwise, make sure that the option is selected.
3. (*Optional*) Click **Edit** next to **Log Selection Filter** to define a filter that is applied to the log data before the data is displayed to the Web Portal user. The Local Filter Properties dialog opens. For information on how to create local filters, see Creating and Editing Local Filters (page 187).
4. (*Optional*) Click **Add** next to **Log Browsing Filters** to select one or more Filter Type(s) that define the groups of Filters that the Web Portal user can use when browsing log data in the Web Portal.

What's Next?

- Continue by Selecting Report Data Permissions for a Web Portal User (page 302).

Selecting Report Data Permissions for a Web Portal User

▼ To select report data permissions for a Web Portal User

1. Switch to the **Reports** tab in the Web Portal User element's properties.

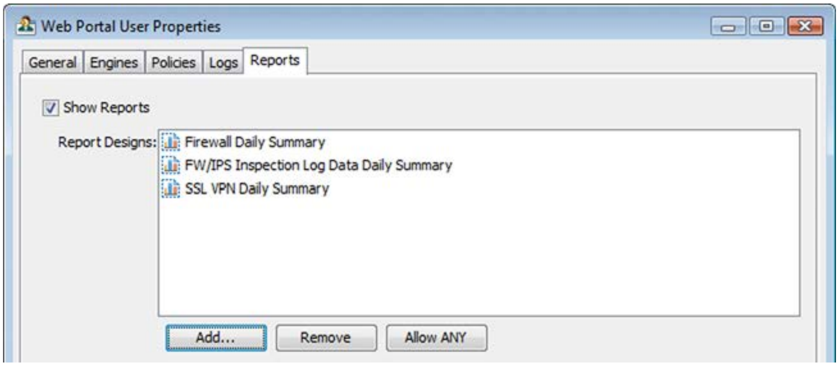

2. Define whether the Web Portal user is allowed to view reports in the Web Portal.
   - Deselect **Show Reports** to deny the Web Portal user access to reports.
   - Otherwise, make sure that **Show Reports** is selected.
3. Click **Add** if you want to select the Reports Designs based on which the Web Portal user is allowed to view reports. The Web Portal user is allowed to view all the reports that are based on the granted Report Resigns (regardless of the Domain in which the reports were created if Domain elements have been configured).
4. Click **OK** to save the Web Portal User element. The Web Portal user account is ready for use.

Related Tasks

▶ Reports (page 167)

Adding a New Web Portal Language

You can add new translations of the Web Portal interface labels in addition to the default languages offered. When you modify the language files, save the file using the UTF-8 or the UTF-16 character encoding. Visit the StoneBlog user community at http://stoneblog.stonesoft.com/stoneblog-community/files/ for some example translation files.

Importing a Web Portal Language File through the Management Client

▼ **To import a language file through the Management Client**

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Browse to **Other Elements→Web Portal Localizations**. A list of existing Web Portal Localizations opens.
3. Right-click the list and select **New Web Portal Localization**. The Web Portal Localization Properties dialog opens.
4. Click **Import** and browse to the location of the language file. You are prompted to confirm the change of Locale.
5. Click **Yes**. The contents of the imported language file are displayed in the dialog.

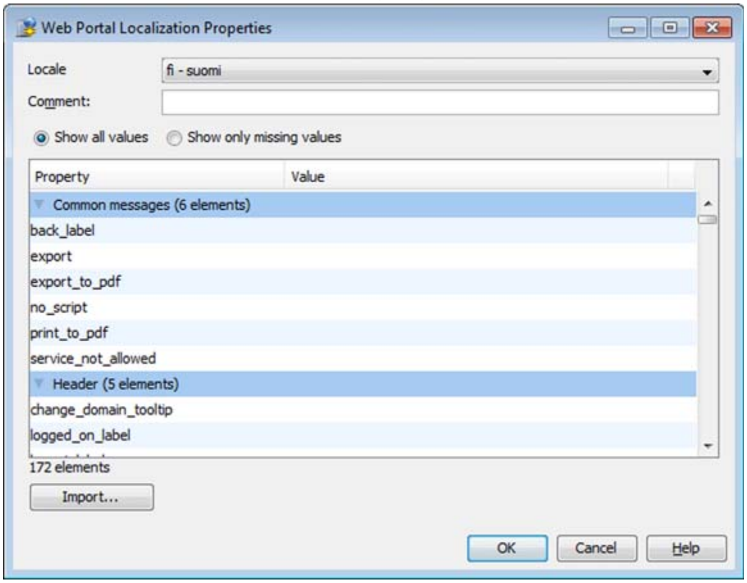

Web Portal Localization Properties

Locale: fi - suomi

Comment:

● Show all values ○ Show only missing values

Property	Value
Common messages (6 elements)	
back_label	
export	
export_to_pdf	
no_script	
print_to_pdf	
service_not_allowed	
Header (5 elements)	
change_domain_tooltip	
logged_on_label	

172 elements

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

Enabling/Disabling a Web Portal Localization

Web Portal Localization languages can be enabled and disabled through the Management Client. Disabled Web Portal Localizations are marked with an icon.

▼ To enable/disable a Web Portal language

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Browse to **Other Elements→Web Portal Localizations**. A list of existing Web Portal Localizations opens.
3. Right-click the Web Portal Localization and select **Enable**/**Disable**. The language is shown/ hidden in the Web Portal language selection options.

Customizing the Look of the Web Portal

The Web Portal pages presented to the users are generated dynamically based on configuration files. The following aspects of the presentation can be adjusted:

- If you are using Domains, the icon in the Web Portal is sourced from the icon defined for the Domain in the Management Client. See Defining a Domain Logo (page 285).
- You can add new languages as explained in Adding a New Web Portal Language (page 303).

In addition to the two items listed above, it is possible to customize the Web Portal source files more extensively, but with some major limitations:

- Any changes to files are overwritten at upgrade.
- It may not be possible to reuse customized files after an upgrade. You may have to re-customize the corresponding new files after the upgrade.
- Customization requires some previous knowledge of HTML and CSS coding.

With the issues mentioned above in mind, the following local files on the Web Portal Server control the display of content to users and are fairly safe to edit (although we recommend that you create copies of the files before you edit them):

- The CSS files are stored in `<installation directory>/webserver/webapps/webclient/resources/css/`.
- The page template files that are stored in `<installation directory>/webserver/webapps/webclient/` as several `.jsp` files.
- The help pages that are stored in `<installation directory>/webserver/webapps/webclient/help/` as several `.jsp` files.

Writing Announcements to Web Portal Users

Prerequisites: See Getting Started with Web Portal Access

You can display announcements to the administrators who log in to the Web Portal. This can be useful, for example, for service break notifications.

▼ To define an announcement for Web Portal users

1. Open the properties of the correct element:
   - Announcements in the Web Portal Server element are shown for all users that connect to that Web Portal Server.
   - Announcements in the Management Server properties are shown for all users.
   - If administrative Domains have been configured, the announcements defined in the properties of the Domain elements are shown to the users of that specific Domain (or to all users if defined in the properties of the Shared Domain).
   - You can define an announcement in any combination of the elements listed above. Each type of announcement adds a bullet point for users who can view several announcements.
2. Switch to the **Announcement** tab.
3. Select **Display Announcement to Web Portal Users**, and enter the announcement in the field below. The length is limited to 160 characters. You can add formatting to the announcement with standard HTML tags (which are also included in the character count).

Example This produces <b>boldface</b> text.

Example This produces <u>underlined</u> text.

Example This text is displayed in the automatically added bullet point.
</li><li>But this text is displayed in a second bullet point underneath.

Example Here is an internal <a href=”policies/policies.jsf”>link</a>.

Example Here is an external <a href=”http://www.example.com”>link</a>.

Example Here is a small image scaled to be roughly proportionate with the surrounding text
<img height=”20px” src=”http://www.example.com/images/note_icon.jpg”>

Note – If you leave out the protocol (HTTP:// or HTTPS:// part) from a URL, the protocol is attached automatically based on the protocol the Web Portal Server is using. This may prevent an otherwise valid link from working.

4. Click **OK**. The announcement is displayed on the home page of the Web Portal (Services page) to affected users the next time they (re)load the page.

The announcement can be removed from the Web Portal by disabling the **Display Announcement to Web Portal Users** option. The announcement you have entered is not deleted, so you can later display the same announcement without typing it in again.

Related Tasks

▶ Defining Web Portal User Accounts (page 298)

CHAPTER 22

DISTRIBUTING MANAGEMENT CLIENTS THROUGH WEB START

The Management Client can be distributed through Web Start. This eliminates the need for each administrator to upgrade their client when the SMC is upgraded to a new version (the version of the client must always match the version of the Management Server).

The following sections are included:

Getting Started with Web Start Distribution

Prerequisites: None

What Web Start Distribution Does

Distributing Management Clients through Web Start allows users to log in to the Management Client through a web browser.

The Management Server can be configured to serve a web page for launching the Management Client. When the Management Center is upgraded, the Web Start files are also updated, and Web Start automatically downloads the updated version when the user logs in. Alternatively, you can make the Management Clients available on some other server, with the limitations described below.

Limitations of Web Start Distribution

The new Management Client version is automatically available to the Web Start users when the Management Center is updated only if you use the Management Server as the Web Start Server. If you install the Web Start package on another server, you must delete the existing files and install a new Web Start package according to the instructions in Distributing Web Start from External Servers (page 310) each time you upgrade the Management Center.

When the Web Start package is installed on a shared network drive, the path to the network, including the drive letter, must be the same for all administrators that use that particular version of the Web Start package. If the network drive paths vary, consider putting the package on a web server instead, for example, in the Intranet of your company.

What Do I Need to Know Before I Begin?

A current version of Java Runtime Environment (JRE) must be installed on each computer where the Web Start installation is to be used. It can be downloaded for free from www.java.com or from java.sun.com.

Configuration Overview

1. Enable Web Start access in one of the following ways depending on where users will access the Web Start login:
    - Enable Web Start access on the Management Server as instructed in Activating Web Start on the Management Server (page 309)
    - Configure Web Start access on a separate web server or shared network drive as instructed in Distributing Web Start from External Servers (page 310).
2. Test the Web Start distribution as instructed in Accessing the Web Start Management Clients (page 311).

What's Next?

- ▶ If you want to use the Management Server as a Web Start Server, proceed to Activating Web Start on the Management Server (page 309).
- ▶ Otherwise, proceed to Distributing Web Start from External Servers (page 310).

Activating Web Start on the Management Server

Prerequisites: None

When you install the Management Server, the files needed for distributing the Management Clients are included in the installation. You can simply enable Web Start access to these files on the Management Server. This activates a web page that administrators can contact using a web browser to launch the Management Client.

▼ To enable a Web Start Server

1. Right-click the Management Server and select **Properties**. The Properties dialog opens.
2. Switch to the **Web Start** tab.

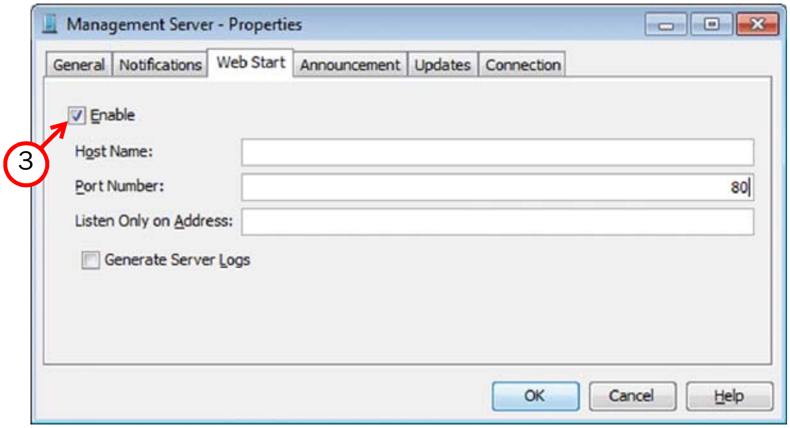

☑ Enable

Host Name:

Port Number: 80

Listen Only on Address:

□ Generate Server Logs

3. Select **Enable**. The Web Start Server options are enabled.
4. (*Optional*) Enter the **Host Name** that the Web Start service uses.
5. (*Optional*) Enter the (TCP) **Port Number** that the service listens to.
   - By default, the standard HTTP port 80 is used on Windows and 8080 on Linux (which does not allow the use of reserved ports for this type of service).

Note – Make sure the listening port is not in use on the server. For ports reserved for Management Center services, see Default Communication Ports (page 1167).

6. (*Optional*) If the Management Server has several addresses and you want to restrict access to one address, specify the IP address to use in the **Listen Only on Address** field.
7. (*Optional*) Select **Generate Server Logs** if you want to log all file load events for further analysis with external web statistics software.
8. Click **OK**.

What's Next?

▶ Test the Web Start installation by following the instructions in Accessing the Web Start Management Clients (page 311).

Distributing Web Start from External Servers

Prerequisites: None

You can use Web Start even if you do not want to use the Management Server as a Web Start Server. In this case, you can put the Web Start package on any web server.

The Web Start package can also be put on a shared network drive. When the Web Start package is installed on a shared network drive, the path to the network, including the drive letter, must be the same for all administrators that use that particular version of the Web Start package.

Note – You must delete the existing files and install a new Web Start package according to these instructions each time you upgrade the Management Center. Otherwise, any administrators that use Management Clients that are installed through Web Start are not able to log in.

▼ To install the Web Start package

1. Browse to **Stonesoft_SW_Installer→Webstart** on the installation DVD.

Caution – The Web Start installation creates an `index.html` **file. Any existing** `index.html` **file will be overwritten. We strongly recommend creating a new directory for the Web Start files.**

2. Copy all files and all directories from the Web Start directory on the installation DVD to the directory where you want the Web Start files to be served.
3. On the command line, change to the directory where the Web Start files are located on your server.
4. Run the Web Start setup script and give the URL or the path of the directory where the Web Start files are located on your server as the parameter:
   - Windows: **`cscript webstart_setup.vbs`** *`<web start directory>`*
   - Linux: `./`**`webstart_setup.sh`** *`<web start directory>`*

Table 22.1 Example Web Start Paths

Installation on	Example Web Start Directory
Web server	`http://www.example.com/webstart/`
Network drive	`file://localhost/C:/support/webstart/`

5. If necessary, modify the configuration of the web server to return the appropriate MIME type for `.jnlp` files (`application/x-java-jnlp-file`). Consult the manual of your web server for instructions on how to configure the MIME type.
6. Delete the `webstart_setup.vbs` and `webstart_setup.sh` files from the directory.
   - For security reasons, the webstart_setup.vbs and webstart_setup.sh script files should not be in the same directory as the generated Web Start files and other shared client files.

What's Next?

▶ Test the Web Start installation by following the instructions in Accessing the Web Start Management Clients.

Accessing the Web Start Management Clients

Prerequisites: Activating Web Start on the Management Server/Distributing Web Start from External Servers

After the Web Start package is installed on a web server or a network drive or the Management Server has been enabled as a Web Start Server, the administrators can install the Management Client using the Web Start package.

To be able to use the Web Start Management Client, there must be a current version of Java Runtime Environment (JRE) installed (the required version is shown on the example login page provided).

Note – If Web Start access is required through the firewall, you must allow these connections in your firewall's policy. They are not allowed by default.

▼ To access the Web Start Management Clients

1. Enter the Web Start download page address in your web browser `http://<server address>:<port>`
   - `:<port>` is only needed if the server is configured to run on a different port from the HTTP standard port 80.
2. Click the link for the Web Start Management Client.
   - Web Start automatically checks if the version on the server is already installed on your local computer. If not, the new client is automatically installed on your computer. This is done each time the client is started this way, automatically upgrading your client installation whenever needed without any action from you.
   - The client starts and displays the login dialog.
3. Log in with your account credentials.

Related Tasks

▶ See Administrator Accounts (page 247) for information on setting up user accounts for Management Clients.

CHAPTER 23

LOG SERVER CONFIGURATION

The sections below contain information on how to define Log Servers, how to forward log data from a Log Server to an external host, how to change Log Server configuration parameters, and how to forward log data from a Log Server to an external syslog server.

The following sections are included:

Defining a Log Server

Prerequisites: None

Usually, the Log Server element is created during the Log Server installation. However, a different Log Server element can also be defined manually. For more information on defining a Log Server element during installation, see the *Stonesoft Management Center Installation Guide*.

Related Tasks

- Defining a Log Server Element
- Certifying the Log Server (page 316)

Defining a Log Server Element

▼ To define a new Log Server element

1. Select **Monitoring→System Status**. The System Status view opens.
2. Right-click **Servers** and select **New→Log Server**. The Log Server Properties dialog opens.

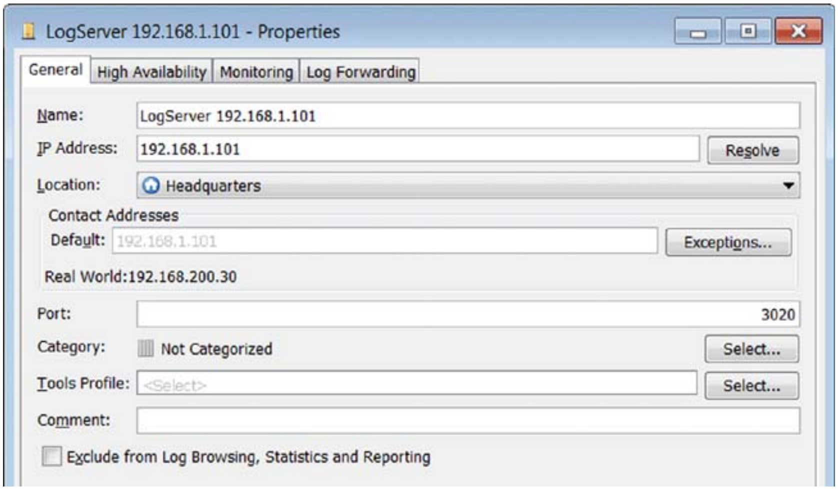

LogServer 192.168.1.101 - Properties

General | High Availability | Monitoring | Log Forwarding

Name: LogServer 192.168.1.101

IP Address: 192.168.1.101 [Resolve]

Location: Headquarters

Contact Addresses

Default: 192.168.1.101 [Exceptions...]

Real World:192.168.200.30

Port: 3020

Category: Not Categorized [Select...]

Tools Profile: <Select> [Select...]

Comment:

□ Exclude from Log Browsing, Statistics and Reporting

3. Specify the **Name** and **IP Address** of the server.
4. A **Location** and **Contact Address** are needed if there is a NAT device between a Security Engine and the Log Server. For instructions, see Defining Contact IP Addresses (page 66).
5. (*Optional*) Define the Log Server's TCP **Port Number**. We recommend that you always use the default port 3020 if possible.

Note – If you want to use a non-standard port for the Log Server, you must always manually add Access rules to allow communications using the new port from the Security Engine(s) to the Log Server.

6. *(Optional)* If you want to use additional commands with the new Log Server, select a Tools Profile for adding the new commands.
7. (*Optional*) Select **Exclude from Log Browsing, Statistics and Reporting** if you do not want the Log Server to gather statistical information for monitoring and you do not want its logging data to be included in Reports. In most situations, it is better to leave this option

deselected. For more information on Reports, see Reports (page 167).

Caution – Be very careful when excluding Log Servers from reporting. If you select this setting for a Log Server that is in use, there is no warning that the generated reports are missing parts of the log data.

8. Click **OK**.

What's Next?

▶ If you want to set backup Log Servers for the Log Server, select the backup Log Servers as described in Selecting Backup Log Servers.

▶ Otherwise, check that the Log Server is licensed and certified. See Generating New Licenses (page 1047) and Certifying the Log Server (page 316).

Selecting Backup Log Servers

You can select several backup Log Servers for a Log Server in a Log Server's properties. The same Log Server can simultaneously be the main Log Server for some components and a backup Log Server for components that primarily use some other Log Server. You can also set Log Servers to be backup Log Servers for each other so that whenever one goes down, the other Log Server is used. If you have not yet defined backup Log Servers, see Installing Additional Log Servers (page 335).

Caution – If the log volumes are very high, make sure that the backup Log Server can handle the traffic load in fail-over situations.

▼ To select a backup Log Server

1. Select **Monitoring→System Status**. The System Status view opens.
2. Expand the **Servers** branch.
3. Double-click to open the properties of the Log Server that you want to select a backup Log Server.
4. Switch to the **High Availability** tab.

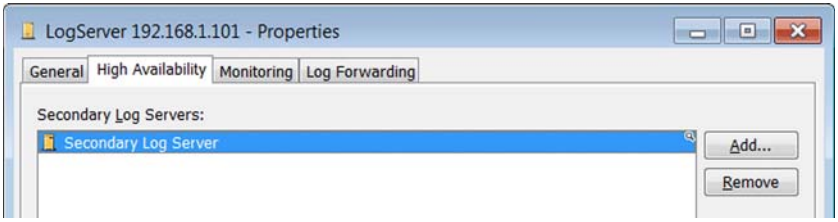

5. Click **Add**. The Select Element dialog opens.
6. Select one or more Log Servers and click **Select**.
7. Click **OK**.

What's Next?

▶ If the Log Server was not certified during the installation, or if it needs a new certificate, continue by Certifying the Log Server.

Certifying the Log Server

If the Log Server was not certified during the installation, or if it needs a new certificate, certify it as directed below.

▼ To certify the Log Server

1. Stop the Log Server:
   - If you run the Log Server as a service in Windows, you can stop it in the Windows Control Panel's **Services** list.
   - In Linux, run the script *<installation directory>*`/bin/sgStopLogSrv.sh`.
2. Request the certificate:
   - In Windows, run the script *<installation directory>*`/bin/sgCertifyLogSrv.bat`.
   - In Linux, run the script *<installation directory>*`/bin/sgCertifyLogSrv.sh`.
3. Enter the credentials for an Administrator account with unrestricted privileges.
4. If there are Domains configured and the Log Server does not belong to the Shared Domain, enter the name of the **Domain**.
5. Wait for the certification to finish and start the Log Server again through the Windows Services list or using the **sgStartLogSrv** script.

Forwarding Log Data to an External Host

Prerequisites: Defining a Log Server

Log Servers can be configured to forward log data to external hosts. You can define which type of log data you want to forward and in which format. You can also use Filters to specify in detail which log data is forwarded.

Defining Log Forwarding Rules

▼ To create and modify Log Forwarding rules

1. Select **Monitoring→System Status**.
2. Expand the **Servers** branch.
3. Right-click the Log Server from which you want to forward log data and select **Properties**. The Log Server Properties dialog opens.
4. Switch to the **Log Forwarding** tab.

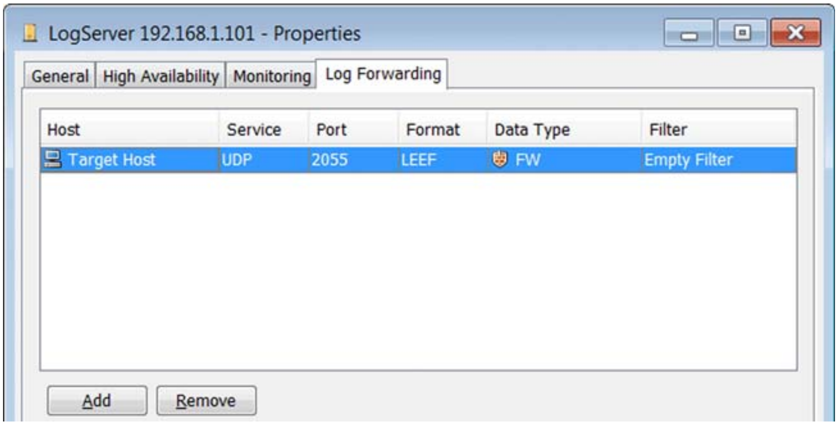

Host	Service	Port	Format	Data Type	Filter
Target Host	UDP	2055	LEEF	FW	Empty Filter

5. Edit the Log Forwarding rules:
   - Select a rule click select **Remove** to remove a Log Forwarding rule.
   - Click **Add** to create a new Log Forwarding rule. A new row is added to the table.
6. Configure the Log Forwarding rule as explained in the table below:

Table 23.1 Log Forwarding Rule Settings

Cell	Value	Explanation
Target Host		The Host element that represents that target host to which the log data is forwarded.
Service	TCP UDP	The network protocol for forwarding the log data. **Note!** If you have to define an Access rule that allows traffic to the target host, make sure that the Service you select is also used as the Service in the Access rule.
Port		The Port that is used for log forwarding. The default port is 2055. **Note!** If you have to define an Access rule that allows traffic to the target host, make sure that the Port you select is also used as the Port in the Access rule.
Format	CEF	Logs are forwarded in CEF format.
	CSV	Logs are forwarded in CSV format.
	LEEF	Logs are forwarded in LEEF format.
	NetFlow	Logs are forwarded in NetFlow format.
	IPFIX	Logs are forwarded in IPFIX format.
	XML	Logs are forwarded in XML format.
Data Type		The type of log data that is forwarded.
Filter *(Optional)*		An optional local filter that defines which log data is forwarded. The local filter is only applied to the log data that matches the Log Forwarding rule. See Creating and Editing Local Filters (page 187).

Note – A local filter selected in a Log Forwarding rule is only used in forwarding the log data that matches the rule. If you have defined any Filters for forwarding syslog data to an external syslog server in the LogServerConfiguration.txt file, both the local filter in the Log Forwarding rule and the Filter(s) defined in LogServerConfiguration.txt are applied to the log data that matches the Log Forwarding rule.

7. Repeat Step 5-Step 6 to define more Log Forwarding rules.
8. Click **OK**.

What's Next?

▶ Continue by Enabling Logging for Traffic That You Want to Monitor (page 318).

Enabling Logging for Traffic That You Want to Monitor

To generate log data that can be forwarded to an external host, you must enable logging for the traffic that you want to monitor.

▼ To enable logging for the traffic that you want to monitor

1. Select **Configuration→Configuration→Security Engine**. The Security Engine Configuration view opens.
2. Expand **Policies** and browse to the type of policy you want to edit.
3. Right-click the Firewall, Layer 2 Firewall, or IPS Policy and select **Edit Firewall Policy**, **Edit Layer 2 Firewall Policy**, or **Edit IPS Policy**. The policy opens for editing.
4. Switch to the **IPv4 Access** or **IPv6 Access** tab and edit the rule that allows the traffic that you want to monitor.
   - If there is no Access rule for the traffic that you want to monitor, create a new Access rule as instructed in Editing Access Rules (page 672).
5. Double-click the **Logging** cell. The logging options dialog opens.

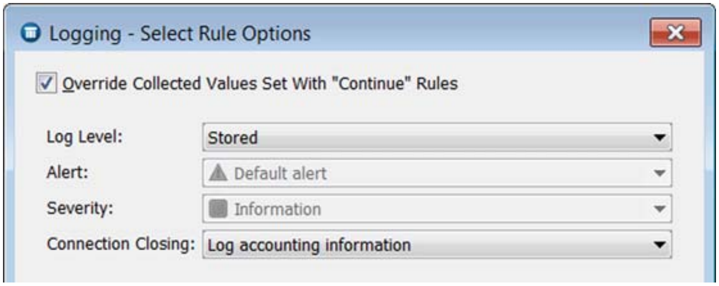

Logging - Select Rule Options

☑ Override Collected Values Set With "Continue" Rules

Log Level: Stored

Alert: Default alert

Severity: Information

Connection Closing: Log accounting information

6. Select **Override Collected Values Set With "Continue" Rules**.
7. Select **Stored** or **Essential** as the **Log Level**.
8. *(Optional)* If you want to forward logs in the NetFlow or IPFIX format, select **Connection Closing** for **Log Accounting Information**.

9. Click **OK**.

What's Next?

- ▶ If the Log Server and the target host to which you want to forward logs are separated by a firewall, continue by Creating an Access Rule Allowing Traffic to an External Host (page 319).
- ▶ If you are finished editing the policy, save and install the policy to start using the new configuration.

Creating an Access Rule Allowing Traffic to an External Host

If the external host and the Log Server are located in different networks separated by a Firewall or Layer 2 Firewall, you must modify the Firewall Policy or Layer 2 Firewall Policy to allow the traffic from the Log Server to the host.

▼ To create an Access rule allowing traffic to an external host

1. Select **Configuration→Configuration→Security Engine**. The Security Engine Configuration view opens.
2. Expand **Policies** and browse to the type of policy you want to edit.
3. Right-click the Firewall or Layer 2 Firewall policy and select **Edit Firewall Policy** or **Edit Layer 2 Firewall Policy**.
4. Switch to the **IPv4 Access** or **IPv6 Access** tab and add an Access rule with the following values:
   - **Source**: your Log Server
   - **Destination**: the target Host element
   - **Service**: UDP or TCP depending on the protocol used. The same Service and Port that was selected in the Log Forwarding rule must be selected here. See Defining Log Forwarding Rules (page 316).
   - **Action**: Allow
   - **Logging**: In most cases, we recommend setting the logging to **None**, since logging the log forwarding can create a loop (a log entry is sent to the target host, creating a log entry that is sent to the target host, creating a new log entry, and so on). If you want to log the log forwarding, create a local filter in the Log Forwarding rule to exclude logs related to forwarding. See Creating and Editing Local Filters (page 187).

What's Next?

- ▶ If you are finished editing the policy, save and install the policy to start using the new configuration.

Changing Log Server Configuration Parameters

Prerequisites: None

To configure the Log Server in detail, you can edit the `LogServerConfiguration.txt` file as explained in this section. The file is located on the Log Server machine in the folder `<`*`installation directory`*`>/data/`.

Normally, it is not necessary to configure the Log Server outside of the Management Client. However, under special circumstances, you may want more control over the way the Log Server behaves.

▼ To change the Log Server configuration parameters

1. Stop the Log Server:
   - If you run the Log Server as a service in Windows, you can stop it in the Windows Control Panel's **Services** list.
   - In Linux, run the script `<`*`installation directory`*`>`**`/bin/sgStopLogSrv.sh`**.
2. Open the `LogServerConfiguration.txt` file in a text editor.
3. Modify the parameter values as explained in the table below.
4. Save the changes.
5. Restart the Log Server.

All configuration parameters are listed in the following table. Note that not all parameters are included in the default configuration file. You may have to add some parameters manually.

Table 23.2 Log Server Configuration Parameters in LogServerConfiguration.txt file

Parameter name	Description
ARCHIVE_DIR_0	Directory that is used for storing the logs archived by the Log Data tasks. By default, ARCHIVE_DIR_0=${SG_ROOT_DIR}/data/archive. You can define up to 32 directories: ARCHIVE_DIR_0 … ARCHIVE_DIR_31.
AUDIT_ARCHIVE_DIR	Directory used for archiving audit logs. By default, ${SG_ROOT_DIR}/data/audit/archive.
AUDIT_DISK_LIMIT	The threshold for minimum available disk space for audit logs. If the free disk space goes below this limit, the Log Server considers the disk full and stops storing audit logs.
AUDIT_LOG_DIR	Directory used for audit logs. By default, ${SG_ROOT_DIR}/data/audit/log.
AUTO_ACK_ALERT_NB	The number of alerts that are automatically acknowledged if the active alert list is full (5000 by default). The alerts by lowest severity and oldest timestamp are acknowledged first.
DISK_THRESHOLD_IN_KBYTES	The threshold for minimum available disk space (in kilobytes). If the free disk space goes below this limit, the Log Server considers the disk full and stops storing log records (100000 by default).
SNMP_COMMUNITY	SNMP community string used for sending SNMP messages from the Log Server (public by default).

Table 23.2 Log Server Configuration Parameters in LogServerConfiguration.txt file (Continued)

Parameter name	Description
SNMP_ENTERPRISE_OID	SNMP Enterprise Object Identifier (OID) used for SNMP messages sent from the Log Server (.1.3.6.1.4.1.1369 by default).
LOG_BACKUP_DIR	Directory used for Log Server backup files. By default, ${SG_ROOT_DIR}/backups. The backup files must be moved to a separate media after creating a backup.
LOG_EXPORT_DIR	Directory used for storing the files exported by Log Data tasks. By default, ${SG_ROOT_DIR}/data/export.
LOG_FW_PORT	Log Server port that listens for connections from the Security Engines (3020 by default). Changing this value requires reinstalling the Log Server software.
LOG_LOGFILE_DIR	Directory used for storing the logfile.txt that logs the task scheduler operations. By default, ${SG_ROOT_DIR}/data.
LOG_QUERY_TIMEOUT	Timeout (in milliseconds) for queries in the Logs view (30000 by default).
LOG_SCRIPT_DIR	Directory for the scripts used in Log Data tasks. By default, ${SG_ROOT_DIR}/data/script.
LOG_SERVER_ADD	IP address of the Log Server. Changing this value requires reinstalling the Log Server software.
MAX_ACTIVE_ALERT_NB	Maximum number of active alerts. If there are more alerts, they are automatically acknowledged (50000 by default).
MGT_SERVER_ADD	IP address of the Management Server. Do not change this parameter value directly to the file. Instead, use the sgChangeMgtIPOnLogSrv.bat (or .sh) script to change this parameter value.
PHY_LOC	Log Server database location. By default, ${SG_ROOT_DIR}/data/db/logserver.
PHY_PORT	Log Server database port that the Log Server connects to (1314 by default).
SNMP_TRAP_RECEPTION_PORT	Defines the port used for receiving SNMP traps. The default port is UDP 162 in Windows and UDP 5162 in Linux. Note, that only the reception of SNMPv1 traps is supported.
SYSLOG_CONF_FILE	Configuration file for syslog data. By default, the file is stored in ${SG_ROOT_DIR}/data/fields/syslog_templates.
SYSLOG_EXPORT_FORMAT	The file format used for syslog forwarding. Either CEF, CSV, IPFIX, LEEF, NetFlow, or XML.

Table 23.2 Log Server Configuration Parameters in LogServerConfiguration.txt file (Continued)

Parameter name	Description
SYSLOG_FILTER_MATCH	Defines how many of the defined filters a log event must match to forward the event to the syslog server. The value “ALL” sends only events that match all defined filters. The value “ONE” sends all events that match any of the defined filters. The value “NONE” sends only events that match none of the defined filters.
SYSLOG_FILTER_TYPE	Defines how the defined filters are used for sending events to the syslog server. The value “KEEP” sends all the matching events. The value “DISCARD” sends only the events that do not match.
SYSLOG_MESSAGE_PRIORITY	The priority (0–191) of the syslog message is included at the beginning of each UDP packet (the default is 6). See RFC 3164.
SYSLOG_PORT	The target UDP port for sending syslog. The default port is 514.
SYSLOG_RECEPTION_PORT	The UDP port for receiving syslog. If this parameter has not been defined, the default port (514 for Windows or 5514 for Linux) is used. Note - In Linux the value of this parameter must always be set > 1024.
SYSLOG_RECEPTION_TCP_PORT	The TCP port for receiving syslog. If this parameter has not been defined, the UDP default port (514 for Windows and 5514 for Linux) is used. Note - In Linux the value of this parameter must be set > 1024.
SYSLOG_SERVER_ADDRESS	The IP address of the syslog server used for sending log events to syslog using the UDP protocol.
SYSLOG_TCP_SERVER_ADDRESS	The IP address of the syslog server used for sending log events to syslog using the TCP protocol.
SYSLOG_TCP_PORT	The target TCP port for sending syslog. The default port is 514.
SYSLOG_USE_DELIMITER	Defines whether to use double quotes (“) in syslog messages to delimit the field values. The default setting “ALWAYS_EXCEPT_NULL” uses double quotes only for non-empty fields. “NEVER” does not use delimiters. “ALWAYS” uses double quotes as delimiters for all empty and non-empty field values.

Forwarding Log Data to Syslog

Prerequisites: None

Log Servers can be configured to forward log data to external syslog servers. The log data to be forwarded can be selected by using filters just as any other log data in the Management Client.

Logs that are deleted by the Immediate Discard log pruning filters are not forwarded to a syslog server. Discard Before Storing pruning does not affect syslog forwarding.

> **Note – The Filters that you use in forwarding log data to a syslog server are also used if you configure the Log Server to forward log data to an external host. See Forwarding Log Data to an External Host (page 316).**

What's Next?

- ▶ Defining General Syslog Settings

Defining General Syslog Settings

▼ To define general syslog settings

1. Stop the Log Server.
   - If you run the Log Server as a service in Windows, you can stop it in the Windows Control Panel's **Services** list.
   - In Linux, run `<`*`installation directory`*`>/bin/sgStopLogSrv.sh`.
2. Create a text file on the Log Server that lists the fields to forward in the correct order.
   - The `<installation directory>/data/fields/syslog_templates/` directory contains example configuration files, which show you the correct syntax to use.
   - See Log Entry Fields (page 1212) for listings of the syslog forwarding names for the log fields.
3. Modify the `LogServerConfiguration.txt` file as shown in the table below. The file is located in `<`*`installation directory`*`>/data/`. For a list of syslog entry message categories, see Syslog Entries (page 1252).
4. Save the file and restart the Log Server.

Table 23.3 Log Server Configuration

Parameter	Value	Description
SYSLOG_CONF_FILE	<File name>	Path to the file you created in Step 2, which defines the fields that are forwarded and their order.
SYSLOG_EXPORT_AUDIT	YES NO	Defines whether to forward Audit entries to syslog.

Table 23.3 Log Server Configuration (Continued)

Parameter	Value	Description
SYSLOG_EXPORT_FORMAT	CEF CSV IPFIX LEEF NetFlow XML	Defines the file format used for syslog forwarding: CEF, CSV, IPFIX, LEEF, NetFlow, or XML.
SYSLOG_FILTER_MATCH	ALL	The log is forwarded if it matches all filters.
	ONE	The log is forwarded if it matches at least one filter.
	NONE	The log is forwarded if it does not match any of the filters.
SYSLOG_FILTER_TYPE	KEEP	The matching logs are sent to the syslog server.
	DISCARD	The matching logs are not sent to the syslog server.
SYSLOG_MESSAGE_PRIORITY	0-191 [a]	The priority of the syslog message is included at the beginning of each UDP packet (the default is 6). a) As defined in RFC 3164 (http://www.ietf.org/rfc/rfc3164.txt).
SYSLOG_PORT	[UDP port]	The target UDP port for sending syslog. The default port is 514.
SYSLOG_SERVER_ADDRESS	[IPv4 address]	The IP address of the syslog server for syslog sending over UDP. If left empty, transfers are not made.
SYSLOG_TCP_SERVER_ADDRESS	[IPv4 address]	The IP address of the syslog server for syslog sending over TCP. This parameter is not listed by default. If left empty, transfers are not made.
SYSLOG_TCP_PORT	[TCP port]	The target TCP port for sending syslog. The default port is 514. This parameter is not listed by default.

Table 23.3 Log Server Configuration (Continued)

Parameter	Value	Description
SYSLOG_USE_DELIMITER	ALWAYS_EXCEPT_NULL NEVER ALWAYS	Defines whether to use double quotes (“) in syslog messages to delimit the field values. The default setting “ALWAYS_EXCEPT_NULL” uses double quotes only for non-empty fields. “NEVER” does not use delimiters. “ALWAYS” uses double quotes as delimiters for all empty and non-empty field values.

What’s Next?

- ▶ If you want to filter the data sent to the syslog server, continue by Exporting Log Filters for Syslog Forwarding.
- ▶ If you do not want to use filters, continue by Creating an Access Rule Allowing Traffic to the Syslog Server (page 328).

Exporting Log Filters for Syslog Forwarding

You can use filters to select the data that is forwarded from the Log Server to the syslog server. First create the filter(s) as instructed in Filtering Data (page 183). You can then export and apply the filters separately to different kinds of logs.

Note – The Filters that you use in forwarding log data to a syslog server are also used if you configure the Log Server to forward log data to an external host. See Forwarding Log Data to an External Host (page 316).

▼ To export log filters for syslog forwarding

1. Select **Configuration→Configuration→Monitoring**.
2. Expand the **Other Elements** tree.
3. Browse to **Filters→All Filters** or **Filters→By Filter Tag** in the left panel.
4. Right-click the filter you want to export and select **Tools→Save for Command Line Tools** from the menu. The Save Filter dialog opens.
5. Select **Local Workstation** and click the **Browse** button to select where the filter is exported. The Export Data dialog opens.
   - To use the filter for filtering Firewall/VPN logs, select `<installation directory>`**`/data/syslog/Firewall`** as the directory.
   - To use the filter for filtering IPS logs, select `<installation directory>`**`/data/syslog/IPS`** as the directory.
   - To use the filter for filtering Layer 2 Firewall logs, select `<installation directory>`**`/data/syslog/L2FW`** as the directory.
   - To use the filter for filtering SSL VPN logs, select `<installation directory>`**`/data/syslog/SSL_VPN`** as the directory.

Note – If you cannot export the filter directly to these directories, export the filter (.flp file) to another location and then copy the filter manually to the `<installation directory>`/`data/syslog/Firewall`, `<installation directory>`/`data/syslog/IPS`, `<installation directory>`/`data/syslog/L2FW`, or `<installation directory>`/`data/syslog/SSL_VPN` directory.

6. Give the export file a name and click **OK** to export the filter.

What's Next?

▶ Continue by Configuring Syslog Filter Settings.

Configuring Syslog Filter Settings

Once you have exported the filter(s) for filtering the log data forwarded to the syslog server, you can define how the log data is matched with the filter(s). Add the filter settings to the `LogServerConfiguration.txt` file by setting values for the `SYSLOG_FILTER_TYPE=` and `SYSLOG_FILTER_MATCH=` parameters. Editing the file is explained in Defining General Syslog Settings (page 323). The parameters and their values are explained in Table 23.3.

The table below shows an example of filter settings.

Table 23.4 Example Filter Setting Combinations

SYSLOG_FILTER_TYPE=	SYSLOG_FILTER_MATCH=	Effect
KEEP	ALL	The logs matching all filters are sent to the syslog server.
KEEP	ONE	The logs matching at least one of the filters are sent to the syslog server.
KEEP	NONE	The logs matching none of the filters are sent to the syslog server.

Note – If you do not use filters for the data to be forwarded, remove the `SYSLOG_FILTER_TYPE=` **and** `SYSLOG_FILTER_MATCH=` **parameters from the** `<`*`installation directory`*`>/data/LogServerConfiguration.txt` **file.**

What's Next?

▶ Continue by Creating an Access Rule Allowing Traffic to the Syslog Server.

Creating an Access Rule Allowing Traffic to the Syslog Server

If the syslog server and the Log Server are located on different networks separated by a Firewall or Layer 2 Firewall, you must modify the Firewall Policy or Layer 2 Firewall Policy to allow the traffic. You may also need to define an appropriate NAT rule in the Firewall Policy.

▼ To create an Access rule allowing traffic to the syslog server

1. Select **Configuration→Configuration→Security Engine**. The Security Engine Configuration view opens.
2. Expand **Policies** and browse to the type of policy you want to edit.
3. Right-click the Firewall or Layer 2 Firewall policy and select **Edit Firewall Policy** or **Edit Layer 2 Firewall Policy**.
4. Switch to the **IPv4 Access** or **IPv6 Access** tab and add an Access rule with the following values:
   - **Source**: your Log Server
   - **Destination**: the syslog server
   - **Service**: Syslog (UDP) or Syslog (TCP) depending on the protocol used.
   - **Action**: Allow
   - **Logging**: In most cases, we recommend setting the logging to **None**, since logging syslog connections can create a loop (a log entry is sent to the syslog server, creating a log entry that is sent to the syslog server, creating a new log entry, and so on). If you want to log the syslog connections, make sure they are filtered out from syslog forwarding as explained in Configuring Syslog Filter Settings (page 327). Make sure the filter accounts for error conditions, for example, the ICMP “Destination unreachable” messages that are generated if the syslog destination server goes down.

What’s Next?

▶ If you are finished editing the policy, save and install the policy to start using the new configuration.

CHAPTER 24

CONFIGURING ADDITIONAL SMC SERVERS

You can install several Management Servers and Log Servers to provide high availability for the Stonesoft Management Center (SMC).

The following sections are included:

About Additional SMC Servers

Prerequisites: None

Although the Security Engines always work independently without the SMC according to their installed configuration, configuration changes and system monitoring are not possible without a Management Server and a Log Server. The Management Server in particular is a critical component, as it is the only place where the full configuration information is stored.

You can install additional Management Servers and Log Servers. The high-availability (HA) solution includes automatic incremental replication of the configuration data stored on the Management Server. This way, manual intervention is minimized, and the system can be fully managed and monitored without going through a manual re-installation and backup restoration process.

Installing Additional Management Servers

Prerequisites: Your organization has purchased a license that allows activating this feature.

To use additional Management Servers, you must have a special Management Server license that includes the high-availability features. The license is a combined license for all Management Servers and it must list the IP addresses of all Management Servers.

Configuration Overview

1. Create the element as explained in Defining an Additional Management Server Element (page 331).
2. License the servers as explained in Installing Licenses for Additional Management Servers (page 332).
3. Allow communications to the new server through firewalls as necessary as explained in Creating Access Rules for Additional Management Servers (page 333).
4. Install the software on the target server as explained in Installing Additional Management Server Software (page 333).

Defining an Additional Management Server Element

Any new Management Server elements you add to an existing Management Center are considered to be additional Management Servers. You can set up several additional Management Servers.

▼ To define an additional Management Server element

1. Select **Monitoring→System Status**. The System Status view opens.
2. Right-click **Servers** and select **New→Management Server**. The Management Server Properties dialog opens.

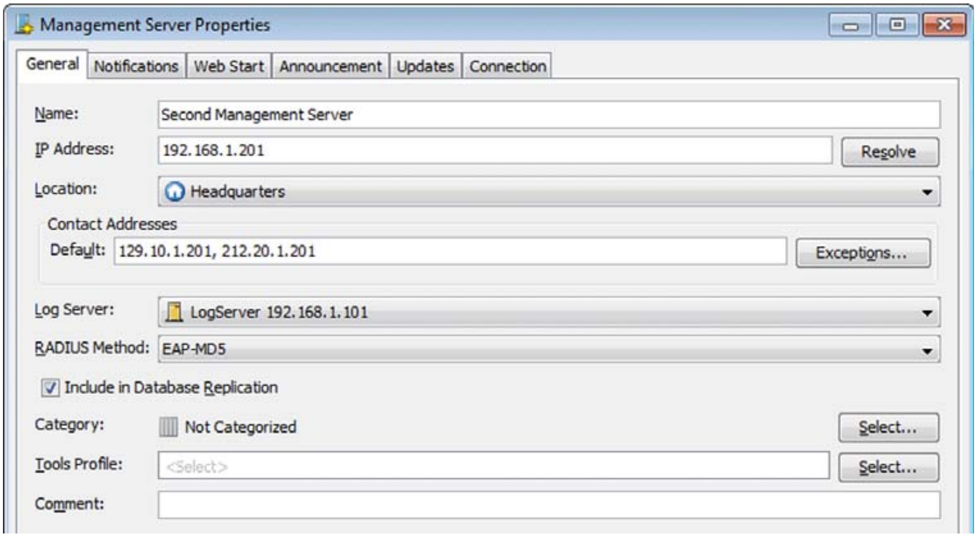

Name: Second Management Server

IP Address: 192.168.1.201

Location: Headquarters

Contact Addresses

Default: 129.10.1.201, 212.20.1.201

Log Server: LogServer 192.168.1.101

RADIUS Method: EAP-MD5

☑ Include in Database Replication

Category: Not Categorized

Tools Profile: <Select>

Comment:

3. Type in the **Name** and **IP Address**.
4. If NAT is used to translate addresses in communications between this server and other system components, define the **Location** and **Contact Address**. See Defining Locations (page 66) and Defining Contact IP Addresses (page 66) for more information.
5. Select the **Log Server** to which any alerts generated by this Management Server are sent.
6. If you use a RADIUS server for authenticating administrators, select the **RADIUS Method** that is used in the authentication.
7. (*Optional*) Deselect **Include in Database Replication** if you want to temporarily disable automatic database replication to prevent excessive failed replication attempts when you know the Management Server will be unreachable.
   - You can disable automatic database replication when, for example, changing the Management Server's hardware or IP address, or for troubleshooting purposes.
   - By default, the Management database changes are replicated incrementally at regular intervals, and no configuration is necessary.

Caution – Re-enable automatic database replication as soon as possible once the Management Server is working normally. Leaving the automatic database replication disabled after the Management Server becomes reachable causes serious mismatches in the Management Server databases.

8. (*Optional*) If you want to send generated reports as e-mail directly from the Management Server, switch to the **Notifications** tab and enter the address of the **SMTP Server** and the **Sender Address** that is shown as the e-mail sender.
   - Remember to allow these connections in the Firewall policy if necessary.
9. (*Optional*) If you want to enable Web Start users to access the Web Start files through a web browser at the address that the additional Management Server uses, switch to the **Web Start** tab and configure the settings (see Activating Web Start on the Management Server (page 309)).
10. Click **OK**.

What's Next?

▶ Continue by Installing Licenses for Additional Management Servers.

Related Tasks

▶ Disabling and Enabling Automatic Database Replication (page 339)

▶ Synchronizing Management Databases Manually (page 339)

Installing Licenses for Additional Management Servers

Using additional Management Servers requires a special combined license that lists the IP addresses of all Management Servers within the same SMC. After receiving the proof-of-license (POL) code from Stonesoft, generate the License file at the Stonesoft website at http://www.stonesoft.com/en/customer_care/licenses/. For more information on generating licenses, see Generating New Licenses (page 1047).

▼ To install a license for an additional Management Server

1. Select **File→System Tools→Install Licenses**. The Install License File(s) dialog opens.
2. Browse to the License file on your local workstation and click **Install**.
3. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
4. Browse to **Licenses→Unlicensed Components→Servers** and make sure the additional Management Server is *not* listed there.
   - If the server is listed, check that the IP address information is the same in the license and the additional Management Server element.

What's Next?

▶ Continue by Creating Access Rules for Additional Management Servers (page 333).

Creating Access Rules for Additional Management Servers

The Firewall Template policy contains rules that allow connections between the Firewall and all Management Servers and Log Servers that the Firewall connects to. However, if other components need to connect through a Firewall to the additional Management Servers, you must add rules that allow this traffic in the Firewall's policy. You may also have to add NAT rules.

The rules in the Firewall Template policy are IPv4 Access rules that allow the system communications. See Default Communication Ports (page 1167) for information on ports and default Service elements for system communications. For more information on adding Access rules, see Editing Access Rules (page 672).

What's Next?

▶ After adding the Access rules, continue by Installing Additional Management Server Software.

Installing Additional Management Server Software

Additional Management Servers are installed using the Management Center Installation Wizard. To install the software, you need the Management Center installation files. See Obtaining the SMC Installation Files (page 1057).

▼ To install additional Management Server software

1. Start the installation in one of the following ways:
   - **From a .zip file**: unzip the file and run **`setup.exe`** on Windows or **`setup.sh`** on Linux.
   - **From a DVD**: insert the installation DVD and run the `setup` executable from the DVD:

Operating System	Path to Executable
Windows 32-bit	\Stonesoft_SW_Installer\Windows\setup.exe
Windows 64-bit	\Stonesoft_SW_Installer\Windows-x64\setup.exe
Linux 32-bit	/Stonesoft_SW_Installer/Linux/setup.sh
Linux 64-bit	/Stonesoft_SW_Installer/Linux-x64/setup.sh

Note – If the DVD is not automatically mounted in Linux, mount the DVD with the following command: `mount /dev/cdrom /mnt/cdrom`

2. Proceed according to instructions in the Installation Wizard until you are prompted to select which components you want to install.

Note – We do not recommend selecting `C:\Program Files\Stonesoft\Management Center` as the installation directory in Windows. Selecting `C:\Program Files\Stonesoft\Management Center` as the installation directory creates an additional `C:\ProgramData\Stonesoft\Management Center` folder, which duplicates some of the folders in the installation directory. Some of the program data is also stored in the `C:\ProgramData\Stonesoft\Management Center` folder.

3. If you also want to install a Log Server and a local Management Client on this computer, leave **Typical** selected and click **Next**. Otherwise, select **Custom**, select the components you want to install and click **Next**.
4. Select the IP address of the Management Server from the list or type it in.
   - This must be the IP address defined for the corresponding Management Server element.
5. Type in the IP address of the Log Server for sending alerts.
6. Select **Install as an Additional Management Server for High Availability**.
7. Click **Next** and follow the instructions to start the installation. A login prompt for Replication opens.
8. Log in using an unrestricted administrator account. The Management Server Selection dialog opens.

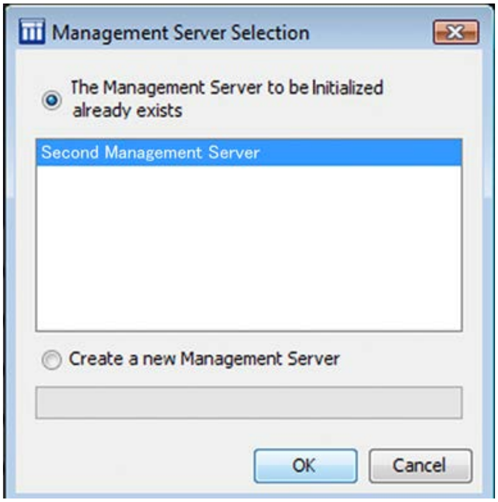

Management Server Selection

● The Management Server to be Initialized already exists

Second Management Server

○ Create a new Management Server

9. Select the correct Management Server from the list and click **OK**. The databases are synchronized.

Note – If the synchronization fails for some reason (such as a network connection problem), run the `sgOnlineReplication` **script on the additional Management Server when connectivity is restored.**

What's Next?

▶ The additional Management Server configuration is complete. You can view replication information in the Info panel when you select the Management Server.

Installing Additional Log Servers

Prerequisites: None

Additional Log Servers can be used to allow continued monitoring of the system if one Log Server fails. Alert escalation proceeds normally, new logs can be browsed, and the engine status and statistics can be examined. However, the log data is not automatically replicated between the Log Servers, so some log data is always unavailable during outages. We recommend that you select the same Log Server for all engines in the same location. If different engines send events to different Log Servers, it is not be possible to correlate events that are detected by different engines, as none of the Log Servers sees all of the events.

Any Log Server can be used both as the main Log Server for some components and as a backup Log Server for one or more other Log Servers. However, you must consider the load of the Log Servers before you set up an actively used Log Server as a backup Log Server, to avoid overloading when a fail-over occurs.

You can set up additional Log Servers with normal Log Server licenses. A separate license for each Log Server is required, even if the Log Server is used only as a backup.

The following overview can be used to install a new Log Server that works as a backup for another Log Server. However, you can alternatively define any existing Log Servers to be used as backups for some other Log Server just by completing Step 3 and refreshing the policies of engines that send their data to the Log Server.

Configuration Overview

1. Add an element for the additional Log Server as explained in Creating Additional Log Server Elements (page 336).
2. License the additional Log Server as explained in Installing Licenses for Additional Log Servers (page 336).
3. Define the Log Server as a backup Log Server for some other Log Server as explained in Setting a Log Server as a Backup Log Server (page 337).
4. Make sure the communications between system components and the additional Log Server are allowed as explained in Creating Access Rules for Additional Log Servers (page 337).
5. Install the Log Server as explained in Installing Additional Log Server Software (page 338).

Creating Additional Log Server Elements

You can set up several backup Log Servers for each Log Server.

▼ To create an additional Log Server element

1. Select **Monitoring→System Status**. The System Status view opens.
2. Right-click **Servers** and select **New→Log Server**. The Log Server Properties dialog opens.

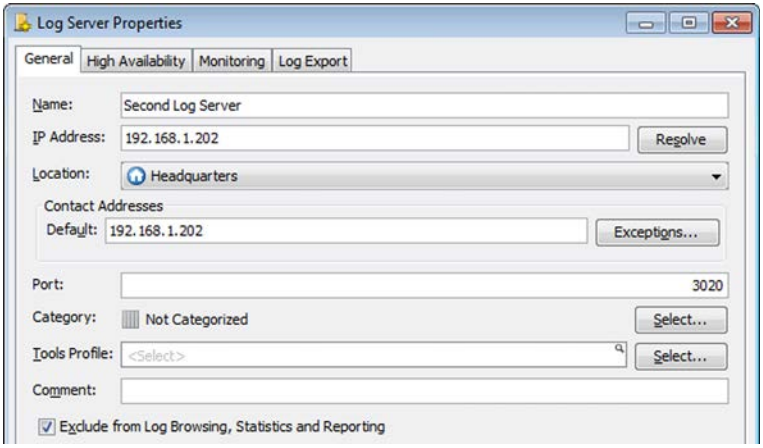

Name: Second Log Server

IP Address: 192.168.1.202

Location: Headquarters

Contact Addresses

Default: 192.168.1.202

Port: 3020

Category: Not Categorized

Tools Profile: <Select>

Comment:

☑ Exclude from Log Browsing, Statistics and Reporting

3. Enter the **Name** and **IP Address** of the server.
4. If NAT is used to translate addresses in communications between this server and other system components, define the **Location** and **Contact Address**. See Defining Locations (page 66) and Defining Contact IP Addresses (page 66) for more information.
5. (*Optional*) Change the Log Server's TCP **Port Number** if necessary. We recommend always using the default port 3020 if possible.

Note – If you want to use a non-standard port for the Log Server, you must manually add rules to allow communications using the new port from the engines to the Log Server even when using the Firewall Template policy.

6. (*Optional*) Select **Exclude from Log Browsing, Statistics and Reporting** if you do not want the Log Server to gather statistical information for monitoring and reports, for example, when the Log Server is not used for daily operations.

What's Next?

- ▶ Continue by Installing Licenses for Additional Log Servers.

Installing Licenses for Additional Log Servers

Each Log Server requires a separate license, even if it is only used as a backup for some other Log Server. After receiving the proof-of-license (POL) code from Stonesoft, generate the License file at the Stonesoft website at http://www.stonesoft.com/en/customer_care/licenses/. For more information on generating licenses, see Generating New Licenses (page 1047).

▼ To install a license for an additional Log Server

1. Select **File→System Tools→Install Licenses**. The Install License File(s) dialog opens.
2. Browse to the License file on your local workstation and click **Install**.
3. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
4. Browse to **Licenses→Unlicensed Components→Servers** and make sure the additional Log Server is *not* listed there.
   - If the server is listed, check that the IP address information is the same in the license and the Log Server element.

What's Next?

▶ Continue by Setting a Log Server as a Backup Log Server.

Setting a Log Server as a Backup Log Server

You can select several backup Log Servers for each Log Server. The same Log Server can be the main Log Server for some components and a backup Log Server for other components at the same time. You can also set Log Servers to be backup Log Servers for each other so that whenever one fails, another Log Server is used. See Selecting Backup Log Servers (page 315).

If Domain elements have been configured, a Log Server and its backup Log Server(s) must belong to the same Domain.

Caution – If the log volumes are very high, make sure that the backup Log Server can handle the traffic load in fail-over situations.

What's Next?

▶ Continue by Creating Access Rules for Additional Log Servers.

Creating Access Rules for Additional Log Servers

The Firewall Template contains rules that allow connections between a Firewall and all Management and Log Servers the Firewall connects to. However, if other components need to connect through a Firewall to the additional Log Server, you must add rules that allow this traffic in the Firewall's policy, along with possible NAT rules when necessary.

The rules in the Firewall Template are IPv4 Access rules that allow the system communications. See Default Communication Ports (page 1167) for information on ports and default Service elements for system communications. For more information on adding Access rules, see Editing Access Rules (page 672).

What's Next?

▶ After adding the Access rules, continue the installation of a new Log Server in Installing Additional Log Server Software.

Installing Additional Log Server Software

Additional Log Servers are installed using the Management Center Installation Wizard in the same way as any other Log Servers. This section presents an outline of the installation. For detailed installation instructions, see the *Management Center Installation Guide*.

To install the software, you need the Management Center installation files, see Obtaining the SMC Installation Files (page 1057).

▼ To install additional Log Server software

1. Start the installation in one of the following ways:
   - **From a .zip file**: unzip the file and run **`setup.exe`** on Windows or **`setup.sh`** on Linux.
   - **From a DVD**: insert the installation DVD and run the `setup` executable from the DVD:

Operating System	Path to Executable
Windows 32 bit	\Stonesoft_SW_Installer\Windows\setup.exe
Windows 64 bit	\Stonesoft_SW_Installer\Windows-x64\setup.exe
Linux 32 bit	/Stonesoft_SW_Installer/Linux/setup.sh
Linux 64 bit	/Stonesoft_SW_Installer/Linux-x64/setup.sh

Note – If the DVD is not automatically mounted in Linux, mount the DVD with the following command: `mount /dev/cdrom /mnt/cdrom`

2. Proceed according to instructions in the Installation Wizard until you are prompted to select which components you want to install.
3. Select **Custom**, and select **Log Server**.
4. Finish the installation according to instructions in the Installation Wizard. You must certify the Log Server before the Log Server can connect to the Management Server.

The additional Log Server is now installed and should be shown in green (Online) in the System Status view.

The logs are not automatically copied between the Log Servers by default, as the volumes can be quite large. You can set up separate scheduled or manually run Tasks if you want to copy logs between the servers. See Getting Started with Log Data Management (page 1016).

What's Next?

- ▶ If you installed the additional Log Server to use as a backup Log Server, continue by Selecting Backup Log Servers (page 315).
- ▶ Otherwise, the additional Log Server installation is complete.

Disabling and Enabling Automatic Database Replication

Prerequisites: Installing Additional Management Servers

By default, changes in the configuration are automatically replicated between all Management Servers. Management database replication is incremental and continuous; only the changed parts of the database are replicated and the replication is done in real time.

Note – The Management database is not synchronized automatically between the Management Servers after a Management Server upgrade. You must synchronize the database manually after the upgrade. See Synchronizing Management Databases Manually**.**

Disabling the automatic database replication is not recommended unless you have a specific need to do so. For example, you may need to prevent excessive failed replication attempts when you know the Management Server will be unreachable for a long time, such as changing the Management Server's hardware or the Management Server's IP address.

Caution – Re-enable automatic database replication in the Management Server's properties as soon as the Management Server is reachable. You must synchronize the Management database manually once the Management Server is working normally. See Synchronizing Management Databases Manually**.**

▼ To disable or enable automatic database replication

1. Double-click an additional Management Server. The Management Server Properties dialog opens.
2. Deselect/select **Include in Database Replication** to disable/enable automatic replication.
3. Click **OK** to close the Management Server Properties dialog.

Synchronizing Management Databases Manually

Prerequisites: Installing Additional Management Servers

You must synchronize the configuration information manually between all Management Servers through the Management Client after upgrading the Management Servers. You can also synchronize the databases between the Management Servers in situations when automatic Management database replication cannot be used.

The manual Management database synchronization is primarily meant for resynchronizing the databases after upgrading the SMC. We do not recommend using the manual database synchronization unless you have a specific need to do so.

▼ To synchronize the management databases manually

1. Select **File→System Tools→Control Management Servers**. The Control Management Servers dialog opens.

Illustration 24.1 Control Management Servers - Synchronizing Databases Manually

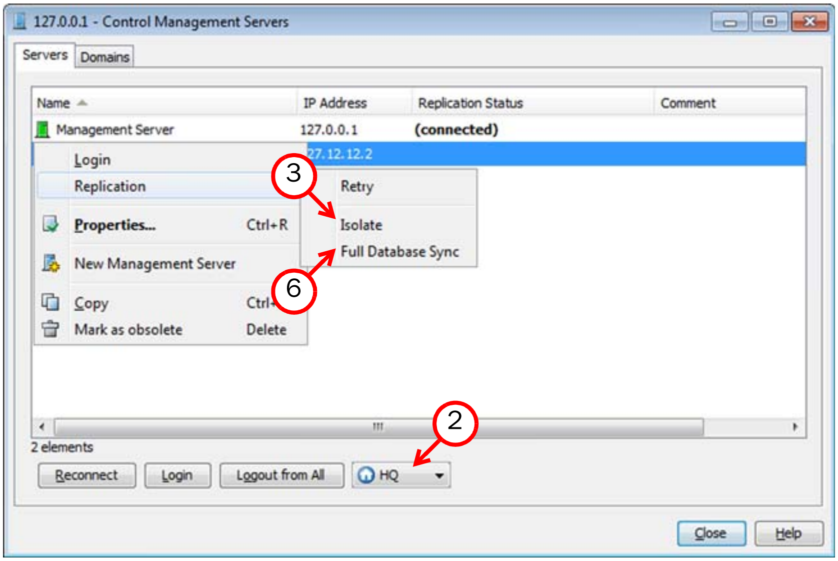

2. If you are logged in to a different Management Server than the one that you are selecting for replication, select the Location from which to send the command. This ensures that the command is sent to the correct Contact Address for the Management Server.
3. Right-click the additional Management Server and select **Replication→Isolate**. You are prompted to confirm the isolation.
4. Click **Yes**. The additional Management Server's database is made incompatible with the other Management Servers' databases.
5. Click **OK** to acknowledge the completion of the isolation.
6. Right-click the isolated Management Server and select **Replication→Full Database Sync**. You are prompted to confirm the synchronization.
7. Click **Yes**. All existing configurations on the additional Management Server are overwritten.
8. Click **OK** to acknowledge the completion of the synchronization and wait for the Management Server to restart.
9. If you need to synchronize more than one additional Management Server, repeat Step 3-Step 8 for each additional Management Server.
10. Click **Close** to close the Control Management Servers dialog.

CHAPTER 25

RECONFIGURING THE MANAGEMENT CENTER AND ENGINES

This section includes tasks related to configuring additional settings for Management Servers, changing your Management Center hardware platform or the IP addresses used in system communications, as well as changing the role of Security Engines.

The following sections are included:

Modifying a Management Server Element

Prerequisites: None

One Management Server element is automatically created during SMC installation. Check and modify the settings as necessary. You can rename the Management Server element freely, but you cannot delete the Management Server that the SMC is connected to. To define additional Management Servers for high availability, see Installing Additional Management Servers (page 330).

To change the management database password that was automatically created during the installation, see Changing the Management Database Password (page 343).

▼ To modify a Management Server element

1. Select **Monitoring→System Status** from the menu. The System Status view opens.
2. Expand **Servers**.
3. Double-click the **Management Server** element you want to modify. The Management Server Properties dialog opens.

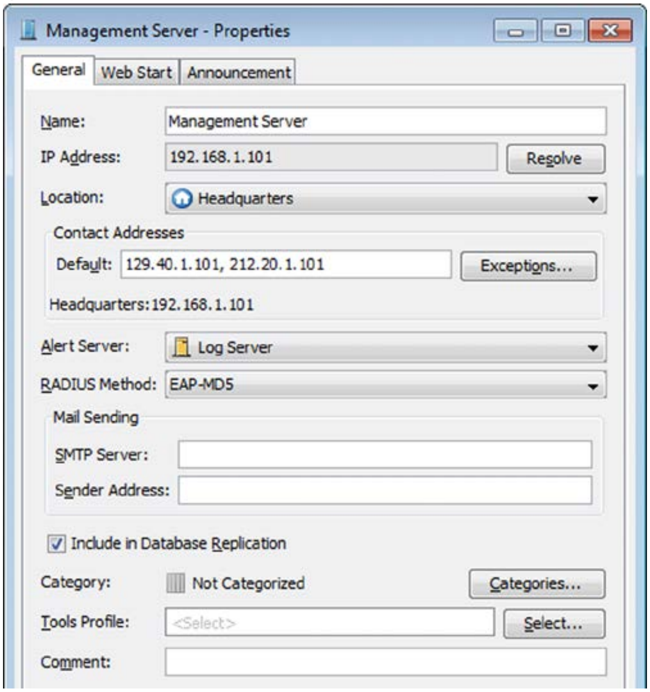

Name: Management Server

IP Address: 192.168.1.101

Location: Headquarters

Contact Addresses

Default: 129.40.1.101, 212.20.1.101

Headquarters: 192.168.1.101

Alert Server: Log Server

RADIUS Method: EAP-MD5

Mail Sending

SMTP Server:

Sender Address:

☑ Include in Database Replication

Category: Not Categorized

Tools Profile: <Select>

Comment:

4. *(Optional)* Change the **Name** of the Management Server.
5. *(Optional)* Add a **Location** and a **Contact Address** if there is a NAT device between a Security Engine and the Management Server so that they cannot connect directly to the IP address defined for the interface. See Defining Contact IP Addresses (page 66) for more information.
6. Select the **Alert Server** to which any alerts generated by this Management Server are sent.
7. If you use a RADIUS server for authenticating administrators, select the **RADIUS Method** that is used in the authentication.

8. If you want to send generated reports as e-mail directly from the Management Server, enter the address of the **SMTP Server** and the **Sender Address** that is shown as the e-mail sender. Remember to allow these connections in the firewall policy if necessary. See E-Mailing Reports (page 182) for more information.

9. (*Optional*) Deselect **Include in Database Replication** to temporarily disable automatic database replication to prevent excessive failed replication attempts when you know the Management Server will be unreachable.
   - For example, for troubleshooting purposes, while changing the Management Server's hardware, or changing the Management Server's IP address.
   - By default, the Management Server database changes are replicated incrementally at regular intervals.

Caution – Re-enable automatic database replication as soon as possible once the Management Server is working normally. Leaving the automatic database replication disabled after the Management Server becomes reachable causes serious mismatches in the Management Server databases.

10. (*Optional*) To enable Web Start users to start the Management Client through a web browser at the IP address that the Management Server uses, switch to the **Web Start** tab and proceed as described in Activating Web Start on the Management Server (page 309).

11. *(Optional)* If you want to use a Tools Profile with the Management Server, click **Select** and select the **Tools Profile**.

12. Click **OK** to apply the changes.

Changing the Management Database Password

Prerequisites: None

The Management Server contains a database for storing the configuration information. The password for the database is automatically created during the installation of the Management Server. In some rare cases, you may need this password to complete an operation. In these cases, you can change the database password.

▼ To change the management database password

1. Select **File→System Tools→Password→Change Database Password** in the menu. The Change Database Password dialog opens.

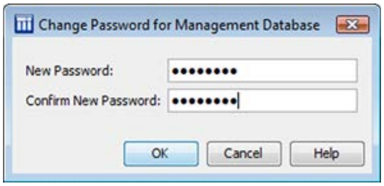

2. Enter the password in the **New Password** and the **Confirm New Password** fields.

3. Click **OK** to confirm the change.

Changing the Management Platform

Prerequisites: None

If you want to change the hardware or the operating system that the Management Server or Log Server components run on, follow the procedure below. For any Web Portal Server, install the new Web Portal Server as a fresh installation as instructed in the *Management Center Installation Guide*.

▼ To change the Management Platform

1. Take backups of your Management Server(s) and your Log Server(s) as instructed in Creating Backups (page 1009).
   - The Web Portal Server has no information to back up.
2. (*Multiple Management Server environments only*) Open the properties of the Management Server as instructed in Modifying a Management Server Element (page 342). Deselect **Include in Database Replication** to temporarily disable automatic database replication to prevent excessive failed replication attempts.

Caution – Re-enable automatic database replication as soon as possible once the Management Server is working normally. Leaving the automatic database replication disabled after the Management Server becomes reachable causes serious mismatches in the Management Server databases.

3. Install new Management Server(s) and Log Server(s) on the new platform(s). See the *Management Center Installation Guide* for instructions.
4. Restore the Management Server and Log Server backups from the old environment to the new installation. See Restoring Backups (page 1011).
5. (*Hardware change*) Shut down and disconnect the old Management Server and/or Log Server hardware.
6. (*Hardware change*) Connect the new Management Server and/or Log Server to the network environment.
7. (*Multiple Management Server environments only*) Open the properties of the Management Server as instructed in Modifying a Management Server Element (page 342). Select **Include in Database Replication** to re-enable automatic database replication.

What's Next?

▶ (*Multiple Management Server environments only*) If the replication status in the Info panel indicates a problem with database replication, synchronize the Management databases as instructed in Synchronizing Management Databases Manually (page 339).

Changing SMC IP Addressing

Prerequisites: None

The following instructions explain how you can change IP addresses without losing management connectivity. When you change IP addressing, other connections between the different components may be temporarily lost. You must make sure that the connections return to normal after the IP address changes.

What's Next?

▶ If you want to change IP addresses and the Management Server and Log Server are on the same machine, proceed to Changing IP Addresses of Combined Management/Log Servers (page 347).

▶ Otherwise, proceed to Changing the Management Server IP Address or Changing the Log Server IP Address (page 346) depending on which component's IP address you want to change.

Changing the Management Server IP Address

This section explains how to change the Management Server's IP address if the Management Server and the Log Server are on different machines. If your Management Server and Log Server are on the same machine, see Changing IP Addresses of Combined Management/Log Servers (page 347).

Before changing the Management Server's IP address, we recommend making a backup of the Management Server and the Log Server as instructed in Creating Backups (page 1009).

Note – If any Firewalls between the Management Server and other components do not use a policy based on the Firewall Template, check that they allow all the necessary connections.

▼ To change the Management Server's IP address

1. Request an IP address change for the Management Server license from Stonesoft's online License Center at https://my.stonesoft.com/managelicense.do.
2. Add Firewall IPv4 Access rules (and possibly NAT rules) that allow policy upload connections from the new IP addresses to the Firewall as instructed in Editing Access Rules (page 672) and refresh the Firewall Policies.
   - The services needed for the communications between the different components are explained in Default Communication Ports (page 1167).
3. (*Multiple Management Server environments only*) Open the properties of the Management Server as instructed in Modifying a Management Server Element (page 342). Deselect **Include in Database Replication** to temporarily disable automatic database replication to prevent excessive failed replication attempts.

Caution – Re-enable automatic database replication as soon as possible once the Management Server is working normally. Leaving the automatic database replication disabled after the Management Server becomes reachable causes serious mismatches in the Management Server databases.

4. Stop the Management Server and all Log Server services.

5. Change the IP address of the host server in the operating system.
6. On the Management Server, run the command `sgChangeMgtIPOnMgtSrv` <new Management Server IP address>.
7. On all Log Servers, run the command `sgChangeMgtIPOnLogSrv` <new Management Server IP address>.
8. Start the Management Server and Log Server services and log in using the Management Client.
9. Import the new Management Server license when prompted.
10. Remove the Firewall IPv4 Access rules that you created in Step 2 and refresh the Firewall Policies.
    - After running the IP address change scripts, the Alias elements in the inherited rules translate to the right IP addresses.
11. (*Multiple Management Server environments only*) Open the properties of the Management Server as instructed in Modifying a Management Server Element (page 342). Select **Include in Database Replication** to re-enable automatic database replication.

What's Next?

▶ (*Multiple Management Server environments only*) If the replication status in the Info panel indicates a problem with database replication, synchronize the Management databases as instructed in Synchronizing Management Databases Manually (page 339).

Changing the Log Server IP Address

When you change the Log Server's IP address, the traffic between the Log Server and the engines is interrupted and the logs are spooled on the engines. Changing the IP address may also mean that the transfer of engine status and statistics information is temporarily interrupted.

Before changing the Log Server's IP address, we recommend making a backup of the Management Server and the Log Server as instructed in Creating Backups (page 1009).

▼ To change the Log Server's IP address

1. Request a license binding change for the Log Server if the license is bound to the Log Server's IP address from Stonesoft's online License Center at https://my.stonesoft.com/managelicense.do and import the new license.
2. Edit the `<`*`installation directory`*`>/data/LogServerConfiguration.txt` file on the Log Server and update the Log Server IP address. See Changing Log Server Configuration Parameters (page 320).
3. In the Management Client, open the Log Server properties and update the Log Server IP address.
4. Stop and restart the Log Server service.
5. Refresh the policies of all engines that send data to the Log Server.

Changing IP Addresses of Combined Management/Log Servers

This section explains how to change the Management Server's IP address if the Management Server and the Log Server are on the same machine. If your Management Server and Log Server are on separate machines, see Changing the Management Server IP Address (page 345) and Changing the Log Server IP Address.

When you change the Log Server's IP address, the traffic between the Log Server and the engines is interrupted and the logs are spooled on the engines. Changing the IP address may also mean that the transfer of engine status and statistics information is temporarily interrupted.

Before changing the IP addresses, we recommend making a backup of the Management Server and the Log Server as instructed in Creating Backups (page 1009).

Note – If any Firewalls between the Management Server and other components do not use a policy based on the Firewall Template, check that they allow all the necessary connections.

▼ To change the IP Address of a combined Management Server and Log Server

1. Request a license binding change to the new IP address for the Management Server license, and also the Log Server if the license is bound to an IP address from Stonesoft's online License Center at https://my.stonesoft.com/managelicense.do.
2. Add Firewall IPv4 Access rules (and possibly NAT rules) that allow policy upload connections from the new IP addresses to the Firewall as instructed in Editing Access Rules (page 672) and refresh the firewall policies.
   - The services needed for the communications between the different components are explained in Default Communication Ports (page 1167).
3. (*Multiple Management Server environments only*) Open the properties of the Management Server as instructed in Modifying a Management Server Element (page 342). Deselect **Include in Database Replication** to temporarily disable automatic database replication to prevent excessive failed replication attempts.

Caution – Re-enable automatic database replication as soon as possible once the Management Server is working normally. Leaving the automatic database replication disabled after the Management Server becomes reachable causes serious mismatches in the Management Server databases.

4. Stop the Management Server and Log Server services.
5. Change the IP address of the host server in the operating system.
6. Run the `sgChangeMgtIPOnMgtSrv` script on the Management Server (see Command Line Tools (page 1145)).
7. Run the `sgChangeMgtIPOnLogSrv` script on the Log Server (see Command Line Tools (page 1145)).
8. Edit the `<`*`installation directory`*`>/data/LogServerConfiguration.txt` file on the Log Server and update the Log Server IP address (see Changing Log Server Configuration Parameters (page 320)).
9. Start the Management Server service and log in using the Management Client.
10. Import the new licenses when prompted.

11. Open the Log Server properties and update the IP address.

12. Start the Log Server service.

13. Remove the Firewall IPv4 Access rules that you created in Step 2 and refresh the Firewall Policies.
    - After running the IP address change scripts, the Alias elements in the inherited rules translate to the right IP addresses.

14. (*Multiple Management Server environments only*) Open the properties of the Management Server as instructed in Modifying a Management Server Element (page 342). Select **Include in Database Replication** to re-enable automatic database replication.

What's Next?

▶ (*Multiple Management Server environments only*) If the replication status in the Info panel indicates a problem with database replication, synchronize the Management databases as instructed in Synchronizing Management Databases Manually (page 339).

Creating a New Internal ECDSA Certificate Authority

The Elliptic Curve Digital Signature Algorithm (ECDSA) is a digital signature algorithm that uses elliptic curve cryptography. You can switch to using an Internal ECDSA Certificate Authority to sign the certificates that components use to identify each other in system communications. Switching to an Internal ECDSA Certificate Authority enables 256-bit encryption on the Management Server for connections between the Management Server and Security Engines.

When you create a new Internal ECDSA Certificate Authority, system components gradually start using the new Internal ECDSA CA to sign certificates. Initially, the new Internal ECDSA CA is in the "Created for Different Certificate Type" state. When the new Internal ECDSA CA is ready to begin signing certificates, it changes to the “Ready to Use for Different Certificate Type” state. At first, only Management Server certificates are signed by the new Internal ECDSA CA. Certificates for other components are signed by whichever Internal CA is currently used by the Management Server. When all Management Servers are using the new Internal ECDSA CA, the new Internal ECDSA CA changes to the "Active" state.

Caution – When you start using an ECDSA Certificate Authority, you must recertify all SMC servers and you may also need to make initial contact between the engines and the Management Server. Engines with versions lower than 5.5 and SSL VPN gateways cannot communicate with the SMC when an ECDSA Certificate Authority is used.

▼ To create a new Internal ECDSA Certificate Authority

1. Select **Configuration→Configuration→Administration** from the menu. The Administration Configuration view opens.
2. Browse to **Other Elements→Internal Certificate Authorities**.
3. Click the Tools icon in the toolbar and select **Create New Internal ECDSA Certificate Authority**. You are prompted to confirm that you want to create a new Internal ECDSA Certificate Authority.
4. Click **Yes**.
   - The Internal ECDSA Certificate Authority creation process begins and a new tab opens to show the progress of the process.
   - When the process is finished, the progress report shows the steps you must take next.
   - The status of the ECDSA CA is **Created for Different Certificate Type**.
5. Restart the Log Server, Authentication Server, and Web Portal Server. See Management Center Commands (page 1146) for more information.
6. Start the **Renew Internal Certificate Authorities Task**. See Starting Tasks Manually (page 1040).
   - When the Task is finished, the status of the ECDSA CA is **Ready to Use for Different Certificate Type**.
7. Right-click the Renew Internal Certificate Authorities Task in the History branch and select **Show Progress Report**. The progress report shows which steps you must take next.
   - Follow the instructions to resolve any issues. For example, you may be prompted to check the status or connectivity of some engines.
8. Recertify the Management Server as explained in Renewing SMC Server Certificates (page 1094).
9. Start the **Renew Internal Certificate Authorities Task** again.
   - When the Task is finished, the status of the ECDSA CA is **Active**.
10. Recertify the Log Server, Authentication Server, and Web Portal Server as explained in Renewing SMC Server Certificates (page 1094).
    - If the engine is not able to communicate with the Management Server, check that 256-bit encryption is enabled on the engine and make initial contact between the engine and the Management Server. See Enabling 256-bit Security Strength for Engines (page 234).
    - For other certificate-related problems, see Troubleshooting Certificates (page 1091).

If Configuration Changes Prevent Managing the Engines

Prerequisites: None

If an engine cannot connect to the Management Server because of changes in the system configuration, you can restore the contact by generating a new initial configuration for the engine (through the engine's right-click menu) and then running the `sg-reconfigure` command on the engine command line. If you also suspect that the engine's configuration is out of date, select the option to return the engine to the initial configuration state. See instructions in Reconfiguring Basic Engine Settings (page 233).

Caution – Returning the engine to the initial configuration state clears the engine's configuration. Only management communications are allowed and no traffic is allowed to pass through the engine.

Changing the Role of Security Engines

Prerequisites: See Limitations

You can change the role of a Security Engine to convert one type of Security Engine to another type if you have a specific need to do so.

Limitations

- You can only change the Security Engine role for engines that currently have Security Engine software installed. To change the role of engines that currently have specific engine software (for example, Firewall/VPN) installed, you must reinstall the engine software as instructed in the engine-specific *Installation Guide*.
- Changing the engine role is only supported on modular appliances, for engines installed on a virtualization platform, or for engines installed on your own hardware. You cannot change the engine role on small appliances.
- You must have a Security Engine license that is valid for all engine roles. You cannot change the role of engines that have a license for a specific type of engine.
- You must connect to the engine's command line through a serial console or VGA console. It is not possible to change the engine's role using an SSH connection to the engine.

Preparing to Change the Security Engine Role

Note – You cannot use the same primary control IP address in multiple elements. You must either change the primary control IP address in the engine's current interface configuration or delete the existing engine element before creating a new engine element for the new Security Engine role.

▼ To prepare to change the Security Engine role

1. Create the correct type of engine element for the new role and define the basic properties as explained in Getting Started with Engine Elements (page 358).
2. Configure the engine's interfaces as explained in Getting Started with Interface Configuration (page 420).
3. Configure the routing as explained in Getting Started with Routing (page 588).
   - Routes to directly connected networks are automatically added for all Security Engine roles.
   - For Firewalls, you must add a default route and any routes through next-hop gateways to networks that are not directly connected to the Firewalls.
   - You may need to define a default route for IPS engines and Layer 2 Firewalls if the Management Center and other system components are not located on a directly connected network.
4. Generate the initial configuration for the engine and save the configuration on a USB stick as explained in Saving an Initial Configuration for Security Engines (page 515).

Clearing the Existing Security Engine Configuration

▼ To clear the existing Security Engine configuration

1. Connect to the engine's command line through a serial console or VGA console and log in as root.
2. Run the **`sg-clear-all`** command. The engine reboots and you are prompted to select the system restore options.

Caution – This command clears all configuration information from the engine. It is not possible to recover the engine's previous configuration when this command is used.

3. Enter **`2`** and press **`Enter`**. You are prompted to select the number of overwrites for the data partition.
4. Enter **`0`** and press **`Enter`**. You are prompted to select the number of overwrites for the swap partition.
5. Enter **`0`** and press **`Enter`**. You are prompted to select whether to remove the spool partition.
6. Enter **`2`** and press **`Enter`**. You are prompted to select the number of overwrites for the spool partition.
7. Enter **`0`** and press **`Enter`**. You are prompted to confirm that you want to clear the configuration.

8. Type **YES** and press **`Enter`**. The system starts clearing the engine's configuration.

> **Caution – Do not power off or reboot the engine while the configuration is being cleared. This can cause a serious error that prevents the engine from booting.**

9. Press **Enter** to reboot the engine when prompted. The engine reboots and the Configuration Wizard starts automatically.

Reconfiguring the Security Engine

▼ To reconfigure the engine

1. Highlight **Role** and press **`Enter`**.

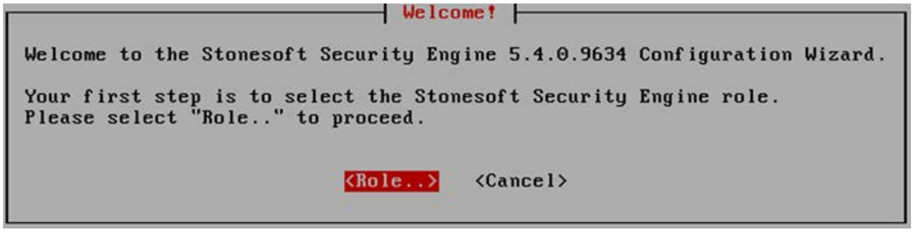

```
Welcome to the Stonesoft Security Engine 5.4.0.9634 Configuration Wizard.

Your first step is to select the Stonesoft Security Engine role.
Please select "Role.." to proceed.
```

2. Select the new role for the Security Engine and press **`Enter`**. The role-specific Configuration Wizard starts.

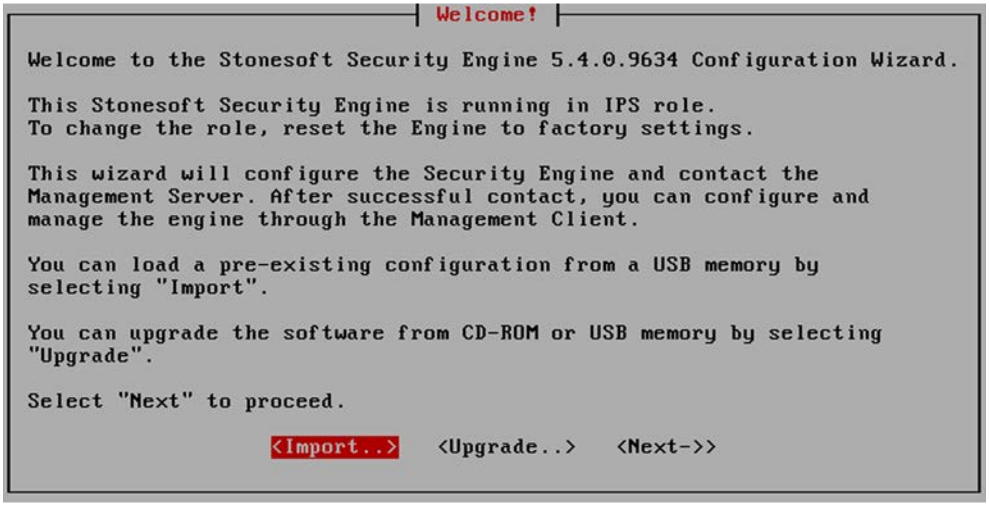

```
Welcome to the Stonesoft Security Engine 5.4.0.9634 Configuration Wizard.

This Stonesoft Security Engine is running in IPS role.
To change the role, reset the Engine to factory settings.

This wizard will configure the Security Engine and contact the
Management Server. After successful contact, you can configure and
manage the engine through the Management Client.

You can load a pre-existing configuration from a USB memory by
selecting "Import".

You can upgrade the software from CD-ROM or USB memory by selecting
"Upgrade".

Select "Next" to proceed.
```

3. Select one of the following configuration methods:
   - Highlight **Import** and press **`Enter`** to import a saved configuration.
   - Highlight **Next** and press **`Enter`** to manually configure the engine's settings.
4. (*Manual configuration only*) Configure the Operating System Settings and Network Interfaces. See the *Firewall/VPN Installation Guide* and the *IPS and Layer 2 Firewall Installation Guide* for detailed instructions.

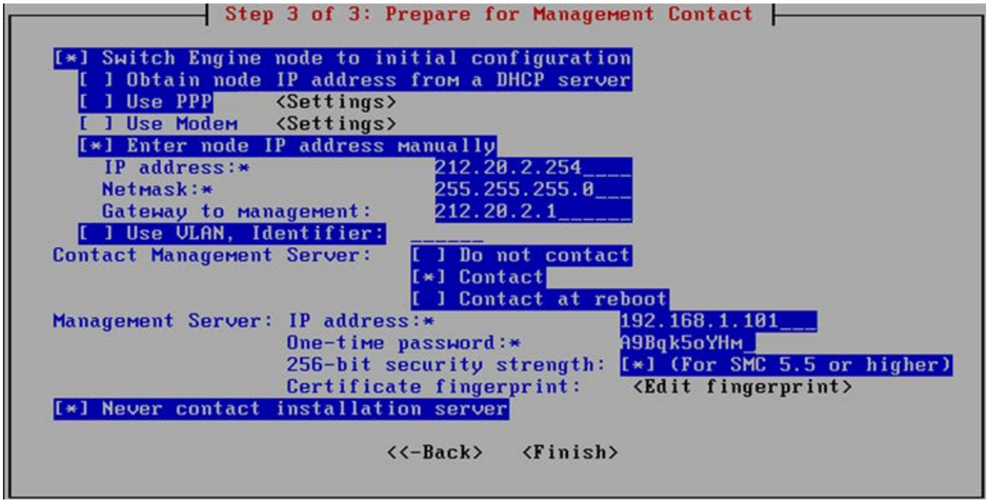

```
              Step 3 of 3: Prepare for Management Contact

[*] Switch Engine node to initial configuration
  [ ] Obtain node IP address from a DHCP server
  [ ] Use PPP       <Settings>
  [ ] Use Modem     <Settings>
  [*] Enter node IP address manually
    IP address:*               212.20.2.254____
    Netmask:*                  255.255.255.0___
    Gateway to management:     212.20.2.1______
  [ ] Use VLAN, Identifier:  ______
Contact Management Server:   [ ] Do not contact
                             [*] Contact
                             [ ] Contact at reboot
Management Server: IP address:*               192.168.1.101___
                   One-time password:*        A9Bqk5oYHm_
                   256-bit security strength: [*] (For SMC 5.5 or higher)
                   Certificate fingerprint:    <Edit fingerprint>
[*] Never contact installation server

                           <<-Back>   <Finish>
```

5. Select **Switch Engine node to initial configuration** and fill in the Management Server information. See the *Firewall/VPN Installation Guide* and the *IPS and Layer 2 Firewall Installation Guide* for detailed instructions.

6. Select **Contact** and enter the **Management Server IP address** and the **one-time password**.

7. (*Optional*) Select **256-bit Security Strength** and press the spacebar to use 256-bit encryption for the connection to the Management Server. 256-bit encryption must also be enabled for the Management Server.

8. (*Optional*) Highlight **Edit Fingerprint** and press `Enter`. Fill in the Management Server's certificate fingerprint (also shown when you saved the initial configuration). Filling in the certificate fingerprint increases the security of the communications.

9. Highlight **Finish** and press `Enter`. The Security Engine makes initial contact with the Management Server.

10. Install a policy on the engine as explained in Installing Policies (page 652).

Engine Element Configuration

In this section:

CHAPTER 26

CREATING AND MODIFYING ENGINE ELEMENTS

Engine elements contain the configuration information that is directly related to the Firewalls, IPS engines, and Layer 2 Firewalls. This includes interface definitions, cluster mode selection, tester settings, and other such engine-specific options. This section explains how to create and modify these elements and lists the tasks you can do in the engine element properties. For information about creating and modifying Master Engine and Virtual Firewall elements, see Getting Started with Virtual Engines (page 404).

The following sections are included:

Getting Started with Engine Elements

Prerequisites: None

What Engine Elements Do

Engine elements are used for the configuration and management of your Firewalls, IPS engines, and Layer 2 Firewalls. They contain settings that cannot be reused in the configuration of other components, such as the network interface configuration. The engine elements also determine which of the reusable elements the configuration of a particular component includes (for example, the Log Server to which the component sends its log data). All engines are centrally configured and controlled through the Management Server.

What Do I Need to Know Before I Begin?

If you are configuring new engine elements for the first time, we recommend that you follow the instructions in the *Firewall/VPN Installation Guide* or the *IPS and Layer 2 Firewall Installation Guide* instead of the instructions here. For background information on how the system works, see the *Firewall/VPN Reference Guide* or the *IPS and Layer 2 Firewall Reference Guide*.

Configuration Overview

The overview below does not cover the Stonesoft SSL VPN product. See Creating a New SSL VPN Gateway Element (page 382). For information on configuring Stonesoft SSL VPN, see the Stonesoft *SSL VPN Administrator's Guide*.

1. Install a license for the engine. See Getting Started with Licenses (page 1044).
2. Create a new engine element and define the basic properties. See Creating New Engine Elements (page 359).
3. Configure the engine's interfaces. See Getting Started with Interface Configuration (page 420).
4. Configure the routing. See Getting Started with Routing (page 588).
5. Generate the initial configuration for the engine and use it to establish a connection between the engine and the Management Server. See Connecting Engines to the Management Center (page 513).
6. Install a policy on the engine. See Installing Policies (page 652).

Related Tasks

- Editing Single Firewall Properties (page 395)
- Editing Firewall Cluster Properties (page 396)
- Editing Single IPS Engine Properties (page 397)
- Editing IPS Cluster Properties (page 398)
- Editing Single Layer 2 Firewall Properties (page 399)
- Editing Layer 2 Firewall Cluster Properties (page 400)

Creating New Engine Elements

Prerequisites: None

You can either create a new element as instructed in the sections listed below or you can copy and modify an existing element as explained in Duplicating an Existing Engine Element (page 383). Before you define a new engine element, make sure you have a license for it. The element can be configured without a license, but you must have a license to make the engine operational.

There are two ways to create Single Firewall and Firewall Cluster elements:

- You can create the Single Firewall and Firewall Cluster elements one by one. See Creating a New Single Firewall Element and Creating a New Firewall Cluster Element (page 373).
- You can create several new Single Firewall and Firewall Cluster elements at the same time with a wizard. See Creating Multiple Single Firewall Elements (page 361) and Creating Multiple Firewall Cluster Elements (page 374).

What's Next?

- ▶ Creating a New Single Firewall Element
- ▶ Creating Multiple Single Firewall Elements (page 361)
- ▶ Creating a New Firewall Cluster Element (page 373)
- ▶ Creating Multiple Firewall Cluster Elements (page 374)
- ▶ Creating a New Single IPS Element (page 378)
- ▶ Creating a New IPS Cluster Element (page 379)
- ▶ Creating a New Single Layer 2 Firewall Element (page 380)
- ▶ Creating a New Layer 2 Firewall Cluster Element (page 381)
- ▶ Creating a New SSL VPN Gateway Element (page 382)

Creating a New Single Firewall Element

Single Firewall elements represent firewalls that consist of one physical device. They can also be converted to Firewall Cluster elements. See Converting a Single Firewall to a Firewall Cluster (page 385).

If you want to create several Single Firewall elements at the same time, see Creating Multiple Single Firewall Elements (page 361).

▼ To create a new Single Firewall element

1. Select **Configuration→Configuration→Security Engine**. The Security Engine Configuration view opens.
2. Right-click **Security Engines** and select **New→Single Firewall**. The Single Firewall Properties dialog opens.

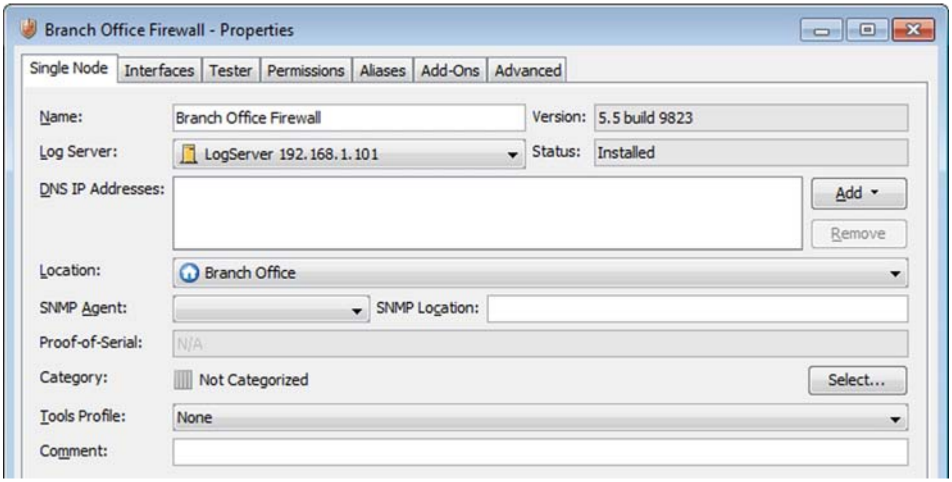

Branch Office Firewall - Properties

Name: Branch Office Firewall

Version: 5.5 build 9823

Log Server: LogServer 192.168.1.101

Status: Installed

DNS IP Addresses:

Location: Branch Office

SNMP Agent:

SNMP Location:

Proof-of-Serial: N/A

Category: Not Categorized

Tools Profile: None

Comment:

3. Give the element a unique **Name**.

4. Select the **Log Server** to which the Firewall sends its event data.

5. (*Optional*) Define one or more **DNS IP Addresses**. These are the IP addresses of the DNS server(s) that the Firewall uses to resolve virus signature mirrors, domain names, and web filtering categorization services (which are defined as URLs). There are two ways to define IP addresses.
   - To enter a single IP address manually, click **Add** and select **IP Address**. Enter the IP address in the dialog that opens.
   - To define an IP address using a network element, click **Add** and select **Network Element**. Select an existing element, or click the New icon and define a new element. See Defining Host Elements (page 727) and Defining an External DNS Server (page 640) for more information.

6. Select the **Location** for this engine if there is a NAT device between this Firewall and other system components. See Defining Locations (page 66) for more information.

7. (*Optional*) If you want to enable the Firewall to send SNMP traps, select an **SNMP Agent** as explained in Activating the SNMP Agent on Engines (page 584).

8. (*Optional, appliances only*) Enter the **Proof-of-Serial** (POS) code delivered with the appliance if you want to configure the engine using plug-and-play configuration.

9. (*Optional*) If you want to include the Firewalls in pre-defined categories, select the appropriate **Categories** as explained in Selecting Categories for Elements (page 77).

10. (*Optional*) If you want to add custom commands to the Firewall's right-click menu, add a **Tools Profile** as explained in Attaching a Tools Profile to an Element (page 56).

What's Next?

▶ Continue the configuration in Firewall Interface Configuration (page 421).

Creating Multiple Single Firewall Elements

You can create multiple Single Firewall elements at the same time using the Create Multiple Single Firewalls Wizard. Using the wizard is a particularly good way to create firewalls for plug-and-play configuration, where the appliances configure themselves automatically when they are plugged in and connected to the network. Only specific Stonesoft appliances can use plug-and-play configuration.

Using the Create Multiple Single Firewalls Wizard offers several benefits:

- It is easy to create several Firewall elements at the same time.
- You can automatically create Internal Security Gateway elements that represent the Firewalls in VPNs.
- You can define a policy that is automatically installed on the Firewalls when they make initial contact to the Management Server.
- (*Plug-and-play configuration only*) You can upload the Firewall configurations to the Stonesoft Installation Server, instead of having to upload the configurations one by one or transfer the initial configuration to the appliances manually.

You can either create a specific number (1-1000) of Single Firewall elements or you can use the Proof-of-Serial (POS) codes delivered to you with the Stonesoft appliances to create the Single Firewall elements. When you use POS codes in the wizard, all the appliances must be the same model. If you have POS codes for different types of appliances, you must run the wizard separately for each appliance model to create the elements. Before you start creating the Firewall elements, make a note of the serial numbers and the geographical locations where the appliances will be used.

Note – The Firewall properties you define in the Create Multiple Single Firewalls Wizard are common to all Firewalls created with the wizard. Consider carefully which properties are shared by all Firewalls and which properties you need to define separately for each Firewall.

Once you have created the Firewalls through the wizard, you can modify the properties of each individual Firewall. See Editing Single Firewall Properties (page 395). You can also modify some shared properties of several Firewalls at the same time. See Modifying Properties of Several Engines at Once (page 385).

▼ To create multiple Single Firewall elements

1. Select **Configuration→Configuration→Security Engine**. The Security Engine Configuration view opens.
2. Right-click **Security Engines** and select **New→Multiple Single Firewalls**. The Create Multiple Single Firewalls Wizard opens.

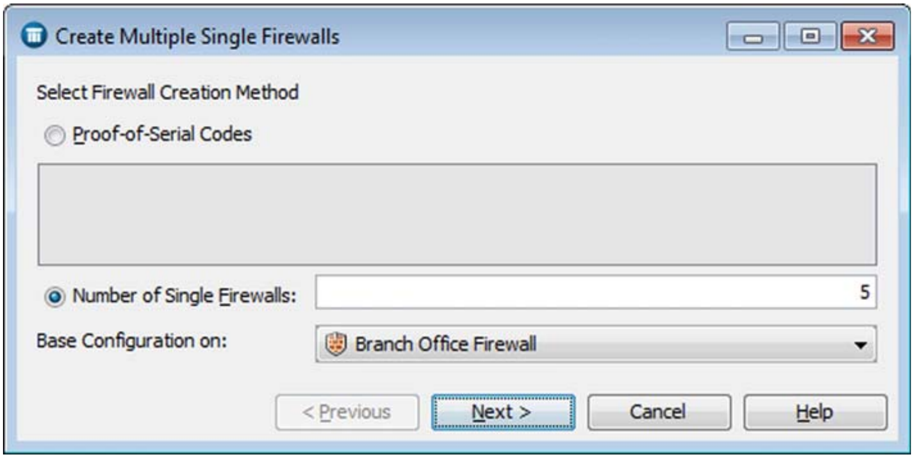

Create Multiple Single Firewalls

Select Firewall Creation Method

○ Proof-of-Serial Codes

● Number of Single Firewalls: 5

Base Configuration on: Branch Office Firewall

3. Select the Firewall creation method.
   - If you do not have POS codes, enter the **Number of Single Firewalls** (1-1000) and proceed to Step 5.
   - If you have received POS codes from Stonesoft, enter the POS codes in the **Proof-of-Serial Codes** field and click **Next**. The Proof-of-Serial Code Information page opens.
4. (*POS codes only*) Check that the details on the Proof-of-Serial Code Information page are correct.
5. (*Optional*) Select a Firewall on which you want to base the Firewalls' configuration from the **Base Configuration on** list and click **Next**. The Define Basic Firewall Information page opens.
6. Enter a common **Name Prefix**. The system adds either a running number or the serial number of the appliance to the name prefix to generate a unique name for each individual Firewall.
7. Select the **Log Server** to which the Firewall sends its event data.

Note – Name Prefix and Log Server are the only mandatory properties you must define at this stage. Review the other properties carefully to see which ones to define as the shared properties for all Firewalls created with the wizard.

8. (*Optional*) Define one or more **DNS IP Addresses**. These are the IP addresses of the DNS server(s) that the Firewalls use to resolve virus signature mirrors, domain names, and web filtering categorization services (which are defined as URLs). There are two ways to define IP addresses.
   - To enter a single IP address manually, click **Add** and select **IP Address**. Enter the IP address in the dialog that opens.
   - To define an IP address using a network element, click **Add** and select **Network Element**. Select an existing element, or click the New icon and define a new element. See Defining Host Elements (page 727) and Defining an External DNS Server (page 640) for more information.

9. Select the **Location** for the Firewalls if there is a NAT device between these Firewalls and other system components. For more information, see Defining Locations (page 66).

Note – Select a Location only if all the Firewalls you create with the wizard belong to the same Location element.

10. Define other settings according to your environment:
    - If you have a Stonesoft appliance, enter the **Proof-of-Serial** (POS) code delivered with the appliance.
    - If you want to enable the Firewalls to send SNMP traps, select an **SNMP Agent** as explained in Activating the SNMP Agent on Engines (page 584).
    - If you want to include the Firewalls in pre-defined categories, select the appropriate **Categories** as explained in Selecting Categories for Elements (page 77).
    - If you want to add custom commands to the Firewalls' right-click menu, add a **Tools Profile** as explained in Attaching a Tools Profile to an Element (page 56).
11. Click **Next**.The Review and Edit Names and Comments page opens.
12. Review the names of the Firewalls. If necessary, right-click the name and select **Edit Name**.

Note – It is recommended to give each Firewall a unique, descriptive name after the common Name Prefix, such as the geographical location where each particular Firewall engine will be used.

13. Click **Next**. The Define Interfaces for the Multiple Single Firewall Elements page opens.

Defining Interfaces for Multiple Single Firewalls

You can define various properties for the interface(s) at this stage or return to them later. The same interface properties are available in the Properties dialog of every Firewall element. However, it is recommended that you configure all the shared interface properties in the wizard. After you exit the wizard, you must configure the properties separately for each Firewall.

The interface properties you define for the first Firewall are used to automatically create the corresponding properties for the rest of the Firewalls. This also includes the IP addresses, which are automatically generated in numerical order. Make sure that the IP addresses that are assigned to the Firewalls are not used by any other components.

If you want to use a physical interface as the control interface, begin the interface definition by defining a Physical Interface with a dynamic IP address, so that the Physical Interface is assigned Interface ID 0. When connecting the cables to the appliance, connect the cable for the control connection to Ethernet port 0. See the relevant *Appliance Installation Guide(s)* for detailed information on mapping the Interface IDs with specific ports on the appliance(s).

▼ To define interfaces for the Firewalls

1. Click **Add** and select the type of interface that you want to add. Add the required number of network interfaces in the following order:
   - Define Physical Interfaces as explained in Defining Physical Interfaces for Firewall Engines (page 422).
     - For settings on the **DHCP** tab, see Activating the Internal DHCP Server on a Firewall Interface (page 511) or, for DHCP relay, Routing DHCP Messages (page 590).
     - For settings on the **Advanced** tab, see Configuring Advanced Interface Properties for Firewalls (page 436).
   - Define integrated ADSL modems as explained in Adding ADSL Interfaces for Single Firewalls (page 427).
   - Define integrated 3G modems as explained in Defining Modem Interfaces for Single Firewalls (page 448).
   - Define Tunnel Interfaces as explained in Defining Tunnel Interfaces for Firewalls (page 435).
   - Define integrated wireless modems as explained in Adding Wireless Interfaces for Single Firewalls (page 429).
   - Define SSID Interfaces as explained in Defining SSID Interfaces for Single Firewalls (page 431).
2. *(Optional, Physical Interfaces only)* Add the required number of VLANs as explained in Adding VLAN Interfaces for Firewall Engines (page 425).
3. (*Not applicable to Modem interfaces*) Configure the IP address settings for the first one of the Firewalls as explained in Configuring Single Firewall IP Addresses (page 439). At least one of the IP addresses must be a dynamic IPv4 address.
4. (*Optional*) Click **Options** to configure Loopback IP addresses as explained in Configuring Loopback IP Addresses for Firewalls (page 450).
5. (*Recommended*) Click **Options** to define which IP addresses are used in particular roles in system communications as explained in Setting Interface Options for Firewalls (page 453).
6. *(Optional)* Configure additional routing settings:
   - For adding **ARP entries**, see Configuring Manual ARP Settings (page 510).
   - For configuring **Multicast Routing**, see Routing Multicast Traffic (page 593).
7. Click **Next**. The Review and Edit Firewall Interfaces page opens.
8. Review the interfaces and edit them, if necessary. You cannot edit any properties of the Modem Interfaces.
9. Click **Next**. The Define Routing for the Multiple Single Firewall Elements page opens.

What's Next?

▶ Defining Routing for Multiple Single Firewalls (page 365)

Defining Routing for Multiple Single Firewalls

Routes to directly connected networks are automatically added. You must add a default route and any routes through next-hop gateways to networks that are not directly connected to the firewalls.

▼ To define the routing for the Firewalls

1. Drag and drop the appropriate Router and Network element(s) under the correct interfaces in the right panel. For detailed information about configuring routing, see Adding Routes for Firewalls (page 589)

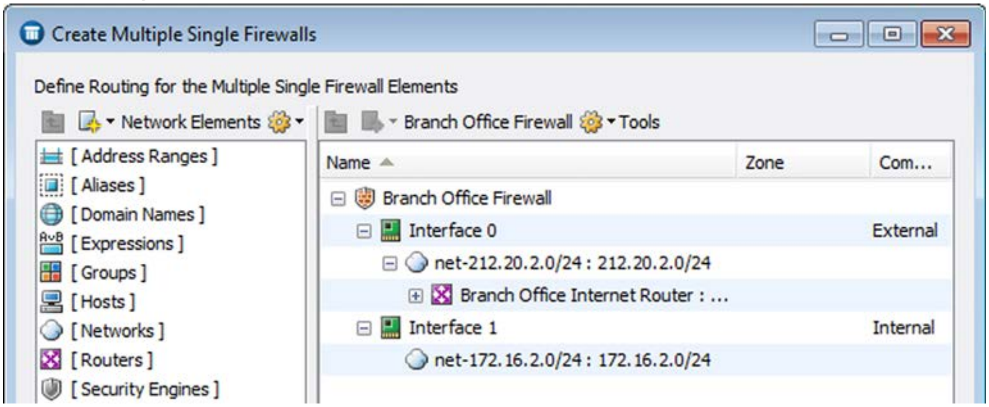

Name	Zone	Com...
Branch Office Firewall		
Interface 0		External
net-212.20.2.0/24 : 212.20.2.0/24		
Branch Office Internet Router : ...		
Interface 1		Internal
net-172.16.2.0/24 : 172.16.2.0/24		

2. Click **Next**. The Review and Edit Routing for the Multiple Single Firewall Elements page opens.
3. Review the routing and edit it, if necessary.
4. Click **Next**. The Select Additional Configuration Options page opens.

What's Next?

▶ Selecting Additional Configuration Options for Multiple Single Firewalls (page 366).

Selecting Additional Configuration Options for Multiple Single Firewalls

After defining the interfaces and routing, you have configured all the mandatory properties for the Firewalls. On the Select Additional Configuration Options page, you can define additional properties for the Firewalls and/or create Internal Security Gateway elements that represent the Single Firewall elements in VPNs. We recommend that you define all the shared properties of the Firewalls in the wizard.

What's Next?

- If you do not want to configure additional options, click **Next** and proceed to Defining Tester Settings for Multiple Single Firewalls (page 366).
- Otherwise, continue as explained below.

▼ To select additional configuration options

1. Select which of the following options, if any, you want to configure:
   - Select **Define Additional Firewall Properties** if you want to define Tester Settings, Permissions, Add-Ons (settings for TLS Inspection, User Agent, Anti-Virus, or Anti-Spam), and/or Advanced settings for the Firewalls.
   - Select **Create an Internal Security Gateway Element for Each Firewall** if you want to create Internal Security Gateway elements for the Firewalls.

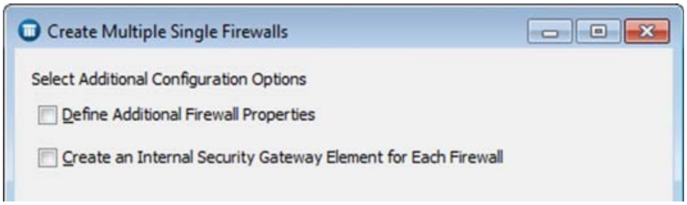

Create Multiple Single Firewalls

Select Additional Configuration Options

□ Define Additional Firewall Properties

□ Create an Internal Security Gateway Element for Each Firewall

2. Click **Next**.

What's Next?

- Defining Tester Settings for Multiple Single Firewalls (page 366)

Defining Tester Settings for Multiple Single Firewalls

On the Define Tester Settings for the Firewalls page, you can set periodic self-tests for the Firewalls to ensure that they function properly. For more information on the engine tests and their settings, see Getting Started with the Engine Tester (page 520).

What's Next?

- If you do not want to configure tester settings, click **Next** and proceed to Defining Permissions for Multiple Single Firewalls (page 367).
- Otherwise, continue as explained below.

▼ To define Tester Settings for the Firewalls

1. If you want to define Global Settings for the Tester, see Specifying Global Engine Tester Settings (page 521).
2. If you want to add Test Entries, see Adding Engine Tests (page 522).
3. Click **Next**. The Define Permissions for the Firewalls page opens.

Defining Permissions for Multiple Single Firewalls

On the Define Permissions for the Firewalls page, you can define which access control list(s) the Firewalls belong to, what kind of permissions individual administrators have, and which policies are allowed for the Firewalls.

For more information on Firewall permissions, see Getting Started with Engine Permissions (page 530). For more information on access control in the Management Center, see Getting Started with Administrator Accounts (page 248).

What's Next?

▶ If you do not want to configure permissions, click **Next** and proceed to Defining Add-Ons for Multiple Single Firewalls (page 367).

▶ Otherwise, continue as explained below.

▼ To define administrator permissions and select permitted policies for the Firewalls

1. If you want to define administrator permissions on the Firewalls, see Defining Administrator Permissions on Engines (page 530).
2. If you want to select the permitted policies for the Firewalls, see Selecting Permitted Policies for Engines (page 531).
3. Click **Next**. The Define Add-Ons for the Firewalls page opens.

Defining Add-Ons for Multiple Single Firewalls

On the Define Add-Ons for the Firewalls page, you can define settings for TLS inspection, User Agent, browser-based user authentication, anti-virus, and/or anti-spam for the Firewalls. Anti-virus and anti-spam are available as separately licensed features on selected platforms.

- For more information about anti-virus, see Configuring Anti-Virus Settings (page 540).
- For more information about anti-spam, see Configuring Anti-Spam Settings (page 541).
- For more information about TLS Inspection, see Getting Started with TLS inspection (page 794).
- For more information about the User Agent, see Enabling Access Control by User (page 838) for more information.

What's Next?

▶ If you do not want to configure Add-Ons, click **Next** and proceed to Defining Advanced Settings for Multiple Single Firewalls (page 368).

▶ Otherwise, continue as explained below.

▼ To define Add-Ons for the Firewalls

1. *(Optional)* Select the appropriate **Client Protection Certificate Authority** element for TLS Inspection.
   - For information on how to configure client protection, see Configuring Client Protection (page 797).
2. *(Optional)* Click **Add** and select the **Server Protection Credentials** element(s). The selected elements are added to the list.
   - For information on how to configure server protection, see Configuring Server Protection (page 796).
3. *(Optional)* Select the **User Agent** element that the Firewall communicates with for Access Control by User.
   - For information on how to create User Agent elements, see Creating User Agent Elements (page 840).
4. *(Optional)* Click **User Authentication** to enable browser-based user authentication.
   - For information on how to enable browser-based user authentication, see Enabling Browser-Based User Authentication (page 874)
5. *(Optional)* Select **Enable** to activate the anti-virus options.
   - For information on how to configure anti-virus settings, see Configuring Anti-Virus Settings (page 540).
6. *(Optional)* Select the appropriate **Anti-Spam Settings**.
   - For information on how to configure anti-spam settings, see Configuring Anti-Spam Settings (page 541).
7. Click **Next**. The Define Advanced Settings for the Firewalls page opens.

Defining Advanced Settings for Multiple Single Firewalls

On the Define Advanced Settings for the Firewalls page, you can define various system parameters and traffic handling parameters for the Firewalls.

Caution – Improper adjustments to some of the advanced settings may seriously degrade the performance of the system.

What's Next?

- ▶ If you do not want to configure advanced settings, and you selected **Create an Internal Security Gateway Element for Each Firewall**, click **Next** and proceed to Defining Internal Security Gateway End-Points for Multiple Single Firewalls (page 369).
- ▶ If you do not want to configure advanced settings, click **Next** and proceed to Uploading the Multiple Single Firewall Initial Configuration to the Installation Server (page 371).
- ▶ Otherwise, continue as explained below.

▼ To define Advanced Settings for the Firewalls

1. For adjusting the system parameters of the Firewalls, see Adjusting Firewall System Parameters (page 553).
2. For adjusting the traffic handling parameters of the Firewalls, see Adjusting Firewall Traffic Handling Parameters (page 554).
3. Click **Next**.

What's Next?

- ▶ If you selected **Create an Internal Security Gateway Element for Each Firewall**, proceed to Defining Internal Security Gateway End-Points for Multiple Single Firewalls (page 369).
- ▶ Otherwise, proceed to Uploading the Multiple Single Firewall Initial Configuration to the Installation Server (page 371).

Defining Internal Security Gateway End-Points for Multiple Single Firewalls

On the Define End-Points for the Internal Security Gateways page, you can define Internal Security Gateway end-points (IP addresses) for the Firewalls. These end-points are used when the Firewall acts as a gateway in a policy-based VPN. The same end-point cannot be used in both a policy-based VPN and the Route-Based VPN. Each IP address can only be used once. For more information on Security Gateway elements, see Defining Security Gateways (page 923).

What's Next?

- ▶ If you do not want to define end-points for Internal Security Gateways, proceed to .

▼ To define end-points for the Internal Security Gateways of the Firewalls

1. Change the selection of the IPv4/IPv6 address(es) that you want to use as end-points. Typically, these are IP address(es) that belong to interface(s) towards the Internet, which are selected by default.
   - Both IPv4 and IPv6 addresses can be used for end-points.
   - If you have more than one Internet connection, select an IP address from each ISP to make Multi-Link load balancing and failover possible.

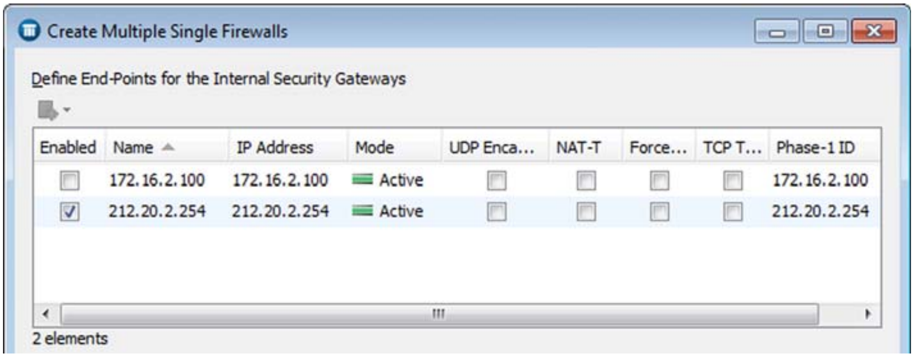

Create Multiple Single Firewalls

Define End-Points for the Internal Security Gateways

Enabled	Name ▲	IP Address	Mode	UDP Enca...	NAT-T	Force...	TCP T...	Phase-1 ID
□	172.16.2.100	172.16.2.100	Active	□	□	□	□	172.16.2.100
☑	212.20.2.254	212.20.2.254	Active	□	□	□	□	212.20.2.254

2 elements

2. *(Optional)* To change the name of an end-point, right-click the name and select **Properties**. Enter a new **Name** and click **OK**.
3. (*Optional*) To select the **Mode** to define how the end-point is used in a Multi-Link configuration, right-click the name and select **Properties**. You can override the mode settings in each individual VPN.
   - **Active**: use the tunnel(s) of this end-point whenever possible.
   - **Standby**: use the tunnel(s) of this end-point only if the Active end-points cannot be used.
   - **Aggregate**: use the tunnel(s) of this end-point together with the tunnel(s) of another end-point.

4. (*Optional, legacy VPN only*) Select **UDP Encapsulation** if you want to encapsulate IPsec communications in UDP packets using Stonesoft's proprietary method.

5. (*Optional*) Select one of the NAT-T options to activate encapsulation for NAT traversal in gateway-to-gateway VPNs, which may be needed to traverse a NAT device at the local or at the remote gateway end.

Option	Description
NAT-T	Select this option if you want to allow encapsulating the IPsec communications in standard NAT-T UDP packets in gateway-to-gateway VPNs when the gateways detect that a NAT operation is applied to the communications. If both gateways do not support this option, the option is ignored.
Force NAT-T	Select this option to force NAT-T even when the gateways do not detect a NAT operation being applied to the communications. If both gateways do not support this option, the VPN fails to establish.

- The gateway always allows VPN clients to use NAT-T regardless of these settings.
- NAT-T always uses the standard UDP port 4500.

6. (*Optional*) Select **TCP Tunneling** if you want to tunnel Stonesoft IPsec VPN Client communications with this Gateway end-point in a TCP connection to bypass a traffic filtering device that does not allow standard IPsec ports to pass or to traverse a NAT device.
   - This option may not be supported by all external gateways. Support is required at both ends of each tunnel.

7. *(Optional)* Review the phase-1 properties of the end-points and edit them, if necessary.
   - To edit phase-1 properties, right-click the **Phase-1 ID** field and select **Properties**.
   - Select the **ID Type** for the ID that identifies the Gateways during the IKE phase-1 negotiations. The **Distinguished Name** type is only valid in certificate-based authentication.
   - Enter an **ID Value**.
     - If the end-point has a dynamic IP address, you must enter a specific IP address as the value for the **IP Address** type.
     - If the end-point has a static IP address, the value for the **IP Address** type is filled in automatically according to the IP address you defined for this end-point.

8. Click **Next**. The Review and Edit Internal Security Gateway End-Points page opens.

9. Review the end-point details and edit them, if necessary.

10. Click **Next**.

What's Next?

▶ If you are creating the Firewall elements using POS codes, proceed to Uploading the Multiple Single Firewall Initial Configuration to the Installation Server (page 371).

▶ Otherwise, proceed to Selecting a Policy to Install on the Firewalls (page 372).

Uploading the Multiple Single Firewall Initial Configuration to the Installation Server

Uploading the initial configuration to the Installation Server makes the configuration available for use in plug-and-play installations.

▼ To upload the Initial Configuration to the Installation Server

1. Make sure **Upload Initial Configuration** is selected.

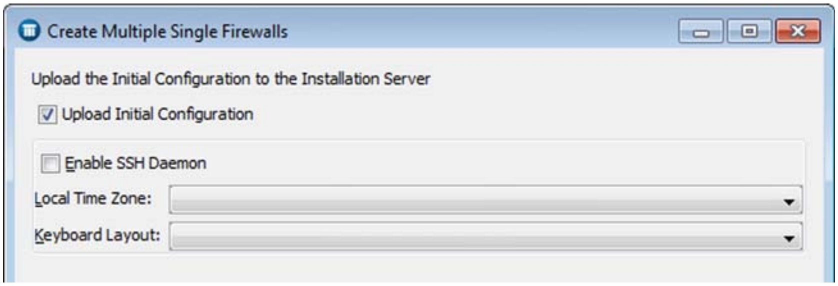

Create Multiple Single Firewalls

Upload the Initial Configuration to the Installation Server

☑ Upload Initial Configuration

☐ Enable SSH Daemon

Local Time Zone:

Keyboard Layout:

2. *(Optional)* Select **Enable SSH Daemon** to allow remote access to the engine command line. SSH access may be helpful for remote troubleshooting.
   - After management contact is established, you can enable and disable remote command line access to the engine at any time through the right-click menu of the engine. We recommend that you disable SSH access whenever it is not needed and that you make sure your Access rules allow SSH access to the engines from the administrators' IP addresses only.
   - The Firewall Template does not allow these connections, but the temporary policy activated right after the engine's initial configuration (in force until you install the working policy) allows SSH access from the Management Server's IP address. Alternatively, you can upload a working policy to be installed on the Firewalls after they have contacted the Management Server. For more information, see Selecting a Policy to Install on the Firewalls (page 372).

Caution – If you enable SSH, set the password for command line access after the initial configuration either through the Management Client or by logging in to the command line. When the password is not set, anyone with SSH access to the engine can set the password.

3. Select the **Local Time Zone** and the **Keyboard Layout** for use on the command line. The time zone setting is only for displaying the time on the local console; the engines always use UTC (GMT) time internally. The clock is automatically synchronized to match the Management Server's time.
4. Click **Next**. The Review and Edit Local Time Zones page opens.
5. Review the local time zones of the Firewalls and change them, if necessary.
6. Click **Next**. The Select a Policy to Install on the Firewalls page opens.

What's Next?

▶ Selecting a Policy to Install on the Firewalls (page 372).

Selecting a Policy to Install on the Firewalls

On the Select a Policy to Install on the Firewalls page, you can select a predefined policy to install on the Firewalls after the Firewalls have contacted the Management Server for the first time.

What's Next?

- If you do not want to select a policy to be installed on the Firewalls, click **Next** and proceed to the Summary page. Click **Finish** to finish the creation of multiple Single Firewall elements.
- Otherwise, continue as explained below.

To select a policy to install on the Firewalls

1. Click **Select** and select the appropriate **Policy** from the list.

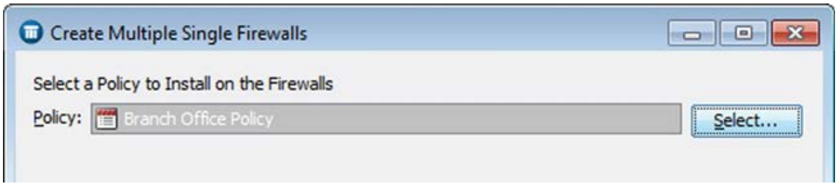
Create Multiple Single Firewalls
Select a Policy to Install on the Firewalls
Policy:
Branch Office Policy
Select...

2. Click **Next**. The Summary page opens.

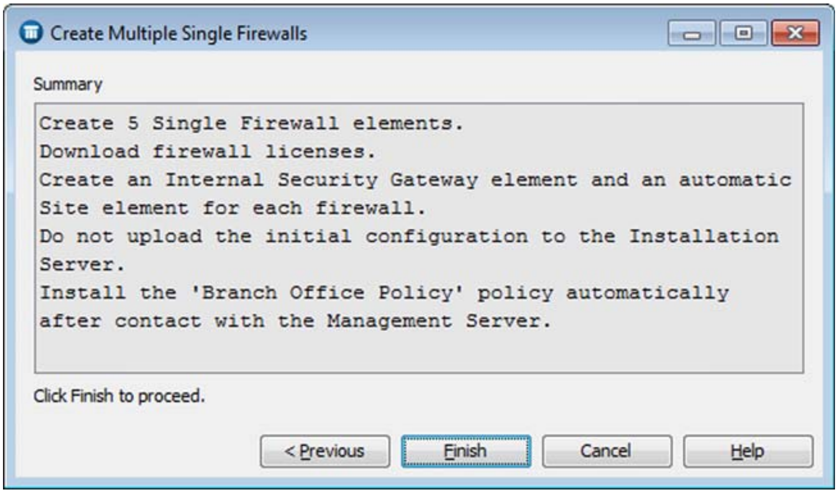

```
Create 5 Single Firewall elements.
Download firewall licenses.
Create an Internal Security Gateway element and an automatic
Site element for each firewall.
Do not upload the initial configuration to the Installation
Server.
Install the 'Branch Office Policy' policy automatically
after contact with the Management Server.
```

Click Finish to proceed.

3. Review the details of the Firewalls you are about to create. If you need to edit a particular detail, use the **Previous** button(s) to navigate back to the correct wizard page, and then use the **Next** button(s) to navigate back to the Summary page.

Note – If you go back to the Define Interfaces for the Multiple Single Firewall Elements page or to any page before that, you must re-define the interface(s) that the Firewalls use for connecting to the Internet.

4. Click **Finish**. The multiple Single Firewall elements are created.

What's Next?

- ▶ If you need to configure some settings separately for each individual Firewall, proceed to Editing Single Firewall Properties (page 395).
- ▶ Otherwise, the creation of multiple Single Firewall elements is done.

Creating a New Firewall Cluster Element

Firewall Cluster elements consist of 2 to 16 physical firewall devices that work together as a single entity.

▼ To create a new Firewall Cluster element

1. Select **Configuration→Configuration→Security Engine** from the menu. The Security Engine Configuration view opens.
2. Right-click **Security Engines** and select **New→Firewall Cluster**. The Firewall Cluster Properties dialog opens.

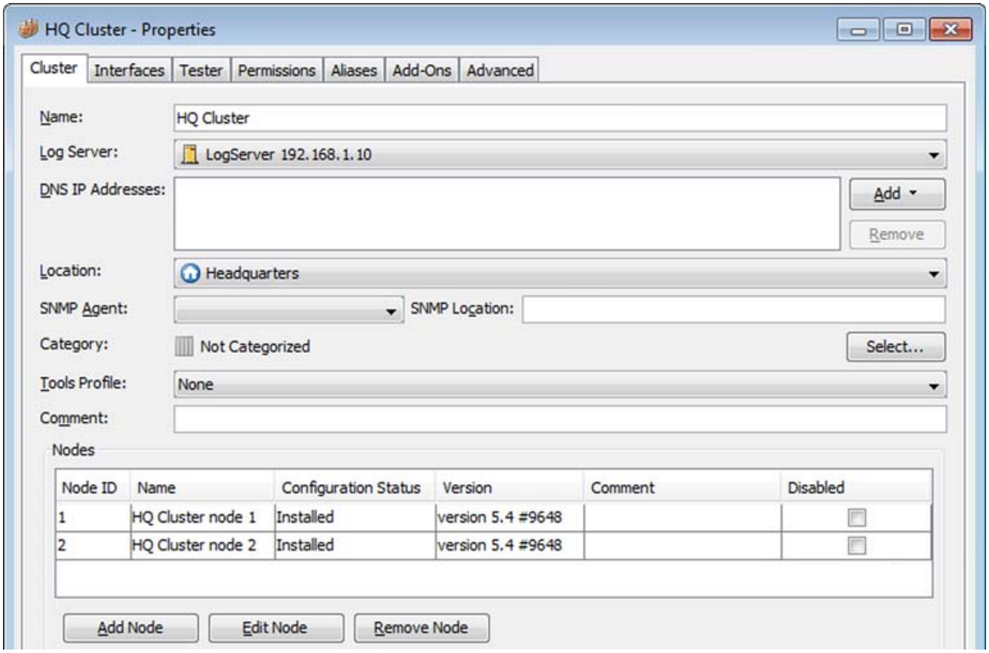

Name: HQ Cluster

Log Server: LogServer 192.168.1.10

DNS IP Addresses:

Location: Headquarters

SNMP Agent: SNMP Location:

Category: Not Categorized

Tools Profile: None

Comment:

Nodes

Node ID	Name	Configuration Status	Version	Comment	Disabled
1	HQ Cluster node 1	Installed	version 5.4 #9648		☐
2	HQ Cluster node 2	Installed	version 5.4 #9648		☐

3. Give the element a unique **Name**.
4. Select the **Log Server** to which the Firewall Cluster sends its event data.
5. (*Optional*) Define one or more **DNS IP Addresses**. These are the IP addresses of the DNS server(s) that the Firewall Cluster uses to resolve virus signature mirrors, domain names, and web filtering categorization services (which are defined as URLs). There are two ways to define IP addresses.
   - To enter a single IP address manually, click **Add** and select **IP Address**. Enter the IP address in the dialog that opens.
   - To define an IP address using a network element, click **Add** and select **Network Element**. Select an existing element, or click the New icon and define a new element. See Defining Host Elements (page 727) and Defining an External DNS Server (page 640) for more

information.

6. Select the **Location** for this Firewall Cluster if there is a NAT device between this Firewall Cluster and other system components. See Defining Locations (page 66) for more information.
7. (*Optional*) If you want to enable the Firewall Clusters to send SNMP traps, select an **SNMP Agent** as explained in Activating the SNMP Agent on Engines (page 584).
8. (*Optional*) If you want to include the Firewall Clusters in pre-defined categories, select the appropriate **Categories** as explained in Selecting Categories for Elements (page 77).
9. (*Optional*) If you want to add custom commands to the Firewall Clusters' right-click menu, add a **Tools Profile** as explained in Attaching a Tools Profile to an Element (page 56).

What's Next?

▶ Continue the configuration in Firewall Interface Configuration (page 421).

Creating Multiple Firewall Cluster Elements

Prerequisites: A Firewall Cluster has been created. See Creating a New Firewall Cluster Element.

You can create multiple Firewall Cluster elements at the same time using the Create Multiple Firewall Clusters Wizard. It is not possible to use plug-and-play configuration with multiple Firewall Clusters.

Once you have created the Firewall Clusters through the wizard, you can modify the properties of each cluster. See Editing Firewall Cluster Properties (page 396). You can also modify some shared properties of several Firewall Clusters at the same time. See Modifying Properties of Several Engines at Once (page 385).

▼ To create multiple Firewall Cluster elements

1. Select **Configuration→Configuration→Security Engine** from the menu. The Security Engine Configuration view opens.
2. Right-click **Security Engines** and select **New→Multiple Firewall Clusters**. The Create Multiple Firewall Clusters Wizard opens.

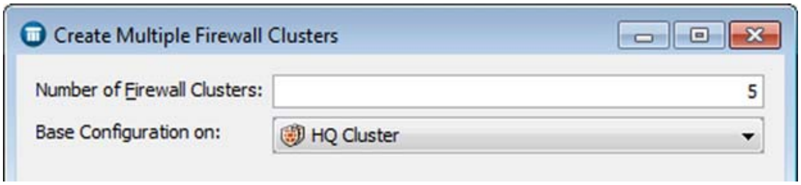

Number of Firewall Clusters: 5

Base Configuration on: HQ Cluster

3. Enter the **Number** of Firewall Cluster elements (1-1000).
4. Select a Firewall Cluster on which you want to base the Firewall Clusters' configuration from the **Base Configuration on** list and click **Next**. The Define Basic Firewall Information page opens.
5. Enter a common **Name Prefix**. The system adds a running number to the name prefix to generate a unique name for each individual Firewall Cluster.

6. Select a **Log Server** to which the Firewall Clusters send their event data.

Note – Name Prefix and Log Server are the only mandatory properties you must define at this stage. Review the other properties carefully to see which ones to define as shared properties for all Firewall Clusters created with the wizard.

7. (*Optional*) Define one or more **DNS IP Addresses**. These are the IP addresses of the DNS server(s) that the Firewall Clusters use to resolve virus signature mirrors, domain names, and web filtering categorization services (which are defined as URLs). There are two ways to define IP addresses.
   - To enter a single IP address manually, click **Add** and select **IP Address**. Enter the IP address in the dialog that opens.
   - To define an IP address using a network element, click **Add** and select **Network Element**. Select an existing element, or click the New icon and define a new element. See Defining Host Elements (page 727) and Defining an External DNS Server (page 640) for more information.
8. Select the **Location** for the Firewall Clusters if there is a NAT device between this Firewall Cluster and other system components. See Defining Locations (page 66) for more information.

Note – Select a Location only if all the Firewall Clusters you create with the wizard belong to the same Location element.

9. Define other settings according to your environment:
   - If you want to enable the Firewall Clusters to send SNMP traps, select an **SNMP Agent** as explained in Activating the SNMP Agent on Engines (page 584).
   - If you want to include the Firewall Clusters in pre-defined categories, select the appropriate **Categories** as explained in Selecting Categories for Elements (page 77).
   - If you want to add custom commands to the Firewall Clusters' right-click menu, add a **Tools Profile** as explained in Attaching a Tools Profile to an Element (page 56).
   - Add, edit, or remove nodes as needed as described in Adding a Node to a Cluster (page 392).
10. Click **Next**. The Review and Edit Names and Comments page opens.
11. Review the names of the Firewall Clusters. If necessary, right-click the name or comment and select **Edit Name** or **Edit Comment**.
12. Click **Next**. The Define Interfaces for the Multiple Firewall Elements page opens.

What's Next?

▶ Defining Interfaces for Multiple Firewall Clusters (page 376)

Defining Interfaces for Multiple Firewall Clusters

The following interface types are available:

- Physical Interface for Ethernet connections
- VLAN Interface for dividing a single physical network link into several virtual links
- Tunnel Interface that defines a tunnel end-point in Route-Based VPN.

You can define various properties for the interface(s) at this stage or return to them later. The same interface properties are available in the Properties dialog of every Firewall Cluster element. However, we recommend that you configure all the shared interface properties in the wizard. After you exit the wizard, you must configure the properties separately for each Firewall Cluster.

The interface properties you define for the first one of the Firewall Clusters are used to automatically create the corresponding properties for the rest of the Firewall Clusters. This also includes the IP addresses, which are automatically generated in numerical order. Make sure that the IP addresses that are assigned to the Firewall Clusters are not used by any other components.

If you want to use a physical interface as the control interface, begin the interface definition by defining a Physical Interface with an IP address, so that the Physical Interface is assigned Interface ID 0. When connecting the cables to the appliance, connect the cable for the control connection to Ethernet port 0. See the relevant *Appliance Installation Guide(s)* for detailed information on mapping the Interface IDs with specific ports on the appliance(s).

▼ To define interfaces for the Firewall Clusters

1. Click **Add** and select the type of interface. Add the required number of network interfaces in the following order:
   - Define Physical Interfaces as explained in Defining Physical Interfaces for Firewall Engines (page 422).
     - For settings on the **DHCP** tab, see Activating the Internal DHCP Server on a Firewall Interface (page 511) or, for DHCP relay, Routing DHCP Messages (page 590).
     - For settings on the **Advanced** tab, see Configuring Advanced Interface Properties for Firewalls (page 436).
   - Define Tunnel Interfaces as explained in Defining Tunnel Interfaces for Firewalls (page 435).
2. (*Optional, Physical Interfaces only)* Add the required number of VLANs as explained in Adding VLAN Interfaces for Firewall Engines (page 425).
3. Configure the IP address settings for the first one of the Firewall Clusters as explained in Configuring Firewall Cluster IP Addresses (page 444).
4. (*Optional*) Click **Options** to configure Loopback IP addresses as explained in Configuring Loopback IP Addresses for Firewalls (page 450).
5. Click **Options** to select which IP addresses are used in particular roles in system communications. See Setting Interface Options for Firewalls (page 453).
6. *(Optional)* Configure additional routing settings:
   - For adding **ARP entries**, see Configuring Manual ARP Settings (page 510).
   - For configuring **Multicast Routing**, see Routing Multicast Traffic (page 593).
7. Click **Next**. The Review and Edit Firewall Interfaces page opens.

8. Review the interfaces and edit them, if necessary.
   - To edit an interface, right-click the interface and select **Edit Interfaces**. For information on editing the IP address properties of an interface, see Configuring Firewall Cluster IP Addresses (page 444).
9. Click **Next**. The Define Routing for the Multiple Firewall Elements page opens.
10. Define and review the routing as explained in Defining Routing for Multiple Single Firewalls (page 365).
11. Click **Next**. The Select Additional Configuration Options page opens.

What's Next?

- ▶ If you want to define additional options, proceed to Selecting Additional Configuration Options for Multiple Firewall Clusters (page 377).
- ▶ Otherwise, click **Next** on the Select Additional Coinfiguration Options page. Proceed to the Summary page and click **Finish** to complete the creation of multiple Firewall Cluster elements.

Selecting Additional Configuration Options for Multiple Firewall Clusters

After defining the interfaces, you have configured all the mandatory properties for the Firewall Clusters. On the Select Additional Configuration Options page, you can select to define additional properties for the Firewall Clusters. It is recommended that you define all the shared properties of the Firewall Clusters in the wizard.

▼ To select additional configuration options

1. Select **Define Additional Firewall Properties** and click **Next**.
2. (*Optional*) Define tester settings as explained in Defining Tester Settings for Multiple Single Firewalls (page 366).
3. (*Optional*) Define permissions as explained in Defining Permissions for Multiple Single Firewalls (page 367).
4. (*Optional*) Define advanced settings as explained in Defining Advanced Settings for Multiple Single Firewalls (page 368).
5. On the Summary page, review the details of the Firewall Clusters you are about to create. If you need to edit a particular detail, use the **Previous** button(s) to navigate back to the correct wizard page, and then use the **Next** button(s) to navigate back to the Summary page.
6. Click **Finish** to add multiple Firewall Cluster elements to your system.

What's Next?

- ▶ If you need to configure some settings separately for each individual Firewall Cluster, proceed to Editing Firewall Cluster Properties (page 396).
- ▶ Otherwise, the creation of multiple Firewall Cluster elements is complete.

Creating a New Single IPS Element

IPS engines look for harmful patterns in traffic. Single IPS elements represent IPS engines that consist of one physical IPS device. They can be later converted to IPS clusters. See Converting a Single IPS Engine to an IPS Cluster (page 390).

▼ To create a new Single IPS element

1. Select **Configuration→Configuration→Security Engine** from the menu. The Security Engine Configuration view opens.
2. Right-click **Security Engines** tree and select **New→Single IPS**. The Single IPS Properties dialog opens.

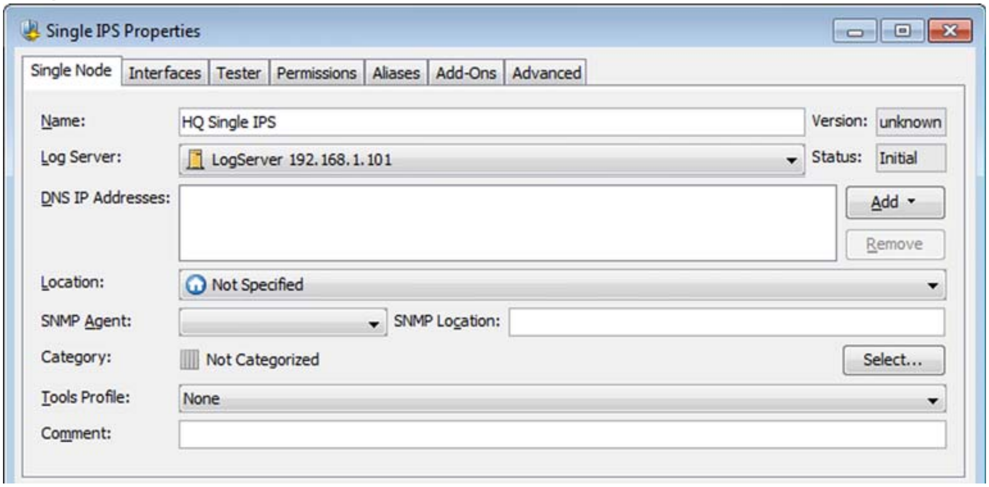

Name: HQ Single IPS
Version: unknown
Log Server: LogServer 192.168.1.101
Status: Initial
DNS IP Addresses:
Location: Not Specified
SNMP Agent:
SNMP Location:
Category: Not Categorized
Tools Profile: None
Comment:

3. Give the element a unique **Name**.
4. Select the **Log Server** to which the IPS engine sends its event data.
5. (*Optional*) Define one or more **DNS IP Addresses** for the IPS engine. These are the IP addresses of the DNS server(s) that the IPS engine uses to resolve domain names and web filtering categorization services (which are defined as URLs). There are two ways to define IP addresses.
   - To enter a single IP address manually, click **Add** and select **IP Address**. Enter the IP address in the dialog that opens.
   - To define an IP address using a network element, click **Add** and select **Network Element**. Select an existing element, or click the New icon and define a new element. See Defining Host Elements (page 727) and Defining an External DNS Server (page 640) for more information.
6. Select the **Location** for this engine if there is a NAT device between this IPS engine and other system components. See Configuring System Communications (page 63) for more information.
7. (*Optional*) If you want to enable the IPS engine to send SNMP traps, select an **SNMP Agent** as explained in Activating the SNMP Agent on Engines (page 584).
8. (*Optional*) If you want to include the IPS engine in pre-defined categories, select the appropriate **Categories** as explained in Selecting Categories for Elements (page 77).

9. (*Optional*) If you want to add custom commands to the IPS engine's right-click menu, add a **Tools Profile** as explained in Attaching a Tools Profile to an Element (page 56).

What's Next?

▶ Continue the configuration in IPS Engine Interface Configuration (page 456).

Creating a New IPS Cluster Element

IPS engines look for harmful patterns in traffic. IPS Cluster elements combine 2 to 16 physical IPS devices into a single entity.

▼ To create a new IPS Cluster element

1. Select **Configuration→Configuration→Security Engine** from the menu. The Security Engine Configuration view opens.
2. Right-click **Security Engines** and select **New→IPS Cluster**. The IPS Cluster Properties dialog opens.

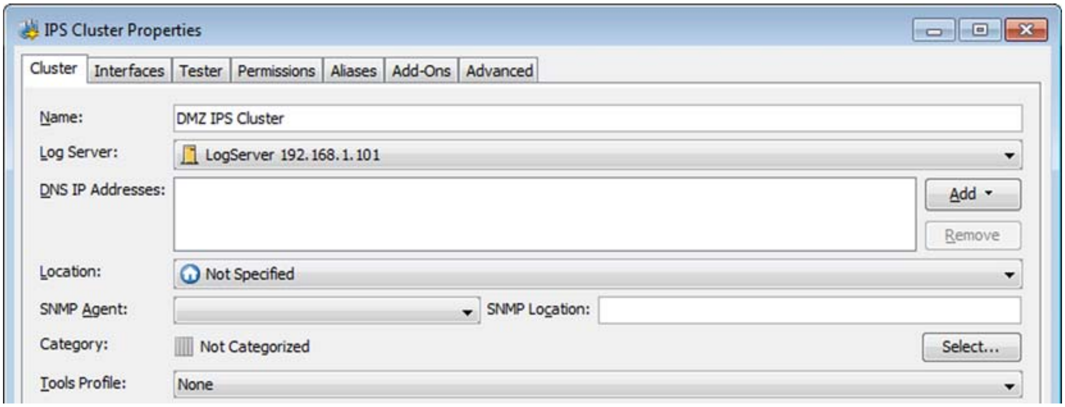

Name: DMZ IPS Cluster

Log Server: LogServer 192.168.1.101

DNS IP Addresses:

Location: Not Specified

SNMP Agent:

SNMP Location:

Category: Not Categorized

Tools Profile: None

3. Give the element a unique **Name**.
4. Select the **Log Server** to which the IPS Cluster sends its event data.
5. (*Optional*) Define one or more **DNS IP Addresses** for the IPS Cluster. These are the IP addresses of the DNS server(s) that the IPS Cluster uses to resolve domain names and web filtering categorization services (which are defined as URLs). There are two ways to define IP addresses.
   - To enter a single IP address manually, click **Add** and select **IP Address**. Enter the IP address in the dialog that opens.
   - To define an IP address using a network element, click **Add** and select **Network Element**. Select an existing element, or click the New icon and define a new element. See Defining Host Elements (page 727) and Defining an External DNS Server (page 640) for more information.
6. Select the **Location** for the IPS cluster if there is a NAT device between this IPS cluster and other system components. See Configuring System Communications (page 63) for more information.
7. (*Optional*) If you want to enable the IPS engine to send SNMP traps, select an **SNMP Agent** as explained in Activating the SNMP Agent on Engines (page 584).
8. (*Optional*) If you want to include the IPS engine in pre-defined categories, select the appropriate **Categories** as explained in Selecting Categories for Elements (page 77).

9. (*Optional*) If you want to add custom commands to the IPS cluster's right-click menu, add a **Tools Profile** as explained in Attaching a Tools Profile to an Element (page 56).

What's Next?

▶ Continue the configuration in IPS Engine Interface Configuration (page 456).

Creating a New Single Layer 2 Firewall Element

Layer 2 Firewalls look for harmful patterns in traffic. Single Layer 2 Firewall elements represent Layer 2 Firewalls that consist of one physical device. They can be later converted to Layer 2 Firewall Clusters. See Converting a Single Layer 2 Firewall to a Cluster (page 391).

▼ To create a new Single Layer 2 Firewall element

1. Select **Configuration→Configuration→Security Engine** from the menu. The Security Engine Configuration view opens.
2. Right-click **Security Engines** and select **New→Single Layer 2 Firewall**. The Single Layer 2 Firewall Properties dialog opens.

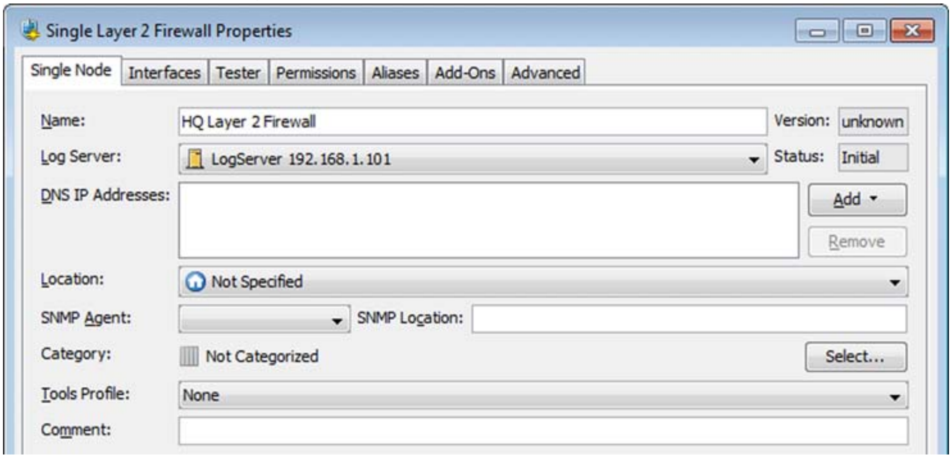

Single Layer 2 Firewall Properties

Name: HQ Layer 2 Firewall
Version: unknown
Log Server: LogServer 192.168.1.101
Status: Initial
DNS IP Addresses:
Location: Not Specified
SNMP Agent:
SNMP Location:
Category: Not Categorized
Tools Profile: None
Comment:

3. Give the element a unique **Name**.
4. Select the **Log Server** to which the Layer 2 Firewall sends its event data.
5. (*Optional*) Define the **DNS IP Addresses** for the Layer 2 Firewall. These are the IP addresses of the DNS server(s) that the Layer 2 Firewall uses to resolve domain names and web filtering categorization services (which are defined as URLs). There are two ways to define IP addresses.
   - To enter a single IP address manually, click **Add** and select **IP Address**. Enter the IP address in the dialog that opens.
   - To define an IP address using a network element, click **Add** and select **Network Element**. Select an existing element, or click the New icon and define a new element. See Defining Host Elements (page 727) and Defining an External DNS Server (page 640) for more information.
6. Select the **Location** for this engine if there is a NAT device between this Layer 2 Firewall and other system components. See Configuring System Communications (page 63) for more information.
7. (*Optional*) If you want to enable the Layer 2 Firewall engine to send SNMP traps, select an **SNMP Agent** as explained in Activating the SNMP Agent on Engines (page 584).

8. (*Optional*) If you want to include the Layer 2 Firewall engine in pre-defined categories, select the appropriate **Categories** as explained in Selecting Categories for Elements (page 77).

9. (*Optional*) If you want to add custom commands to the Layer 2 Firewall engine's right-click menu, add a **Tools Profile** as explained in Attaching a Tools Profile to an Element (page 56).

What's Next?

▶ Continue the configuration in Layer 2 Firewall Interface Configuration (page 473).

Creating a New Layer 2 Firewall Cluster Element

Layer 2 Firewalls look for harmful patterns in traffic. Layer 2 Firewall Cluster elements combine 2 to 16 physical Layer 2 Firewall devices into a single entity.

▼ To create a new Layer 2 Firewall Cluster element

1. Select **Configuration→Configuration→Security Engine** from the menu. The Security Engine Configuration view opens.

2. Right-click **Security Engines** and select **New→Layer 2 Firewall Cluster**. The Layer 2 Firewall Cluster Properties dialog opens.

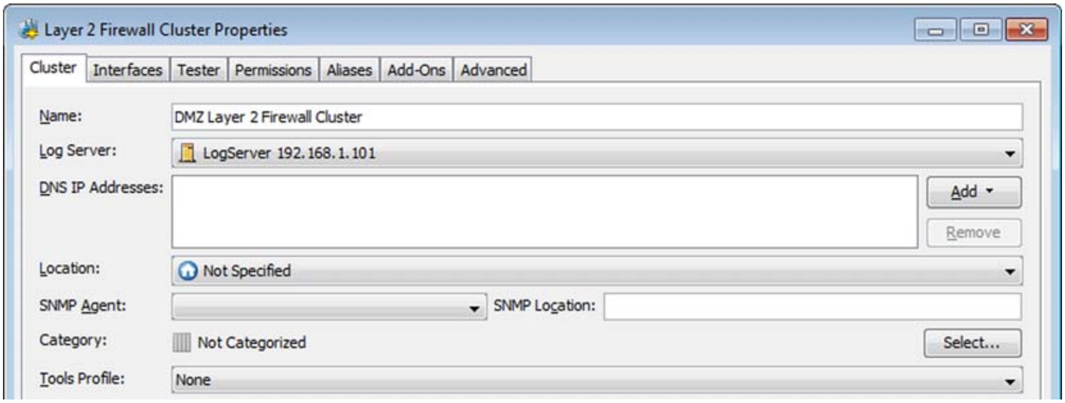

Name: DMZ Layer 2 Firewall Cluster
Log Server: LogServer 192.168.1.101
DNS IP Addresses:
Location: Not Specified
SNMP Agent:
SNMP Location:
Category: Not Categorized
Tools Profile: None

3. Give the element a unique **Name**.

4. Select the **Log Server** to which the Layer 2 Firewall Cluster sends its event data.

5. (*Optional*) Define one or more **DNS IP Addresses** for the Layer 2 Firewall Cluster. These are the IP addresses of the DNS server(s) that the Layer 2 Firewall Cluster uses to resolve domain names and web filtering categorization services (which are defined as URLs). There are two ways to define IP addresses.
   - To enter a single IP address manually, click **Add** and select **IP Address**. Enter the IP address in the dialog that opens.
   - To define an IP address using a network element, click **Add** and select **Network Element**. Select an existing element, or click the New icon and define a new element. See Defining Host Elements (page 727) and Defining an External DNS Server (page 640) for more information.

6. Select the **Location** for the Layer 2 Firewall Cluster if there is a NAT device between this Layer 2 Firewall Cluster and other system components. See Configuring System Communications (page 63) for more information.

7. (*Optional*) If you want to enable the Layer 2 Firewall Cluster to send SNMP traps, select an **SNMP Agent** as explained in Activating the SNMP Agent on Engines (page 584).
8. (*Optional*) If you want to include the Layer 2 Firewall Cluster in pre-defined categories, select the appropriate **Categories** as explained in Selecting Categories for Elements (page 77).
9. (*Optional*) If you want to add custom commands to the Layer 2 Firewall Cluster's right-click menu, add a **Tools Profile** as explained in Attaching a Tools Profile to an Element (page 56).

What's Next?

▶ Continue the configuration in Layer 2 Firewall Interface Configuration (page 473).

Creating a New SSL VPN Gateway Element

Stonesoft SSL VPN appliances can be monitored through the Management Client. For information on configuring Stonesoft SSL VPN, see the *Appliance Installation Guide* delivered with the appliance and the *Stonesoft SSL VPN Administrator's Guide*.

Note – You cannot change any of the SSL VPN gateway's settings in the Management Client. The configuration in the Management Client is only for establishing contact between the Management Server and the SSL VPN gateway.

▼ To create a new SSL VPN Gateway element

1. Select **Configuration→Configuration→Security Engine** from the menu. The Security Engine view opens.
2. Expand the **Network Elements** branch in the Security Engine tree.
3. Right-click **SSL VPN Gateways** and select **New SSL VPN Gateway**. The SSL VPN Gateway Properties dialog opens.

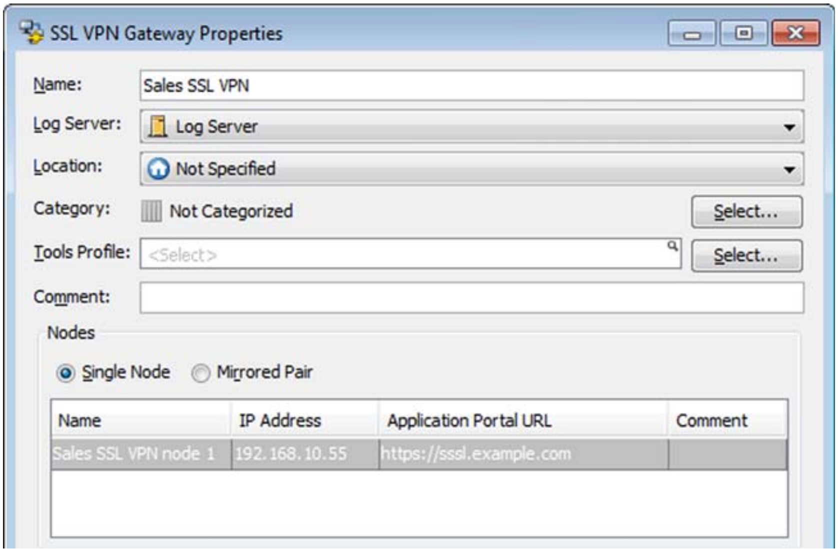

Name: Sales SSL VPN

Log Server: Log Server

Location: Not Specified

Category: Not Categorized

Tools Profile: <Select>

Comment:

Nodes

● Single Node ○ Mirrored Pair

Name	IP Address	Application Portal URL	Comment
Sales SSL VPN node 1	192.168.10.55	https://ssl.example.com	

4. Give the element a unique **Name**.
5. Select the **Log Server** to which the engine sends its event data.

6. Select the **Location** for this engine if there is a NAT device between this SSL VPN appliance and other system components. See Configuring System Communications (page 63) for more information.
7. If you have a mirrored pair of SSL VPN appliances, select **Mirrored Pair** above the Nodes table. In this element, it does not matter which of the nodes is the primary appliance.
8. Double-click a node in the Nodes table to open its properties.
9. Fill in the **Name**, local **IP Address**, and (*optional*) **Application Portal URL** for the node and click **OK**.
   - The Application Portal URL is used when you open the portal through the SSL VPN Gateway element's right-click menu in the Management Client.
   - This is also where you define the **Contact Address** information when necessary (see Configuring System Communications (page 63) for more information).
10. (*Mirrored Pair only*) Repeat Step 8-Step 9 for the other node.
11. Click **OK**.

What's Next?

▶ Continue by Connecting SSL VPN Gateways to the SMC (page 518).

Duplicating an Existing Engine Element

If you have similar configurations at several sites, you can create a duplicate of an existing single or clustered engine element to reduce the need for manual configuration work.

▼ To duplicate an existing engine element

1. Right-click the existing engine element and select **New→Duplicate**. The properties dialog for the new engine opens.
2. Give the engine element a unique **Name**.
3. Adjust the rest of the properties as needed (see Modifying Existing Engine Elements (page 384)). For all engine elements except SSL VPN Gateways, you must change at least the IP address of the control interface (or configure the duplicate to use a different control interface). The control interface must be unique for each engine. See Getting Started with Interface Configuration (page 420).
4. Click **OK**. The new engine element is added to the element tree.

Modifying Existing Engine Elements

Prerequisites: Creating New Engine Elements

Existing engine elements are shown in the System Status view under engine-specific branches. Firewall, IPS, and Layer 2 Firewall elements are shown in the Security Engine Configuration view under the Security Engines branch. SSL VPN Gateways are shown in the Security Engine Configuration view under the **Network Elements→SSL VPN Gateways** branch. You can modify the properties of an individual engine element or the common properties shared by several engine elements. Changing the common properties of several engine elements restricts you to actions and values that are applicable to all of them at the same time. You can also convert a Single Firewall, a Single IPS engine, or a Single Layer 2 Firewall into a cluster.

Related Tasks

- Modifying the Properties of Single Engine Elements (page 384)
- Modifying Properties of Several Engines at Once (page 385)
- Converting a Single Firewall to a Firewall Cluster (page 385)
- Converting a Single IPS Engine to an IPS Cluster (page 390)
- Converting a Single Layer 2 Firewall to a Cluster (page 391)
- Adding a Node to a Cluster (page 392)
- Changing Engine Control IP Address (page 393)

Modifying the Properties of Single Engine Elements

You can modify the properties specific to one individual engine element, such as IP address definitions, by editing the properties of the single engine element. You can also change some properties of multiple engines simultaneously. See Modifying Properties of Several Engines at Once (page 385).

▼ To view or modify any properties of an engine element

1. Right-click an engine element and select **Properties**. The properties dialog for the element opens.
2. Proceed to the one of the following sections depending on the type of the element:
   - Editing Single Firewall Properties (page 395)
   - Editing Firewall Cluster Properties (page 396)
   - Editing Single IPS Engine Properties (page 397)
   - Editing IPS Cluster Properties (page 398)
   - Editing Single Layer 2 Firewall Properties (page 399)
   - Editing Layer 2 Firewall Cluster Properties (page 400)
   - Creating a New SSL VPN Gateway Element (page 382)

Modifying Properties of Several Engines at Once

You can select several engine elements to change properties that are common to all of the elements.

Limitations

- Properties specific to one individual engine element, such as IP address definitions, are never available in the common properties. See Modifying the Properties of Single Engine Elements (page 384) to edit these properties.
- If you select both single and clustered engine elements, the cluster-specific options are not available.
- If you select elements of different types, you can only set the Log Server, Location, SNMP Agent, Comment options, and some system parameters.

▼ To modify common properties for several Firewalls, IPS Engines, or Layer 2 Firewalls

1. Shift-select or Ctrl-select the engine elements you want to modify.
2. Right-click one of the selected items and select **Common Properties**. The Common Engine Properties dialog opens.
3. Select the options you want to set for all of the selected engines. The options you see depend on how similar in type the selected engines are:
   - Editing Single Firewall Properties (page 395)
   - Editing Firewall Cluster Properties (page 396)
   - Editing Single IPS Engine Properties (page 397)
   - Editing IPS Cluster Properties (page 398)
   - Editing Single Layer 2 Firewall Properties (page 399)
   - Editing Layer 2 Firewall Cluster Properties (page 400)
   - Creating a New SSL VPN Gateway Element (page 382)
4. Click **OK**. All selected engines now use the same new settings for the options you changed.

Converting a Single Firewall to a Firewall Cluster

There is a conversion tool that allows you to change an existing Single Firewall engine into a Firewall Cluster. This maintains the relationship of the engine element with other configurations in the system (for example, VPNs), allows you to maintain some of the existing interface configurations (such as VLANs defined on the interfaces), and minimizes service interruptions. The conversion requires you to select one Single Firewall element to convert to a cluster.

Make sure that enough IP addresses are available, especially if the Single Firewall is managed remotely. Each clustered engine node needs at least one dedicated IP address for its management communications. Additionally, the traffic that the nodes inspect requires at least one dedicated IP address per cluster.

Limitations

- It is not possible to combine two Single Firewall elements into a Firewall Cluster.
- A Single Firewall engine can only be converted to a two-node Firewall Cluster. If you want to add more nodes to the cluster, you must add the nodes separately after the conversion as explained in Adding a Node to a Cluster (page 392).

- Due to differences in the supported configurations, there are some configurations that prevent you from converting from a Single Firewall to a Firewall Cluster. These are listed in the table below.

Table 26.1 Unsupported Configurations on Firewall Clusters

Configuration	Notes
ADSL interfaces	Firewall clusters do not support integrated ADSL modems. To convert to a cluster, you must switch to an external ADSL modem that the firewall engines access through an Ethernet connection.
Wireless interfaces	Firewall clusters do not support wireless interfaces.
Internal DHCP Server on older engine versions	Clustered firewalls support an internal DHCP server starting from software version 5.2. Upgrade the engine as necessary before conversion.
Dynamic IP addresses	Firewall clusters can only have static IP addresses. Clusters cannot use a dynamically assigned (DHCP or PPPoE) IP address.
Modem interfaces	Firewall clusters do not support integrated 3G modems. You must switch to a configuration that uses an external 3G modem through an Ethernet connection to convert to a cluster.

Caution – If you change the control IP address of the existing node in this process, the connection between the engine and the SMC is lost. See Changing Engine Control IP Address (page 393) to change the control IP address without losing contact.

Preparing for Converting a Single Firewall to a Firewall Cluster

▼ To prepare for converting a Single Firewall to a Firewall Cluster

1. If you are not using identical hardware, check that the performance levels match your needs. Equipment with different performance levels can be used for balancing the load between the clustered engines. For high availability when one engine is offline, the other engine must be capable of handling all traffic alone.
2. Make sure both firewall engines have their own license. Clustered firewall engines are licensed in the same way as two single firewall engines. All current firewall engine licenses allow clustering the nodes, so no license changes are needed to activate clustering.
   - For more information, see Getting Started with Licenses (page 1044).
3. Make sure the engines are running software versions that are compatible with the SMC, and preferably that both engines are running the same version. Although the cluster can be installed with the engines running different software versions (unless otherwise stated in the *Release Notes*), long-term use with mismatched versions is not supported.
   - For more information, see Upgrading Engines Remotely (page 1064) or the local upgrade instructions in the *Firewall/VPN Installation Guide*.
4. If the new Firewall engine you want to add to the cluster already has a working configuration from previous use, return it to the initial configuration state in the boot menu (factory reset) or through the `sg-reconfigure` wizard on the command line. Do not establish a connection with the Management Server before the Firewall Cluster element is ready.

- For more information, see Getting Started with the Engine Command Line (page 232).

Caution – If the firewall engine has a working configuration, it will go online and process traffic when you power it on to configure it for the firewall cluster.

5. Connect the network cables to the new firewall engine and power it on.

Converting a Single Firewall Element to a Firewall Cluster

Firewall Clusters have additional IP addressing requirements compared to Single Firewalls. The requirements are due to the two types of IP addresses that clusters need to function:

- An NDI (Node Dedicated IP Address) is used for communications between the engine itself and some other host in the network, such as the other nodes in the cluster, the Management Server, hosts you ping from the engine's command line, etc.
- A CVI (Cluster Virtual IP Address) is used for handling traffic that the cluster examines. If other network devices point to the firewall's IP address (as a default gateway or as a VPN endpoint, for example), converting the IP address to a CVI will allow those external configurations to remain unchanged.

The conversion tool requires you to convert the IP addresses used for these roles to the correct types. You must change the IP address that is used for a particular role if the new interface type is not compatible with that role.

The IP address requirements and related important considerations are listed in the table below.

Table 26.2 Interface Type Requirements by Role on Firewall Clusters

Role	Type Required	Notes
Control interface (Management connections)	NDI	Each node requires its own NDI address. Often, the same IP address on a Single Firewall is used for both the engine's own communications and the traffic that the engine processes. In such cases, you can convert the IP address that processes the traffic to a CVI to avoid reconfiguring external equipment, and add new NDI addresses for the nodes. Make sure that there are enough IP addresses available, especially if the firewall is managed remotely.
DHCP relay	CVI	Configured in the Physical Interface properties.
DHCP relay for VPN clients	NDI	Configured in the VPN Gateway properties.
Heartbeat interface	NDI	Heartbeat and state synchronization communications between clustered engines. We recommend using a dedicated interface for the heartbeat, as reliable transmissions are absolutely critical to the operation of the cluster. If the heartbeat traffic passes through a switch, make sure that the switch does not throttle or block multicast traffic between the clustered engines.

Table 26.2 Interface Type Requirements by Role on Firewall Clusters (Continued)

Role	Type Required	Notes
Routing	CVI	Traffic that is sent to an NDI address is not routed to any other destinations. Surrounding network devices that use the firewall as a default gateway must use a CVI address. If the internal DHCP server is used and configured to assign the firewall itself as the default gateway for clients, the default gateway IP address must be a CVI (configured in the Physical Interface properties).
VPN endpoints	CVI	Configured in the VPN Gateway properties.

▼ To modify a Single Firewall element for conversion

1. If you plan to convert the current IP address for management connections to a CVI:
   - **1a.** Add the new NDI IP address as a backup Control Interface address in the Single Firewall element.
   - **1b.** Adjust the Access and NAT rules of any firewalls on the communications path to allow both the current and new control IP addresses to be used and refresh the policies of these firewalls.
   - **1c.** Refresh the policy of the Single Firewall you plan to convert.
   - **1d.** Clear the Backup Control Interface selection and configure the new NDI control IP address as the Primary Control Interface address.
2. Add any new IP addresses that are required for the selected interface roles (see table above) and configure the settings to use those IP addresses.
3. If configured, remove dynamic IP addresses, Modem Interfaces, and ADSL Interfaces. These configurations are not supported on clusters.

▼ To run the element conversion tool and map the interface types

1. Right-click the Single Firewall element and select **Configuration→Upgrade to Cluster**. An interface mapping dialog opens.

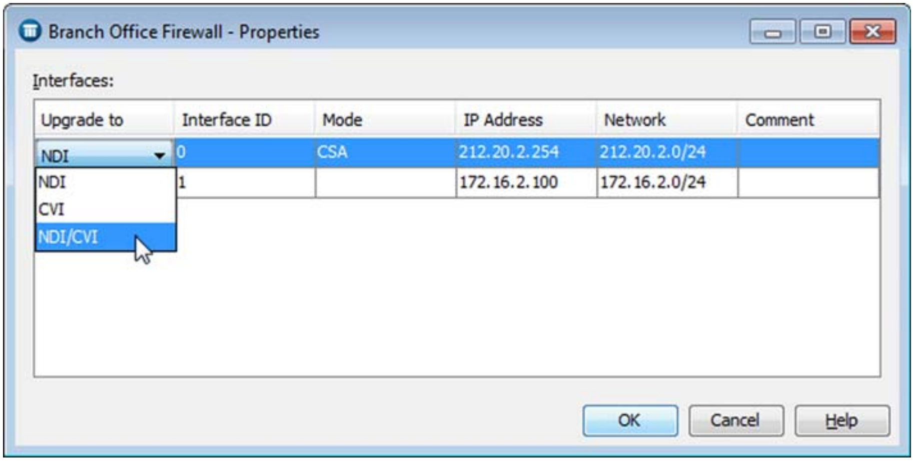

Branch Office Firewall - Properties

Interfaces:

Upgrade to	Interface ID	Mode	IP Address	Network	Comment
NDI	0	CSA	212.20.2.254	212.20.2.0/24	
	1		172.16.2.100	172.16.2.0/24	

2. Click the **Upgrade to** cell for each interface and select the IP address type(s) for the interfaces.

- You can select both a CVI and an NDI to be created for the same physical interface. This is the recommended working configuration for all interfaces, but may not be appropriate for all interfaces at this stage, since you cannot select which role the current IP address takes. Additional IP addresses are generated automatically to create the CVIs and NDIs.
- Each selection is validated and you may not be able to select a type if it is incompatible with the selected role of the interface. See the table above for a summary of requirements.

3. Click **OK**. The Cluster Properties dialog for the new Firewall Cluster element opens.
4. Switch to the **Interfaces** tab.
   - For help on defining the interfaces, see Getting Started with Interface Configuration (page 420).
5. Add the Physical interfaces and IP addresses that are needed in addition to those used on the single firewall. Check that the IP addresses on all interfaces are unique and unassigned, and change them if necessary.
6. Select **Packet Dispatch** as the CVI mode and enter the related unicast MAC address in the properties of all Physical Interfaces that have CVI definitions.
7. Click **Options** and define which IP addresses are used in particular roles in system communications as explained in Setting Interface Options for Firewalls (page 453).
8. If the internal DHCP server is used and configured to assign the firewall itself as the default gateway for clients, make sure the default gateway IP address is a CVI (Physical Interface properties, **DHCP** tab).
9. Click **OK**. The Single Firewall element is converted to a Firewall Cluster.
   - You can still click **Cancel** to return to the previous configuration and undo the conversion.

Activating the Clustered Configuration After Conversion

▼ To activate the Firewall Cluster configuration

1. If any external devices use the firewall as a default gateway or a VPN end-point and the previously used IP address is converted to an NDI, change the configurations of the external equipment to refer to a CVI address instead.
2. Connect to the engine's command line and run the sg-reconfigure command.
3. Make sure the interface IDs are mapped correctly to network ports according to the engine’s cabling.
4. Make initial contact between the engine node and the Management Server.
   - For more information, see Saving an Initial Configuration for Security Engines (page 515). Install and configure any new engine nodes as part of the cluster as in a completely new installation.
5. Install the policy on the cluster.
   - If you want to refresh the policy of the existing node before the new node(s) are initialized, you must mark the inactive nodes in the cluster element’s properties as disabled (on the **Cluster** tab). Otherwise, the policy installation fails due to a lack of connectivity to all nodes.

If there are problems with the clustered configuration, you can command one node offline through the right-click menu or through the command line to return to single-node operation.

Converting a Single IPS Engine to an IPS Cluster

There is a conversion tool that allows you to change an existing Single IPS engine into an IPS Cluster. This maintains the relationship of the engine element with other configurations in the system. The conversion requires you to select one Single IPS element to convert to an IPS Cluster.

For more information on how clustered IPS engines fit in the network architecture, see the *IPS and Layer 2 Firewall Reference Guide*.

Limitations

- It is not possible to combine two Single IPS elements into an IPS Cluster element.
- A Single IPS engine can only be converted to a two-node IPS Cluster. If you want to add more nodes to the cluster, you must add the nodes separately after the conversion as explained in Adding a Node to a Cluster (page 392).

Caution – If you change the control IP address of the existing node in this process, the connection between the engine and the SMC is lost. See Changing Engine Control IP Address (page 393) to change the control IP address without losing contact.

▼ To convert a Single IPS engine to an IPS Cluster

1. Make sure both engines are licensed. The licensing of clustered engine nodes is done in the same way as the licensing of two Single IPS engines. All current IPS engine licenses allow clustering the nodes, so no license changes are needed to activate the feature.
   - For more information, see Getting Started with Licenses (page 1044).
2. Make sure the engines are running software versions that are compatible with the SMC, and preferably that both engines are running the same version. Although the cluster can be installed with the engines running different software versions (unless otherwise stated in the *Release Notes*), long-term use with mismatched versions is not supported.
   - For more information, see Upgrading Engines Remotely (page 1064) or the local upgrade instructions in the *IPS and Layer 2 Firewall Installation Guide*.
3. If the new IPS engine you want to add to the IPS Cluster already has a working configuration from previous use, return it to the initial configuration state in the boot menu (factory reset) or through the `sg-reconfigure` wizard on the command line. Do not establish a connection with the Management Server before the IPS Cluster element is ready.
   - For more information, see Getting Started with the Engine Command Line (page 232).
4. Connect the network cables to the new node and power it on.
5. Right-click the Single IPS element that you want to upgrade to an IPS Cluster and select **Configuration→Upgrade to Cluster**. The IPS Cluster properties dialog opens.
6. Switch to the **Interfaces** tab.
7. Click **Options** and define which IP addresses are used in particular roles in system communications as explained in Setting Interface Options for IPS Engines (page 471).
8. Click **OK**. The Single IPS engine is converted to an IPS Cluster element.
   - You can still click **Cancel** to return to the previous configuration and undo the conversion.
9. Make initial contact between each node and the Management Server.
   - For more information, see Saving an Initial Configuration for Security Engines (page 515). Install and configure any new engine nodes as part of the cluster as in a completely new installation.

10. Install the policy on the IPS Cluster.

    - If you want to refresh the policy of the existing node before the new node(s) are initialized, you must disable the inactive nodes in the IPS Cluster element's properties (on the Cluster tab). Otherwise, the policy installation fails due to a lack of connectivity to all nodes.

Converting a Single Layer 2 Firewall to a Cluster

There is a conversion tool that allows you to change an existing Single Layer 2 Firewall into a Layer 2 Firewall Cluster. This maintains the relationship of the Single Layer 2 Firewall element and other configurations in the system. The conversion requires you to select one Single Layer 2 Firewall element that you convert to a Layer 2 Firewall Cluster.

For more information on how clustered Layer 2 Firewall engines fit in the network architecture, see the *IPS and Layer 2 Firewall Reference Guide*.

Limitations

- It is not possible to combine two Single Layer 2 Firewall elements into a Layer 2 Firewall Cluster element.
- A Single Layer 2 Firewall can only be converted to a two-node Layer 2 Firewall Cluster. If you want to add more nodes to the cluster, you must add the nodes separately after the conversion as explained in Adding a Node to a Cluster (page 392).

Caution – If you change the control IP address of the existing node in this process, the connection between the engine and the SMC is lost. See Changing Engine Control IP Address (page 393) to change the control IP address without losing contact.

▼ To convert a Single Layer 2 Firewall to a Layer 2 Firewall Cluster

1. Make sure both engines are licensed. The licensing of clustered engine nodes is done in the same way as the licensing of two Single Layer 2 Firewall engines. All current Layer 2 Firewall engine licenses allow clustering the nodes, so no license changes are needed to activate the feature.
    - For more information, see Getting Started with Licenses (page 1044).
2. Make sure the engines are running software versions that are compatible with the SMC, and preferably that both engines are running the same version. Although the cluster can be installed with the engines running different software versions (unless otherwise stated in the *Release Notes*), long-term use with mismatched versions is not supported.
    - For more information, see Upgrading Engines Remotely (page 1064) or the local upgrade instructions in the *IPS and Layer 2 Firewall Installation Guide*.
3. If the new Layer 2 Firewall engine you want to add to the Layer 2 Firewall Cluster already has a working configuration from previous use, return it to the initial configuration state in the boot menu (factory reset) or through the `sg-reconfigure` wizard on the command line. Do not establish a connection with the Management Server before the Layer 2 Firewall Cluster element is ready.
    - For more information, see Getting Started with the Engine Command Line (page 232).
4. Connect the network cables to the new node and power it on.
5. Right-click the Single Layer 2 Firewall element that you want to upgrade to a Layer 2 Firewall Cluster and select **Configuration→Upgrade to Cluster**. The Layer 2 Firewall Cluster properties dialog opens.
6. In the IPS Cluster element properties, switch to the **Interfaces** tab.

7. Click **Options** and define which IP addresses are used in particular roles in system communications as explained in Setting Interface Options for Layer 2 Firewalls (page 484).
8. Click **OK**. The Single Layer 2 Firewall engine is converted to a Layer 2 Firewall Cluster element.
   - You can still click **Cancel** to return to the previous configuration and undo the conversion.
9. Make initial contact between each node and the Management Server.
   - For more information, see Saving an Initial Configuration for Security Engines (page 515). Any new engine nodes are installed and configured as part of the cluster as in a completely new installation.
10. Install the policy on the Layer 2 Firewall Cluster.
    - If you want to refresh the policy of the existing node before the new node(s) are initialized, you must disable the inactive nodes in the Layer 2 Firewall Cluster element's properties (on the Cluster tab). Otherwise, the policy installation fails due to a lack of connectivity to all nodes.

Adding a Node to a Cluster

Nodes represent the physical Firewall, IPS, or Layer 2 Firewall engines. By default, the Cluster Properties dialog displays two nodes on the Cluster tab. You can add new nodes to the cluster as described below. Each cluster supports up to 16 nodes.

Before you add a node to the configuration, install the additional physical engine device and connect the cables for at least the control interface and the interface for communications between the clustered engines.

If the device already has a working configuration from previous use, return it to the initial configuration state in the `sg-reconfigure` wizard on the command line before connecting it to the network. Do not make initial contact with the Management Server.

▼ To add a node to a cluster

1. Right-click the cluster and select **Properties**. The Cluster Properties dialog opens.
2. Click the **Add Node** button at the bottom left corner of the dialog. The Engine Node Properties dialog opens.
3. Give the node a unique **Name**.
4. In the Nodes table, check and modify the information in the **IP Address**, **Contact IP Address**, and **Comment** columns for each NDI (double-click the value you want to change).
   - Firewall CVI details are identical between the nodes, so adding a node to a Firewall Cluster does not require changing the CVI configuration in any way.
5. Click **OK** in both open dialogs.
6. Save the initial configuration for the new engine node to create a one-time password. See Saving an Initial Configuration for Security Engines (page 515).
7. Make initial contact between the new engine and the Management Server. See the *Firewall/VPN Installation Guide* or the *IPS and Layer 2 Installation Guide*.
8. Refresh the policy of the cluster to transfer the changes to the engines.
   - If you want to refresh the policy of the existing node(s) before the new node(s) are initialized, you must disable the inactive node(s) in the cluster element's properties (on the Cluster tab).

Changing Engine Control IP Address

The following instructions explain how you can change an engine's control IP address without losing management connectivity.

When you change IP addressing, other connections between the different components may be temporarily lost. You must make sure that the connections return to normal after the IP address changes.

Proceed to one of the following:

- Changing Engine Control Address Within the Same Network
- Changing Firewall Control Address to a Different Network (page 394)

Changing Engine Control Address Within the Same Network

The section explains how you can change the control IP address of an engine if the new control IP address belongs to the same network as the old control IP address. The new control IP addresses of IPS engines and Layer 2 Firewalls must belong to the same network as the existing control IP addresses. The new control IP addresses of Firewalls can belong to the same network as the existing control IP address, or to a different network. If the new control IP address of a Firewall/VPN engine belongs to a different network than the engine's existing control IP address, see Changing Firewall Control Address to a Different Network (page 394).

If there is no need to maintain management connectivity, you can alternatively just change the control IP address in the SMC and then re-initialize the engine through the command line using a new one-time password. See Connecting Engines to the Management Center (page 513).

▼ To change the control IP address to an IP address from the same network

1. If you have an IP-address-bound license for the engine, request a new Management Server POL code bound license at the Stonesoft License center.
   - You must switch to a Management Server POL code bound license because IP-address-bound licences are no longer supported. See Getting Started with Licenses (page 1044).
2. Import and bind the new license to the engine.
3. In the engine properties, create a new interface for the new IP address and set the address as the backup control IP address (see Getting Started with Interface Configuration (page 420)).
4. Install the policy on the engine.
   - From this point on, you can start using the new address in the network.
5. In the engine properties, set the old control IP address as the backup control IP address and the new control IP address as the primary control IP address.

Note – If your engine cannot use both the old control IP address and a new control IP address simultaneously, proceed to Step 7.

6. Refresh the policy.
7. Remove the old control IP address from the engine properties and the corresponding network from the Routing view (see Getting Started with Routing (page 588)).

8. Refresh the policy.

Note – If the connection with the Management Server is lost when you try to change IP addressing, run the `sg-reconfigure` command on the engine command line to return the engine to the initial configuration state and to re-establish initial contact between the engine and the Management Server. See Connecting Engines to the Management Center (page 513).

Changing Firewall Control Address to a Different Network

The section explains how you can change the control IP address of a Firewall/VPN engine when the new IP address is in a different network than the old one. Because these steps require the configuration of Outbound Multi-Link, you can only change the control IP address of Firewalls to a different network. For all other engine role, or if the Firewall/VPN engine's new control IP address belongs to the same network as the old control IP address, see Changing Engine Control Address Within the Same Network (page 393).

If there is no need to maintain management connectivity, you can alternatively just change the control IP address in the SMC and then re-initialize the engine through the command line using a new one-time password (see Connecting Engines to the Management Center (page 513)).

▼ To change the Firewall/VPN engine's control IP address to an address from a different network

1. If you have an IP-address-bound license for the engine, request a new Management Server POL code bound license at the Stonesoft License center.
   - You must switch to a Management Server POL code bound license because IP-address-bound licences are no longer supported. See Getting Started with Licenses (page 1044).
2. Import and bind the new license to the engine.
3. Edit the properties of the Single Firewall or Firewall Cluster element and add a new interface (see Modifying Existing Engine Elements (page 384)).
   - Define the new primary control address as the backup control IP address.
   - If your firewall is a cluster and you do not want to lose any connections, define also a new CVI for the cluster (see Configuring Firewall Cluster IP Addresses (page 444)).
4. Configure Outbound Multi-Link as explained in Outbound Traffic Management (page 609).
   - Create two NetLinks: one for the old control IP address and one for the new control IP address.
5. Install the policy on the firewall.
   - From this point on, you can start using the new address in the network.
6. Edit the properties of the Single Firewall or Firewall Cluster element and set the new control IP address as the primary control IP address and the old control IP address as the backup control IP address.

Note – If your engine cannot use both the old control IP address and a new control IP address simultaneously, proceed to Step 8.

7. Refresh the policy.
8. Remove the interface with the old control IP address from the firewall's properties.
9. Remove the elements and rules you created for the Multi-Link configuration in Step 4.

10. Refresh the policy.

> **Note – If the connection with the Management Server is lost when you try to change IP addressing, run the `sg-reconfigure` command on the engine command line to return the engine to the initial configuration state and to re-establish initial contact between the engine and the Management Server. See Connecting Engines to the Management Center (page 513).**

Editing Single Firewall Properties

Prerequisites: Creating New Engine Elements

If you are creating a new Single Firewall, proceed as explained in Creating a New Single Firewall Element (page 359).

Table 26.3 Tasks to Do in Single Firewall Properties Dialog

Tab	Tasks
Single Node	For defining the basic properties, see Creating a New Single Firewall Element (page 359). For defining Locations and Contact addresses, see Defining Engine Location (page 67). For selecting an SNMP Agent, see Activating the SNMP Agent on Engines (page 584). For selecting Categories, see Using Categories (page 76). For defining a Tools Profile, see Adding Custom Commands to Element Menus (page 55).
Interfaces	For adding or modifying interfaces, see Firewall Interface Configuration (page 421). For defining interface options, see Setting Interface Options for Firewalls (page 453). For adding or modifying ARP entries, see Configuring Manual ARP Settings (page 510). For configuring multicast routing, see Routing Multicast Traffic (page 593).
Tester	See Getting Started with the Engine Tester (page 520).
Permissions	See Getting Started with Engine Permissions (page 530). For more information on Administrator privilege levels, see Getting Started with Administrator Accounts (page 248).
Aliases	See Getting Started with Alias Translations (page 534).

Table 26.3 Tasks to Do in Single Firewall Properties Dialog (Continued)

Tab	Tasks
Add-Ons	For configuring TLS inspection, see Activating TLS inspection on the Engine (page 803). To select the User Agent for Access Control by User, see Enabling Access Control by User (page 838). For configuring browser-based user authentication, see Enabling Browser-Based User Authentication (page 874). For defining anti-virus settings, see Configuring Anti-Virus Settings (page 540). For defining anti-spam settings, see Configuring Anti-Spam Settings (page 541).
Advanced	See Adjusting Firewall System Parameters (page 553), and Adjusting Firewall Traffic Handling Parameters (page 554).

Editing Firewall Cluster Properties

Prerequisites: Creating New Engine Elements

If you are creating a new Firewall Cluster, proceed as explained in Creating a New Firewall Cluster Element (page 373).

Table 26.4 Tasks to Do in Firewall Cluster Properties Dialog

Tab	Tasks
Cluster	For defining the basic properties, see Creating a New Firewall Cluster Element (page 373). For defining Locations and Contact addresses, see Defining Engine Location (page 67). For selecting an SNMP Agent, see Activating the SNMP Agent on Engines (page 584). For selecting Categories, see Using Categories (page 76). For defining a Tools Profile, see Adding Custom Commands to Element Menus (page 55). To add a node to the cluster, see Adding a Node to a Cluster (page 392).
Interfaces	For adding or modifying interfaces, see Firewall Interface Configuration (page 421). For defining interface options, see Setting Interface Options for Firewalls (page 453). For adding or modifying ARP entries, see Configuring Manual ARP Settings (page 510). For configuring multicast routing, see Routing Multicast Traffic (page 593).
Tester	See Getting Started with the Engine Tester (page 520).
Permissions	See Getting Started with Engine Permissions (page 530). For more information on Administrator privilege levels, see Getting Started with Administrator Accounts (page 248).
Aliases	See Getting Started with Alias Translations (page 534).

Table 26.4 Tasks to Do in Firewall Cluster Properties Dialog (Continued)

Tab	Tasks
Add-Ons	For configuring TLS inspection, see Activating TLS inspection on the Engine (page 803). To select the User Agent for Access Control by User, see Enabling Access Control by User (page 838). For configuring browser-based user authentication, see Enabling Browser-Based User Authentication (page 874). For defining anti-virus settings, see Configuring Anti-Virus Settings (page 540). For defining anti-spam settings, see Configuring Anti-Spam Settings (page 541).
Advanced	See Adjusting Firewall System Parameters (page 553), and Adjusting Firewall Traffic Handling Parameters (page 554).

Editing Single IPS Engine Properties

Prerequisites: Creating New Engine Elements

If you are creating a new Single IPS engine, proceed as explained in Creating a New Single IPS Element (page 378).

Table 26.5 Tasks to Do in Single IPS Properties Dialog

Tab	Tasks
Single Node	For defining the basic properties, see Creating a New Single IPS Element (page 378). For defining Locations and Contact addresses, see Defining Engine Location (page 67). For selecting an SNMP Agent, see Activating the SNMP Agent on Engines (page 584). For selecting Categories, see Using Categories (page 76). For defining a Tools Profile, see Adding Custom Commands to Element Menus (page 55).
Interfaces	For adding or modifying interfaces, see IPS Engine Interface Configuration (page 456). For defining interface options, see Setting Interface Options for IPS Engines (page 471). For adding or modifying ARP entries, see Configuring Manual ARP Settings (page 510).
Tester	See Getting Started with the Engine Tester (page 520).
Permissions	See Getting Started with Engine Permissions (page 530). For more information on Administrator privilege levels, see Getting Started with Administrator Accounts (page 248).
Aliases	See Defining Alias Translation Values (page 534).

Table 26.5 Tasks to Do in Single IPS Properties Dialog (Continued)

Tab	Tasks
Add-Ons	For configuring TLS inspection, see Activating TLS inspection on the Engine (page 803). To select the User Agent for Access Control by User, see Enabling Access Control by User (page 838).
Advanced	See Adjusting IPS Engine System Parameters (page 560) and Adjusting IPS Engine Traffic Parameters (page 561).

Editing IPS Cluster Properties

Prerequisites: Creating New Engine Elements

If you are creating a new IPS Cluster, proceed as explained in Creating a New IPS Cluster Element (page 379).

Table 26.6 Tasks to Do in IPS Cluster Properties Dialog

Tab	Tasks
Cluster	For defining the basic properties, see Creating a New IPS Cluster Element (page 379). For defining Locations and Contact addresses, see Defining Engine Location (page 67). For selecting an SNMP Agent, see Activating the SNMP Agent on Engines (page 584). For selecting Categories, see Using Categories (page 76). For defining a Tools Profile, see Adding Custom Commands to Element Menus (page 55). To add a node to the cluster, see Adding a Node to a Cluster (page 392).
Interfaces	For adding or modifying interfaces, see IPS Engine Interface Configuration (page 456). For defining interface options, see Setting Interface Options for IPS Engines (page 471). For adding or modifying ARP entries, see Configuring Manual ARP Settings (page 510).
Tester	See Getting Started with the Engine Tester (page 520).
Permissions	See Getting Started with Engine Permissions (page 530). For more information on Administrator privilege levels, see Getting Started with Administrator Accounts (page 248).
Aliases	See Defining Alias Translation Values (page 534).
Add-Ons	For configuring TLS inspection, see Activating TLS inspection on the Engine (page 803). To select the User Agent for Access Control by User, see Enabling Access Control by User (page 838).

Table 26.6 Tasks to Do in IPS Cluster Properties Dialog (Continued)

Tab	Tasks
Advanced	See Adjusting IPS Engine System Parameters (page 560) and Adjusting IPS Engine Traffic Parameters (page 561).

Editing Single Layer 2 Firewall Properties

Prerequisites: Creating New Engine Elements

If you are creating a new Single Layer 2 Firewall, proceed as explained in Creating a New Single Layer 2 Firewall Element (page 380).

Table 26.7 Tasks to Do in Single Layer 2 Firewall Properties Dialog

Tab	Tasks
Single Node	For defining the basic properties, see Creating a New Single Layer 2 Firewall Element (page 380). For defining Locations and Contact addresses, see Defining Engine Location (page 67). For selecting an SNMP Agent, see Activating the SNMP Agent on Engines (page 584). For selecting Categories, see Using Categories (page 76). For defining a Tools Profile, see Adding Custom Commands to Element Menus (page 55).
Interfaces	For adding or modifying interfaces, see Layer 2 Firewall Interface Configuration (page 473). For defining interface options, see Setting Interface Options for Layer 2 Firewalls (page 484). For adding or modifying ARP entries, see Configuring Manual ARP Settings (page 510).
Tester	See Getting Started with the Engine Tester (page 520).
Permissions	See Getting Started with Engine Permissions (page 530). For more information on Administrator privilege levels, see Getting Started with Administrator Accounts (page 248).
Aliases	See Defining Alias Translation Values (page 534).
Add-Ons	For configuring TLS inspection, see Activating TLS inspection on the Engine (page 803). To select the User Agent for Access Control by User, see Enabling Access Control by User (page 838).
Advanced	See Adjusting Layer 2 Firewall System Parameters (page 564) and Adjusting Layer 2 Firewall Traffic Parameters (page 565).

Editing Layer 2 Firewall Cluster Properties

Prerequisites: Creating New Engine Elements

If you are creating a new Layer 2 Firewall Cluster, proceed as explained in Creating a New Layer 2 Firewall Cluster Element (page 381).

Table 26.8 Tasks to Do in Layer 2 Firewall Cluster Properties Dialog

Tab	Tasks
Cluster	For defining the basic properties, see Creating a New Layer 2 Firewall Cluster Element (page 381). For defining Locations and Contact addresses, see Defining Engine Location (page 67). For selecting an SNMP Agent, see Activating the SNMP Agent on Engines (page 584). For selecting Categories, see Using Categories (page 76). For defining a Tools Profile, see Adding Custom Commands to Element Menus (page 55). To add a node to the cluster, see Adding a Node to a Cluster (page 392).
Interfaces	For adding or modifying interfaces, see Layer 2 Firewall Interface Configuration (page 473). For defining interface options, see Setting Interface Options for Layer 2 Firewalls (page 484). For adding or modifying ARP entries, see Configuring Manual ARP Settings (page 510).
Tester	See Getting Started with the Engine Tester (page 520).
Permissions	See Getting Started with Engine Permissions (page 530). For more information on Administrator privilege levels, see Getting Started with Administrator Accounts (page 248).
Aliases	See Defining Alias Translation Values (page 534).
Add-Ons	For configuring TLS inspection, see Activating TLS inspection on the Engine (page 803). To select the User Agent for Access Control by User, see Enabling Access Control by User (page 838).
Advanced	See Adjusting Layer 2 Firewall System Parameters (page 564) and Adjusting Layer 2 Firewall Traffic Parameters (page 565).

Adjusting the Global Contact Policy for Single Engines

Prerequisites: A single Firewall, IPS engine, or Layer 2 Firewall with a dynamic Control IP Address

Communication between the Management Server and a single Firewall, IPS, or Layer 2 Firewall can be reversed so that the engine opens a connection to the Management Server and keeps it open to wait for any commands. This can be necessary if the engine does not have a static IP address that the Management Server can contact (dynamic IP address on the interface or intermediate dynamic NAT) or if the Management Server's connections are blocked because of a traffic filtering device between the components.

The settings for communication between the Management Server and the engines are set in the `SGConfiguration.txt` file stored on the Management Server. You can either use the default values for each setting or modify the settings by adding parameters and values to the `SGConfiguration.txt` file.

Note – If you have edited a Firewall engine's contact policy settings prior to SMC 5.4, adjust the Firewall engine's contact settings directly through the Advanced tab in the engine properties. See the *StoneGate Administrator's Guide* for SMC 5.3 for more information.

▼ To define global contact policy settings for single engines

1. Browse to the `<installation directory>/data` directory on the Management Server.

Note – If you installed the Management Server in the `C:\Program Files\Stonesoft\Management Center` directory in Windows, some program data may be stored in the `C:\ProgramData\Stonesoft\Management Center\data` directory.

2. Edit `SGConfiguration.txt` and add the following parameters as needed:

Parameter Name	Description
`DCP_INITIAL_DELAY`	Time (in seconds) to wait after initialization before the first connection attempt to the Management Server. The default value is 5 seconds.
`DCP_CONNECTION_INTERVAL`	Time (in seconds) to wait before connecting again to the Management Server after a successful connection. The default value is 25 seconds.
`DCP_RETRY_INTERVAL`	Time (in seconds) to wait before connecting again to the Management Server after a failed connection attempt. The default value is 25 seconds.
`DCP_IDLE_TIMEOUT`	Time (in seconds) before an idle connection is closed. The default value is 1800 seconds (30 minutes).

3. Save and close the file.

What's Next?

▶ Refresh the policies of the engines to transfer the changes.

Related Tasks

- For the option to change the management contact direction, see Setting Interface Options for Firewalls (page 453), Setting Interface Options for IPS Engines (page 471), and Setting Interface Options for Layer 2 Firewalls (page 484).

About Engine Time Synchronization

Prerequisites: None

Engine times are automatically synchronized to match the time of the Management Server. If you want to use an NTP server to synchronize engine times, you must set the computer on which you have installed the Management Server to synchronize its time with the NTP server. If the Log Server runs on a different computer, set it to use the same NTP server.

If the Management Server, Log Server, and the engines do not have the same time, there may be problems with logging and monitoring. Also, make sure that the computer you use for Management Client access has the time and time zone set correctly to avoid time synchronization problems when you view statistics or logs, generate reports, or schedule automatic maintenance tasks.

CHAPTER 27

CREATING AND MODIFYING VIRTUAL SECURITY ENGINES

Virtual Security Engines are logically-separate engines that run as virtual machines on a physical engine device. A Master Engine is a physical engine device that provides resources for Virtual Security Engines.

The following sections are included:

Getting Started with Virtual Engines

Prerequisites: None

What Virtual Engines Do

Using Virtual Security Engines allows the same physical engine device to support multiple policies or routing tables, or policies that involve overlapping IP addresses. This is especially useful in a Managed Security Service Provider (MSSP) environment, or in a network environment that requires strict isolation between networks.

Limitations

- Only the Firewall/VPN role is supported on Virtual Security Engines.
- Only 64-bit appliances can be used as a Master Engine.
- Virtual Security Engines do not support Applications, URL filtering, Anti-Virus, Anti-Spam, Dynamic IP addresses, or Wireless Interfaces.
- If there are multiple administrative Domains, the Master Engine must either belong to the Shared Domain or to the same Domain as the Virtual Security Engines. Getting Started with Domains (page 284) for more information about Domains.
- Virtual Security Engines handle only the traffic routed through the engine for inspection. All other traffic, including communication between the Virtual Security Engines and the SMC, is proxied by the Master Engine. Virtual Security Engines do not communicate directly with other Virtual Security Engines.

What Should I Know Before I Begin?

Before you define a new Master Engine element, make sure you have a license for it. The element can be configured without a license, but you must have a license to make the engine operational. Virtual Security Engines do not require individual licenses. Instead, the Security Engine license for the Master Engine defines how many Virtual Resources can be created. The number of Virtual Resources limits the number of Virtual Security Engines: one Virtual Security Engine at a time can be associated with each Virtual Resource.

Configuration Overview

1. Create a Master Engine as explained in Creating New Master Engine Elements (page 405).
2. Create a Virtual Resource element as explained in Creating New Virtual Resource Elements (page 406).
3. Configure Physical or VLAN interfaces for the Master Engine and assign Virtual Resources to the interfaces as explained in Master Engine Interface Configuration (page 486).
4. Create a Virtual Security Engine as explained in Creating New Virtual Security Engines (page 407).
5. Configure Physical, VLAN, and/or Tunnel Interfaces for the Virtual Security Engine as explained in Virtual Security Engine Interface Configuration (page 498).
6. Install or refresh the policy on the Master Engine to transfer changes to Physical/VLAN interfaces and interface mapping for Virtual Security Engines.
7. Install or refresh the policy on the Virtual Security Engine(s).

Creating New Master Engine Elements

Prerequisites: None

Master Engines are physical engine devices that provide the resources for Virtual Security Engines. One physical Master Engine can support multiple Virtual Security Engines.

▼ To create new Master Engine elements

1. Select **Configuration→Configuration→Security Engine**. The Security Engine Configuration view opens.
2. Right-click **Security Engines** and select **New→Master Engine**. The Master Engine Properties dialog opens.

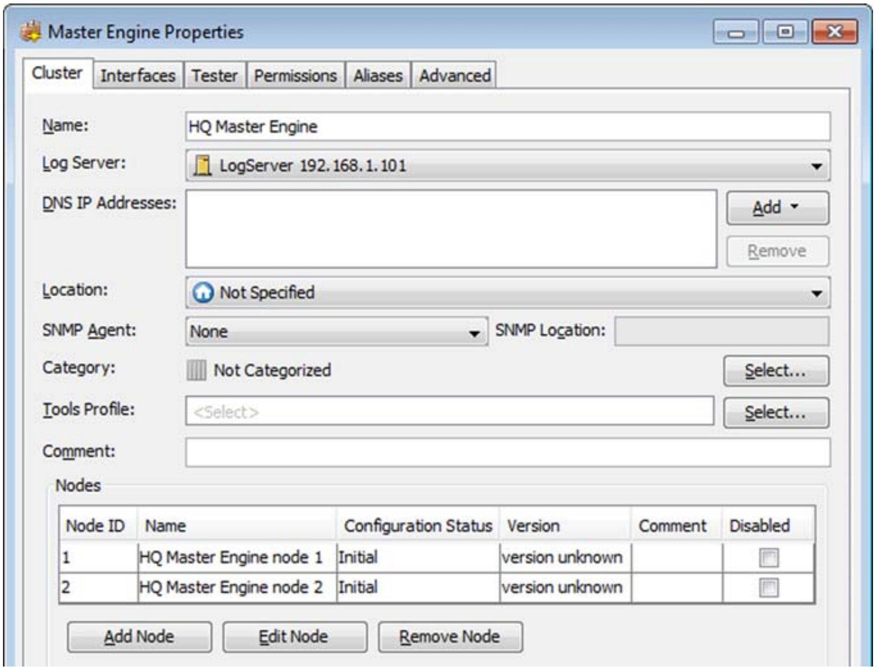

Name: HQ Master Engine

Log Server: LogServer 192.168.1.101

DNS IP Addresses:

Location: Not Specified

SNMP Agent: None

SNMP Location:

Category: Not Categorized

Tools Profile: <Select>

Comment:

Node ID	Name	Configuration Status	Version	Comment	Disabled
1	HQ Master Engine node 1	Initial	version unknown		□
2	HQ Master Engine node 2	Initial	version unknown		□

3. Give the element a unique **Name**.
4. Select the **Log Server** to which the Master Engine sends its event data. The Master Engine also sends log data from the Virtual Security Engines to the same Log Server.
5. (*Optional*) Define one or more **DNS IP Addresses**. These are the IP addresses of the DNS server(s) that the Master Engine uses to resolve domain names. There are two ways to define IP addresses.
   - To enter a single IP address manually, click **Add** and select **IP Address**. Enter the IP address in the dialog that opens.
   - To define an IP address using a network element, click **Add** and select **Network Element**. Select an existing element, or click the New icon and define a new element. See Defining Host Elements (page 727) and Defining an External DNS Server (page 640) for more information.
6. Select the **Location** for this Master Engine if there is a NAT device between this Master Engine and other system components. See Defining Locations (page 66) for more information.

7. (*Optional*) If you want to enable the Master Engine to send SNMP traps, select an **SNMP Agent** as explained in Activating the SNMP Agent on Engines (page 584).
8. (*Optional*) If you want to include the Master Engine in pre-defined categories, select the appropriate **Categories** as explained in Selecting Categories for Elements (page 77).
9. (*Optional*) If you want to add custom commands to the Master Engine's right-click menu, add a **Tools Profile** as explained in Attaching a Tools Profile to an Element (page 56).
10. (*Optional*) If you do not need to use clustering on the Master Engine, select one of the nodes and click **Remove Node**. You are prompted to confirm that you want to delete the selected node. Click **Yes**.

What's Next?

▶ Continue by Creating New Virtual Resource Elements

Creating New Virtual Resource Elements

Prerequisites: Creating New Master Engine Elements

Virtual Resources define the set of resources on the Master Engine that are allocated to each Virtual Security Engine.

▼ To create Virtual Resource elements

1. Switch to the **Interfaces** tab of the Master Engine Properties and click **Virtual Resources**. The Virtual Resources dialog opens.
2. Click **Add**. The Virtual Resource Properties dialog opens.
3. Enter a unique **Name** for the Virtual Resource.
4. Select the **Domain** to which the Virtual Resource belongs.
5. (*Optional*) Select **Show Master Interface IDs in Virtual Engine** if you want the Physical Interface IDs of the Master Engine to be shown in the Interface properties of the Virtual Security Engine.
6. Click **OK**.

What's Next?

▶ If you are creating a new Master Engine, continue the configuration as explained in Master Engine Interface Configuration (page 486).

▶ Otherwise, associate the Virtual Resource with a Master Engine interface as explained in Master Engine Interface Configuration (page 486) and with a Virtual Security Engine as explained in Creating New Virtual Security Engines.

Creating New Virtual Security Engines

Prerequisites: Creating New Virtual Resource Elements

Only the Firewall/VPN role is supported on Virtual Security Engines. Virtual Security Engines do not require individual licenses.

▼ To create new Virtual Security Engines

1. Select **Configuration→Configuration→Security Engine**. The Security Engine Configuration view opens.
2. Right-click **Security Engines** and select **New→Virtual Firewall**. The Virtual Firewall Properties dialog opens.
3. Give the element a unique **Name**.
4. Click **Select** and select a Virtual Resource on the Master Engine to which you want to add the Virtual Firewall.
5. (*Optional*) Define one or more **DNS IP Addresses**. These are the IP addresses of the DNS server(s) that the Virtual Security Engine uses to resolve domain names. There are two ways to define IP addresses.
   - To enter a single IP address manually, click **Add** and select **IP Address**. Enter the IP address in the dialog that opens.
   - To define an IP address using a network element, click **Add** and select **Network Element**. Select an existing element, or click the New icon and define a new element. See Defining Host Elements (page 727) and Defining an External DNS Server (page 640) for more information.
6. (*Optional*) If you want to include the Virtual Security Engine in pre-defined categories, select the appropriate **Categories** as explained in Selecting Categories for Elements (page 77).

What's Next?

▶ Continue the configuration in Virtual Security Engine Interface Configuration (page 498).

Modifying Existing Master Engines and Virtual Security Engines

Prerequisites: See Configuration Overview

Existing Master Engine and Virtual Firewall elements are shown in the System Status view under engine-specific branches, and in the Security Engine Configuration view under the Security Engines branch. You can also convert an existing Single Firewall or Firewall Cluster into a Master Engine and Virtual Firewalls.

Adding a Node to a Master Engine

By default, the Master Engine Properties dialog displays two nodes on the Cluster tab. You can add new nodes to the Master Engine as described below. Each Master Engine supports up to 16 nodes.

Before you add a node to the configuration, install the additional physical engine device and connect the cables for at least the control interface and the interface for communications between the clustered engines.

If the device already has a working configuration from previous use, return it to the initial configuration state in the `sg-reconfigure` wizard on the command line before connecting it to the network. Do not make initial contact with the Management Server.

▼ To add a node to a Master Engine

1. Right-click the Master Engine and select **Properties**. The Master Engine Properties dialog opens.
2. Click the **Add Node** button at the bottom left corner of the dialog. The Engine Node Properties dialog opens.
3. Give the node a unique **Name**.
4. In the Nodes table, check and modify the information in the **IP Address**, **Contact IP Address**, and **Comment** columns for each NDI (double-click the value you want to change).
5. Click **OK** in both open dialogs.
6. Save the initial configuration for the new engine node to create a one-time password. See Saving an Initial Configuration for Security Engines (page 515).
7. Make initial contact between the new engine and the Management Server. See Reconfiguring Basic Engine Settings (page 233).
8. Refresh the policy of the Master Engine to transfer the changes to the engines.

Editing Master Engine Properties

If you are creating a new Master Engine, proceed as explained in Creating New Master Engine Elements (page 405).

Table 27.1 Tasks to Do in Master Engine Properties Dialog

Tab	Tasks
Cluster	For defining the basic properties, see Creating New Master Engine Elements (page 405). For defining Locations and Contact addresses, see Defining Engine Location (page 67). For selecting an SNMP Agent, see Activating the SNMP Agent on Engines (page 584). For selecting Categories, see Using Categories (page 76). For defining a Tools Profile, see Adding Custom Commands to Element Menus (page 55). To add a node to the cluster, see Adding a Node to a Master Engine (page 408).
Interfaces	For adding or modifying interfaces, see Master Engine Interface Configuration (page 486). For defining interface options, see Setting Interface Options for Master Engines (page 497). For adding or modifying ARP entries, see Configuring Manual ARP Settings (page 510). For managing Virtual Resources, see Editing Virtual Resources (page 411).
Tester	See Getting Started with the Engine Tester (page 520).
Permissions	See Getting Started with Engine Permissions (page 530). For more information on Administrator privilege levels, see Getting Started with Administrator Accounts (page 248).
Aliases	See Getting Started with Alias Translations (page 534).
Advanced	See Adjusting Master Engine System Parameters (page 568), Adjusting Master Engine Traffic Handling Parameters (page 569), and Adjusting Master Engine Clustering Options (page 571).

Editing Virtual Security Engine Properties

If you are creating a new Virtual Security Engine, proceed as explained in Creating New Virtual Security Engines (page 407).

Table 27.2 Tasks to Do in Virtual Firewall Properties Dialog

Tab	Tasks
Single Node	For defining the basic properties, see Creating New Virtual Security Engines (page 407). For defining Locations and Contact addresses, see Defining Engine Location (page 67). For selecting Categories, see Using Categories (page 76).
Interfaces	For adding or modifying interfaces, see Virtual Security Engine Interface Configuration (page 498). For defining interface options, see Setting Interface Options for Virtual Security Engines (page 509). For adding or modifying ARP entries, see Configuring Manual ARP Settings (page 510). For configuring multicast routing, see Routing Multicast Traffic (page 593).
Permissions	See Getting Started with Engine Permissions (page 530). For more information on Administrator privilege levels, see Getting Started with Administrator Accounts (page 248).
Aliases	See Getting Started with Alias Translations (page 534).
Add-Ons	For configuring TLS inspection, see Activating TLS inspection on the Engine (page 803). To select the User Agent for Access Control by User, see Enabling Access Control by User (page 838). For configuring browser-based user authentication, see Enabling Browser-Based User Authentication (page 874).
Advanced	See Adjusting Virtual Security Engine System Parameters (page 575) and Adjusting Virtual Security Engine Traffic Handling Parameters (page 575).

Editing Virtual Resources

Editing Virtual Resources allows you to change which Domain is associated with each Virtual Resource, and to change how interfaces in the Master Engine are mapped to Virtual Security Engine Interface IDs.

▼ To edit Virtual Resource elements

1. Right-click the Master Engine and select **Properties**. The Master Engine Properties dialog opens.
2. Switch to the **Interfaces** tab and Click **Virtual Resources**. The Virtual Resources dialog opens.
3. Select a Virtual Resource and click **Edit**. The Virtual Resource Properties dialog opens.

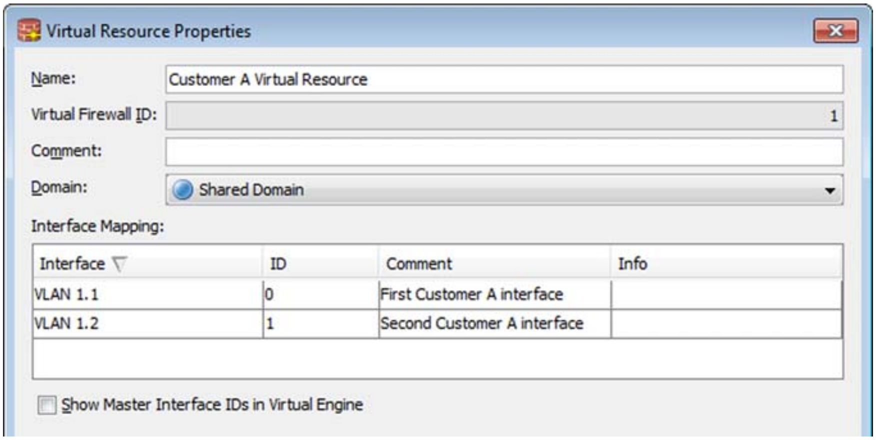

Name: Customer A Virtual Resource

Virtual Firewall ID: 1

Comment:

Domain: Shared Domain

Interface Mapping:

Interface	ID	Comment	Info
VLAN 1.1	0	First Customer A interface	
VLAN 1.2	1	Second Customer A interface	

□ Show Master Interface IDs in Virtual Engine

4. (*Optional*) Select a **Domain** to change the Domain to which the Virtual Resource belongs.

Note – If the VIrtual Resource is already associated with a Virtual Security Engine, you cannot change the Domain here. Instead, you must move the Virtual Security Engine to another Domain using the Move to Domain tool in the Virtual Security Engine's right-click menu. This automatically moves the Virtual Resource to the same Domain.

5. (*Optional*) Select the Interface **ID** in the Virtual Security Engine to which the Interface in the Master Engine is mapped.
6. (*Optional*) Select **Show Master Interface IDs in Virtual Engine** to show the Physical Interface IDs of the Master Engine in the Interface properties of the Virtual Security Engine.
7. Click **OK** to close the Virtual Resource Properties dialog.
8. When you are finished editing Virtual Resources, click **OK** to close the Virtual Resources dialog.

Converting Existing Firewalls to Master Engines and Virtual Security Engines

There is a conversion tool that allows you to change an existing Single Firewall or Firewall Cluster into a Master Engine and Virtual Firewalls.

Limitations

You must upgrade the Single Firewall or Firewall Cluster to version 5.5 or higher before starting the conversion.

Single Firewalls and Firewall Clusters with the following configurations cannot be converted to Master Engines and Virtual Firewalls:

- Firewall Clusters that use a CVI as the Control Interface or the Heartbeat Interface
- Firewall Clusters that have NDI addresses on a Tunnel Interface
- Single Firewalls that have ADSL Interfaces
- Single Firewalls that have Wireless Interfaces
- Single Firewalls that have a dynamic IP address
- Single Firewalls or Firewall Clusters that have a dynamic Contact Address
- Single Firewalls or Firewall Clusters that have DHCP settings on the Control Interface

▼ To convert existing Firewall Clusters to Master Engines and Virtual Security Engines

1. Select **Configuration→Configuration→Security Engine**. The Security Engine Configuration view opens.
2. Right-click a Single Firewall or Firewall Cluster and select **Configuration→Convert to Master and Virtual**. The Convert Engine to Master Engine and Virtual Security Engines wizard starts.

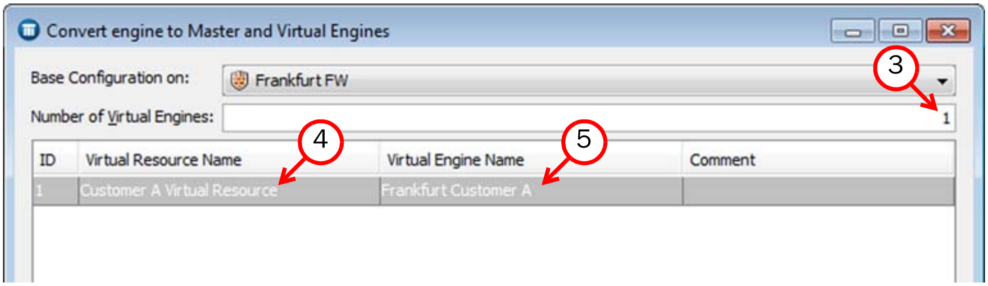

Base Configuration on: Frankfurt FW

Number of Virtual Engines: 1

ID	Virtual Resource Name	Virtual Engine Name	Comment
1	Customer A Virtual Resource	Frankfurt Customer A	

3. Enter the **Number of Virtual Engines** to create. The specified number of Virtual Resources are added to the table, and a Virtual Engine is associated with each Virtual Resource.
4. (*Recommended*) Double-click the **Virtual Resource Name** field edit the automatically generated Virtual Resource Name for each Virtual Security Engine.
5. (*Recommended*) Double-click the **Virtual Security Engine Name** field and edit the automatically generated Virtual Security Engine Name for each Virtual Security Engine.
6. Click **Next**. The Define Basic Master Engine Information page opens.

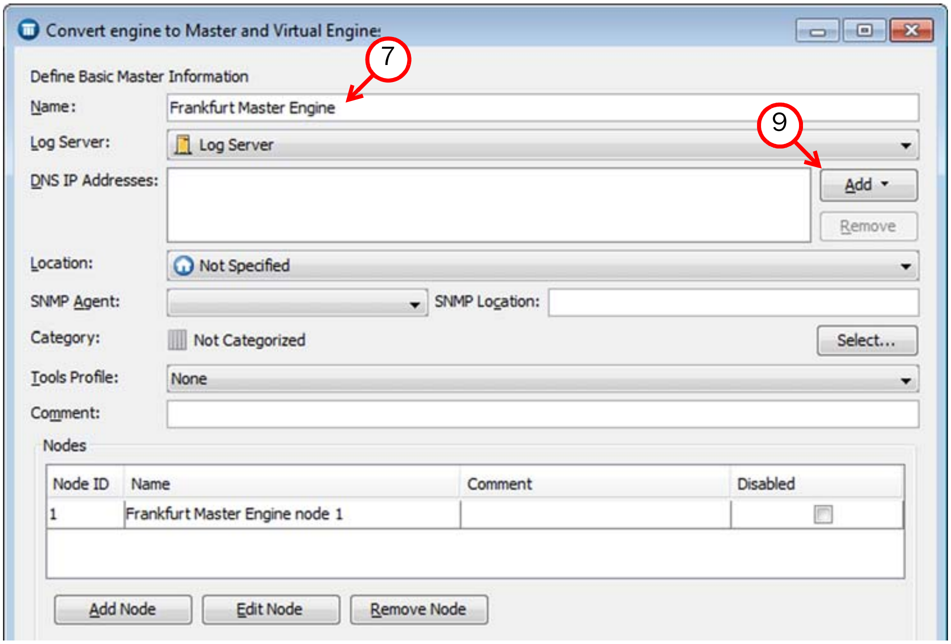

Convert engine to Master and Virtual Engine

Define Basic Master Information

Name: Frankfurt Master Engine

Log Server: Log Server

DNS IP Addresses:

Location: Not Specified

SNMP Agent:

SNMP Location:

Category: Not Categorized

Tools Profile: None

Comment:

Nodes

Node ID	Name	Comment	Disabled
1	Frankfurt Master Engine node 1		□

7. Enter a **Name** for the Master Engine. The name is also used to automatically generate the names of the nodes.

8. Select the **Log Server** to which the Master Engine sends its event data.

Note – Name Prefix and Log Server are the only mandatory properties you must define at this stage.

9. (*Optional*) Define one or more **DNS IP Addresses**. These are the IP addresses of the DNS server(s) that the Master Engine uses to resolve domain names. There are two ways to define IP addresses.
   - To enter a single IP address manually, click **Add** and select **IP Address**. Enter the IP address.
   - To define an IP address using a network element, click **Add** and select **Network Element**. Select an existing element, or click the New icon and define a new element. See Defining Host Elements (page 727) and Defining an External DNS Server (page 640) for more information.

10. Select the **Location** for this Master Engine if there is a NAT device between this Master Engine and other system components. See Defining Locations (page 66) for more information.

11. Define other settings according to your environment:
    - If you want to enable the Master Engine to send SNMP traps, select an **SNMP Agent** as explained in Activating the SNMP Agent on Engines (page 584).
    - If you want to include the Master Engine in pre-defined categories, select the appropriate **Categories** as explained in Selecting Categories for Elements (page 77).
    - If you want to add custom commands to the Master Engine’s right-click menu, add a **Tools Profile** as explained in Attaching a Tools Profile to an Element (page 56).
    - Add, edit, or remove nodes as needed as described in Adding a Node to a Master Engine (page 408).

12. Click **Next**. The Define Interfaces for the Master Engine page opens.

Defining Interfaces for Master Engines

All Master Engine interfaces must either have Interface Options defined or be associated with a Virtual Resource.

▼ To define interfaces for the Master Engine

1. Select a Physical Interfaces and click **Edit**.

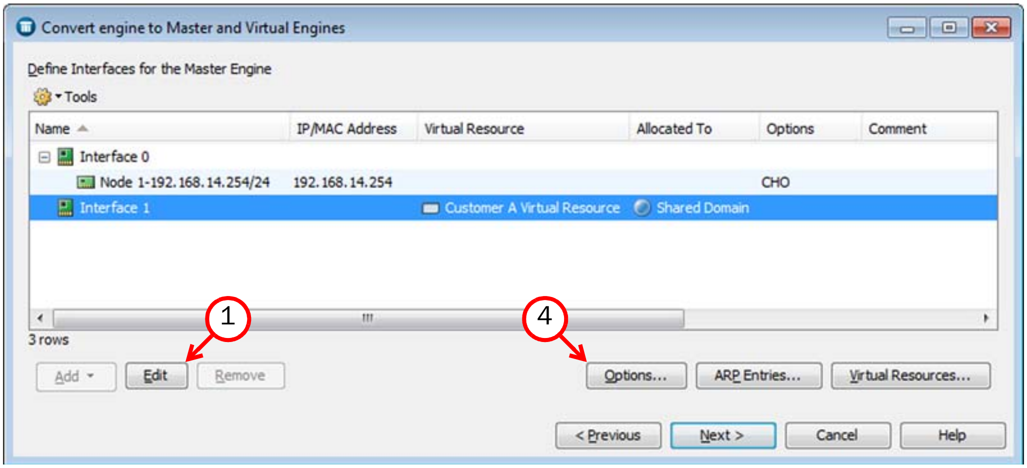

Name	IP/MAC Address	Virtual Resource	Allocated To	Options	Comment
Interface 0					
Node 1-192.168.14.254/24	192.168.14.254			CHO	
Interface 1		Customer A Virtual Resource	Shared Domain		

2. Edit the interface as explained in Defining Physical Interfaces for Master Engines (page 486).
   - If you want to use the interface to host a Virtual Security Engine, select the **Virtual Resource** to associate with the interface.
   - If you want to use the interface for Master Engine communications, do not select a Virtual Resource.
   - For settings on the **Advanced** tab, see Adjusting Master Engine System Parameters (page 568), Adjusting Master Engine Traffic Handling Parameters (page 569), and Adjusting Master Engine Clustering Options (page 571).
3. Configure the IP address settings as explained in Configuring Master Engine IP Addresses (page 493).
4. *(Optional)* Click **Options** to select which IP addresses are used in particular roles in system communications as explained in Setting Interface Options for Master Engines (page 497).
5. *(Optional)* Add **ARP entries** as explained in Configuring Manual ARP Settings (page 510).
6. Click **Next**.

What’s Next?

- If the Single Firewall or Firewall Cluster on which the configuration is associated with an Internal Security Gateway element, the Review Distribution of Internal Gateways to Virtual Security Engines page opens. Continue by Distributing Internal Security Gateways to Virtual Security Engines.
- If Tunnel Interfaces are defined for the Single Firewall or Firewall Cluster on which the configuration is based, the Distribute Tunnel Interfaces to Virtual Security Engines page opens. Continue by Distributing Tunnel Interfaces to Virtual Security Engines.
- Otherwise, the Define Routing for the Master Engine page opens. Continue by Defining Routing for the Master Engine (page 416).

Distributing Internal Security Gateways to Virtual Security Engines

The end-points of the same Internal Security Gateway must belong to the same Virtual Security Engine.

To distribute Internal Security Gateways to Virtual Security Engines

1. Review the distribution of Internal Security Gateways to Virtual Security Engines and modify it if necessary.
   - Double-click the **End-Point 1 IP Address** cell or the **End-Point 1 Phase 1 ID** cell to open the End-Point Properties dialog.
   - See Defining End-Points for Internal Security Gateways (page 924) for detailed instructions.
2. Click **Next**. The Define Routing for the Master Engine page opens.

What’s Next?

- If Tunnel Interfaces are defined for the Single Firewall or Firewall Cluster on which the configuration is based, the Distribute Tunnel Interfaces to Virtual Security Engines page opens. Continue by Distributing Tunnel Interfaces to Virtual Security Engines.
- Otherwise, the Define Routing for the Master Engine page opens. Continue by Defining Routing for the Master Engine (page 416).

Distributing Tunnel Interfaces to Virtual Security Engines

To distribute Tunnel Interfaces to Virtual Security Engines

1. Review the distribution of Tunnel Interfaces to Virtual Security Engines and modify it if necessary.
   - Double-click the Tunnel Interface **Name** to open the Tunnel Interface Properties dialog.
   - See Defining Tunnel Interfaces for Virtual Security Engines (page 502) for detailed instructions.
2. Click **Next**.

Defining Routing for the Master Engine

▼ To define routing for the Master Engine

1. Review and edit the routing as explained in Adding Routes for Master Engines (page 598).
2. Click **Next**. The Select Additional Configuration Options page opens.

What's Next?

- ▶ If you want to define additional options, proceed to Selecting Additional Configuration Options for Master Engines.
- ▶ Otherwise, click **Next** on the Select Additional Configuration Options page. Proceed to Defining Basic Information for Virtual Security Engines.

Selecting Additional Configuration Options for Master Engines

▼ To select additional configuration options for Master Engines

1. Select **Define Additional Master Engine Properties** and click **Next**.
2. (*Optional*) Define tester settings as explained in Getting Started with the Engine Tester (page 520).
3. Click **Next**.
4. (*Optional*) Define permissions as explained in Getting Started with Engine Permissions (page 530).
5. Click **Next**.
6. (*Optional*) Define advanced settings as explained in the following sections:
   - Adjusting Master Engine System Parameters (page 568)
   - Adjusting Master Engine Traffic Handling Parameters (page 569)
   - Adjusting Master Engine Clustering Options (page 571).
7. Click **Next**.

Defining Basic Information for Virtual Security Engines

Most of the settings are inherited from the Master Engine. Some settings can be adjusted for the Virtual Security Engines.

▼ To define basic information for Virtual Security Engines

1. (*Optional*) Edit the **Name** for the Virtual Security Engine.
2. (*Optional*) Define one or more **DNS IP Addresses**. These are the IP addresses of the DNS server(s) that the Master Engine uses to resolve domain names. There are two ways to define IP addresses.
   - To enter a single IP address manually click **Add** and select **IP Address**. Enter the IP address.
   - To define an IP address using a network element, click **Add** and select **Network Element**. Select an existing element, or click the New icon and define a new element. See Defining Host Elements (page 727) and Defining an External DNS Server (page 640) for more information.
3. (*Optional*) Select the appropriate **Categories** as explained in Selecting Categories for Elements (page 77).

4. Click **Next**. The Review Interfaces for Virtual Security Engines page opens.

Configuring Interfaces and Routing for Virtual Security Engines

The interface configuration is generated automatically based on the Master Engine interface configuration. Some settings can be adjusted for the Virtual Security Engines.

▼ To configure interfaces for Virtual Security Engines

1. Click **Options** to select which IP addresses are used in particular roles in system communications as explained in Setting Interface Options for Master Engines (page 497).
2. *(Optional)* Add **ARP entries** as explained in Configuring Manual ARP Settings (page 510).
3. Click **Next**. The Review and Edit Routing for Virtual Security Engines page opens.
4. Review and edit the routing as explained in Adding Routes for Virtual Firewalls (page 599).
5. Click **Next**. The Select Additional Configuration Options page opens.

What's Next?

- ▶ If you want to define additional options, proceed to Selecting Additional Configuration Options for Virtual Security Engines (page 417).
- ▶ Otherwise, click **Next** on the Select Additional Coifiguration Options page. The Summary page opens.

Selecting Additional Configuration Options for Virtual Security Engines

▼ To select additional configuration options for Virtual Security Engines

1. (*Optional*) Define permissions as explained in Getting Started with Engine Permissions (page 530).
2. Click **Next**.
3. (*Optional*) Define add-ons as explained in the following sections:
   - Activating TLS inspection on the Engine (page 803)
   - Enabling Access Control by User (page 838)
   - Enabling Browser-Based User Authentication (page 874).
4. Click **Next**.
5. (*Optional*) Define advanced settings as explained in the following sections:
   - Adjusting Virtual Security Engine System Parameters (page 575)
   - Adjusting Virtual Security Engine Traffic Handling Parameters (page 575).
6. Click **Next**. The Summary page opens.

Finishing the Convert Engine to Master Engine and Virtual Security Engines Wizard

▼ To finish the Convert Engine to Master Engine and Virtual Security Engines wizard

1. Click **Finish**. A new tab opens to show the progress of the conversion.
2. Click **Close** to close the tab when the conversion is complete.
3. Install or refresh the policy on the Master Engine.
4. Install a policy on the Virtual Security Engine(s).

Related Tasks

- ▶ Editing Virtual Security Engine Properties (page 410)

CHAPTER 28

NETWORK INTERFACE CONFIGURATION

The network interface configuration for all security engines is stored on the Management Server in the properties of Single Firewall, Firewall Cluster, Single IPS, IPS Cluster, Single Layer 2 Firewall, Layer 2 Firewall Cluster, Master Engine, and Virtual Firewall elements.

The following sections are included:

Getting Started with Interface Configuration

Prerequisites: Creating New Engine Elements / Modifying Existing Engine Elements

The Interface configuration is done using the Management Client. The interface configuration stored on the Management Server contains all settings related to all supported types of network interfaces except the network card driver selection, the mapping of the operating system port numbers to the numbers used in the Management Client, and the network card speed/duplex settings.

Note – If you use automatic configuration for appliances, the interface numbers (Interface IDs) that you select for the interfaces in the Management Client must match the port numbers on the physical appliances. Check the port numbers in the relevant *Appliance Installation Guide(s)*.

The configuration transferred from the Management Server overwrites the settings that can be defined through the engine command line (the details for initial contact to the Management Server to establish a trusted communications channel).

Limitations

- Clusters cannot have dynamic IP addresses.
- You cannot add both VLAN Interfaces and IP Addresses directly to a Physical Interface. If an IP Address is already configured for a Physical Interface, adding a VLAN Interface replaces the IP Address. If you plan to use VLAN Interfaces, configure the VLAN Interfaces first and then add IP Addresses to the VLAN Interfaces.
- You cannot add VLAN Interfaces on top of other VLAN Interfaces (nested VLANs).
    - You cannot create valid VLAN Interfaces in a Virtual Security Engine if the Master Engine interface that hosts the Virtual Security Engine is a VLAN Interface.
- Only IPv4 addresses are supported in the following interface configurations:
    - As system communications interfaces
    - In Manual ARP Entries
    - For DHCP relay
    - For VRRP
    - As dynamic IP addresses on the engines' interfaces
- An ADSL Interface is only supported on Single Firewall engines that run on specific pre-installed Stonesoft appliances that have an ADSL network interface card.
- A Wireless Interface is only supported on Single Firewall engines that run on specific pre-installed Stonesoft appliances that have a wireless network interface card.
- Modem interfaces, ADSL interfaces, and Tunnel interfaces do not support VLAN tagging.
- Tunnel Interfaces are only supported on Firewalls.
- Tunnel Interfaces can only have static IP addresses.

Configuration Overview

The interface configuration proceeds as follows:

1. Define the interfaces and IP addresses according to the engine role. See the engine-specific section for a detailed configuration overview of interface configuration for each engine role:
   - Firewall Interface Configuration
   - IPS Engine Interface Configuration (page 456)
   - Layer 2 Firewall Interface Configuration (page 473)
   - Master Engine Interface Configuration (page 486)
   - Virtual Security Engine Interface Configuration (page 498)
2. Configure additional related settings depending on the features you want to use.
   - Configuring Manual ARP Settings (page 510)
   - Activating the Internal DHCP Server on a Firewall Interface (page 511)

What's Next?

▶ Firewall Interface Configuration

▶ IPS Engine Interface Configuration (page 456)

▶ Layer 2 Firewall Interface Configuration (page 473)

▶ Master Engine Interface Configuration (page 486)

▶ Virtual Security Engine Interface Configuration (page 498)

Firewall Interface Configuration

Prerequisites: Creating New Engine Elements / Modifying Existing Engine Elements

The interface configuration for Single Firewalls and Firewall Clusters consists of the following main steps:

1. Add the required number of network connections:
   - Define Ethernet links as explained in Defining Physical Interfaces for Firewall Engines (page 422).
   - (*Single Firewalls only*) Define integrated ADSL modems as explained in Adding ADSL Interfaces for Single Firewalls (page 427).
   - (*Single Firewalls only*) Define integrated wireless routers explained in Adding Wireless Interfaces for Single Firewalls (page 429).
   - (*Single Firewalls only*) Define integrated 3G modems as explained in Defining Modem Interfaces for Single Firewalls (page 448).
2. (*Physical Interfaces only)* Add the required number of VLANs as explained in Adding VLAN Interfaces for Firewall Engines (page 425).
3. (*Optional*) Define Tunnel Interfaces for Route-Based VPNs as explained in Defining Tunnel Interfaces for Firewalls (page 435).
4. (*Not applicable to Modem interfaces*) Configure the IP address settings as explained in Configuring Single Firewall IP Addresses (page 439) or Configuring Firewall Cluster IP Addresses (page 444).
5. (*Optional*) Define Loopback IP addresses to allow the firewall to communicate with itself as explained in Configuring Loopback IP Addresses for Firewalls (page 450).

6. Select the interfaces that are used for system communications as explained in Setting Interface Options for Firewalls (page 453).

Related Tasks

- Configuring Manual ARP Settings (page 510)
- Routing Multicast Traffic (page 593)
- IPS Engine Interface Configuration (page 456)
- Layer 2 Firewall Interface Configuration (page 473)
- Master Engine Interface Configuration (page 486)
- Virtual Security Engine Interface Configuration (page 498)

Defining Physical Interfaces for Firewall Engines

Physical Interfaces correspond to network ports on the Firewall. By default, the numbering of the Physical Interfaces in the Management Client corresponds to the operating system interface numbering on the engine (that is, Interface ID 0 is mapped to eth0, ID 1 to eth1, etc.). However, the mapping is not fixed and you can change it through the engine command line. See the relevant *Appliance Installation Guide* for details on which Interface IDs to map with which network ports.

For settings on the **DHCP** tab, see Activating the Internal DHCP Server on a Firewall Interface (page 511) or (for DHCP relay) Routing DHCP Messages (page 590).

For settings on the **Advanced** tab, see Configuring Advanced Interface Properties for Firewalls (page 436).

▼ To define a Physical Interface

1. In the properties dialog for the Firewall, switch to the **Interfaces** tab.
2. Right-click the empty space and select **New→Physical Interface** or **New Physical Interface**. The Physical Interface Properties dialog opens.

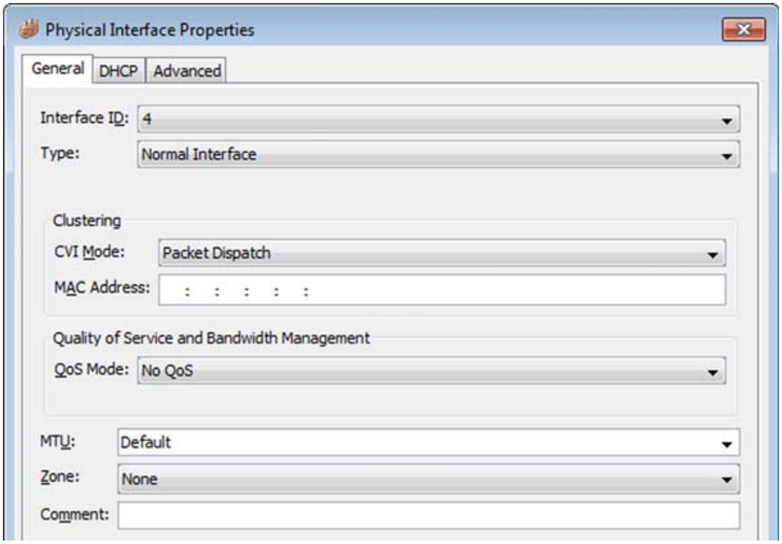
Physical Interface Properties
General
DHCP
Advanced
Clustering
Quality of Service and Bandwidth Management

Interface ID: 4
Type: Normal Interface
CVI Mode: Packet Dispatch
MAC Address: : : : : :
QoS Mode: No QoS
MTU: Default
Zone: None
Comment:

3. Define the Physical Interface properties as explained in the table below.

Table 28.1 Physical Interface Properties - General Tab

Options		Explanation
Interface ID		The Interface ID automatically maps to a physical interface of the same number during the initial configuration of the engine, but the mapping can be changed as necessary through the engine's command line interface.
Type	Normal Interface	Corresponds to a single network interface on the Firewall engine.
	Aggregated Link in High-Availability Mode	Represents two interfaces on the Firewall engine. Only the first interface in the aggregated link is actively used. The second interface becomes active only if the first interface fails. If you configure an Aggregated Link in High-Availability mode, connect the first interface to one switch and the second interface to another switch.
	Aggregated Link in Load-Balancing Mode	Represents two interfaces on the Firewall engine. Both interfaces in the aggregated link are actively used and connections are automatically balanced between the two interfaces. Link aggregation in the load-balancing mode is implemented based on the IEEE 802.3ad Link Aggregation standard. If you configure an Aggregated Link in Load-Balancing Mode, connect both interfaces to a single switch. Make sure that the switch supports the Link Aggregation Control Protocol (LACP) and that LACP is configured on the switch.
Second Interface ID (*Only if interface type is Aggregated Link*)		The second interface in the aggregated link.
CVI Mode (*Optional,* *Firewall Clusters only*)		Unless you have a specific reason to use some other setting, use the default Packet Dispatch setting. The Packet Dispatch mode is the best choice in most environments. See the *Firewall/VPN Reference Guide* for more information.
MAC Address (*Firewall Clusters only*)		The MAC address for the Cluster Virtual IP Address. Do not use the MAC address of any actual network card on any of the nodes. *Packet Dispatch and Unicast MAC modes*: enter a unicast MAC address (even number as the first octet). *Multicast MAC mode*: enter a multicast MAC address (odd number as the first octet). *Multicast MAC with IGMP:* enter a Multicast Address, that is, an IP address from the range 224.0.0.0-239.255.255.255. The address is used for automatically calculating a MAC address.

Table 28.1 Physical Interface Properties - General Tab (Continued)

Options	Explanation
QoS Mode (*Optional*)	Defines how QoS is applied to the link on this interface. If **Full QoS** or **DSCP Handling and Throttling** is selected as the QoS Mode, a QoS Policy must also be selected. If **Full QoS** is selected as the QoS Mode, the Throughput must also be defined.
QoS Policy (*Full QoS or DSCP Handling and Throttling modes only*)	The QoS Policy for the link on this interface. For more information, see Getting Started with QoS (page 782).
Throughput (*Full QoS mode only)*	Enter the throughput for the link on this interface as kilobits per second (for example, 2048). The same throughput is automatically applied to any VLANs created under this Physical Interface. See Adding VLAN Interfaces for Firewall Engines (page 425). The throughput is for uplink speed (outgoing traffic) and typically must correspond to the speed of an Internet link (such as an ADSL line), or the combined speeds of several such links when they are connected to a single physical interface). **Caution!** Make sure you set the Interface speed correctly. When the bandwidth is set, the Firewall always scales the total amount of traffic on this interface to the bandwidth you defined. This happens even if there are no bandwidth limits or guarantees defined for any traffic.
MTU (*Optional*)	The MTU (maximum transmission unit) size on the connected link. Either enter a value between 400-65535 or select a common MTU value from the list. The same MTU is automatically applied to any VLANs created under this Physical Interface. See Adding VLAN Interfaces for Firewall Engines (page 425). The default value (also the maximum standard MTU in Ethernet) is 1500. Do not set a value larger than the standard MTU unless you know that all the devices along the communications path support it.
Zone (*Optional*)	Select the network zone to which the Physical interface belongs from the list or select **Other** to select another Zone. If the Zone is not listed, create a new Zone element through the New icon at the top of the dialog. See Defining Zone Elements (page 730) for more information.

4. Click **OK**. The Physical Interface is added to the interface list.
5. Repeat from Step 2 to add more Physical Interfaces.

What’s Next?

- If you want to use VLANs, add the VLANs before adding IP addresses. Proceed to Adding VLAN Interfaces for Firewall Engines.
- Otherwise, add IP addresses directly to the Physical Interfaces as instructed in Configuring Single Firewall IP Addresses (page 439) or Configuring Firewall Cluster IP Addresses (page 444).

Related Tasks

- Activating the Internal DHCP Server on a Firewall Interface (page 511)
- For DHCP relay, see Routing DHCP Messages (page 590).
- Configuring Advanced Interface Properties for Firewalls (page 436)

Adding VLAN Interfaces for Firewall Engines

VLANs divide a single physical network link into several virtual links. VLANs can be defined for both Single Firewalls and Firewall Clusters. The maximum number of VLANs for a single Physical Interface is 4094. The VLANs must also be defined in the configuration of the switch/router to which the interface is connected.

Note – If an IP Address is already configured for a Firewall Physical Interface, adding a VLAN Interface replaces the IP Address. If you plan to use VLAN Interfaces, configure the VLAN Interfaces first and then add IP Addresses to the VLAN Interfaces.

For settings on the **DHCP** tab, see Activating the Internal DHCP Server on a Firewall Interface (page 511) or (for DHCP relay) Routing DHCP Messages (page 590).

For settings on the **Advanced** tab, see Configuring Advanced Interface Properties for Firewalls (page 436).

▼ To add a VLAN Interface for a Firewall

1. In the properties dialog for the Firewall, switch to the **Interfaces** tab.
2. Right-click a Physical Interface and select **New→VLAN Interface**. The VLAN Interface Properties dialog opens.

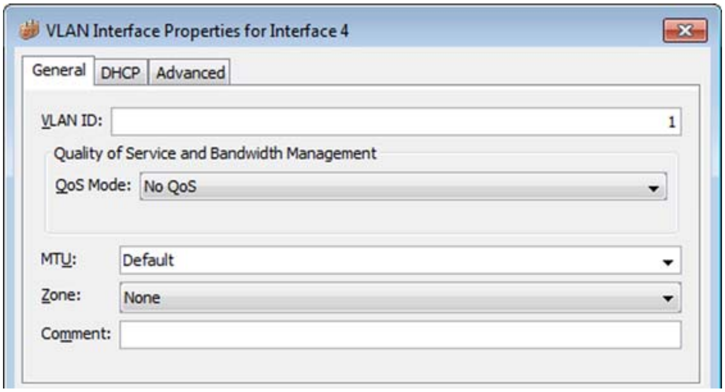

VLAN Interface Properties for Interface 4

VLAN ID: 1

Quality of Service and Bandwidth Management

QoS Mode: No QoS

MTU: Default

Zone: None

Comment:

3. Define the VLAN interface properties as explained in the table below.

Table 28.2 VLAN Interface Properties - General Tab

Option	Explanation
VLAN ID	Enter the VLAN ID (1-4094). The VLAN IDs you add must be the same as the VLAN IDs that are used in the switch at the other end of the VLAN trunk.
QoS Mode (*Optional*)	Defines how QoS is applied to the link on this interface. If **Full QoS** or **DSCP Handling and Throttling** is selected as the QoS Mode, a QoS Policy must also be selected. If **Full QoS** is selected as the QoS Mode, the Throughput must also be defined.
QoS Policy (*Full QoS or DSCP Handling and Throttling modes only*)	The QoS Policy for the link on the interface. For more information, see Getting Started with QoS (page 782).
Throughput (*Full QoS mode only)*	Enter the throughput for the link on this interface as kilobits per second (for example, 2048). If throughput is defined for the Physical Interface to which the VLAN Interface belongs, the throughput value is automatically inherited from the Physical Interface properties. **Caution!** The throughput for each VLAN Interface must not be higher than the throughput for the Physical Interface to which the VLAN Interface belongs. The throughput is for uplink speed (outgoing traffic) and typically must correspond to the speed of an Internet link (such as an ADSL line), or the combined speeds of several such links when they are connected to a single Physical Interface. **Caution!** Make sure you set the interface speed correctly. When the bandwidth is set, the Firewall always scales the total amount of traffic on this interface to the bandwidth you defined. This happens even if there are no bandwidth limits or guarantees defined for any traffic.
MTU (*Optional*)	The MTU (maximum transmission unit) size on the connected link. Either enter a value between 400-65535 or select a common MTU value from the list. **Caution!** The MTU for each VLAN Interface must not be higher than the MTU for the Physical Interface to which the VLAN Interface belongs. The default value (also the maximum standard MTU in Ethernet) is 1500. Do not set a value larger than the standard MTU unless you know that all the devices along the communications path support it.
Zone (*Optional*)	Select the network zone to which the VLAN interface belongs from the list or select **Other** to select another Zone. If the Zone is not listed, create a new Zone element through the New icon at the top of the dialog. See Defining Zone Elements (page 730) for more information.

4. Click **OK**. The specified VLAN ID is added to the Physical Interface.

5. Repeat from Step 2 to add further VLANs on the same or other Physical Interfaces.

Each VLAN interface is identified as `Interface-ID.VLAN-ID`, for example `2.100` for Interface ID 2 and VLAN ID 100.

What's Next?

- Add IP addresses to the VLAN Interfaces as instructed in Configuring Single Firewall IP Addresses (page 439) or Configuring Firewall Cluster IP Addresses (page 444).

Related Tasks

- Activating the Internal DHCP Server on a Firewall Interface (page 511)
- For DHCP relay, see Routing DHCP Messages (page 590)
- Configuring Advanced Interface Properties for Firewalls (page 436)

Adding ADSL Interfaces for Single Firewalls

You can configure one ADSL Interface on each Single Firewall. The supported ADSL standards are ANSI T1.413 issue 2n, G.dmt, G.lite, ADSL2 DMT, ADSL2 G.lite, Annex A, and Annex B. ADSL is only supported on specific Stonesoft appliances that have an integrated ADSL network interface card. ADSL Interfaces are not supported on Firewall Clusters.

For information on which Interface ID to map with the ADSL port on the appliance, see the relevant *Appliance Installation Guide*.

For settings on the **Advanced** tab, see Configuring Advanced Interface Properties for Firewalls (page 436).

▼ To define an ADSL Interface

1. In the properties dialog for the Single Firewall, switch to the **Interfaces** tab.
2. Right-click the empty space and select **New→ADSL Interface**. The ADSL Interface Properties dialog opens.

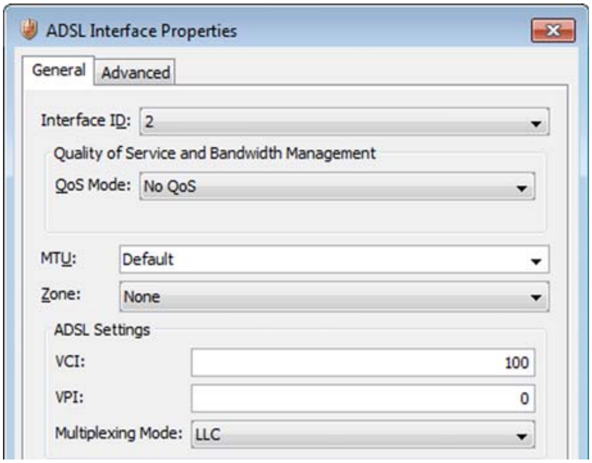

Interface ID: 2

Quality of Service and Bandwidth Management

QoS Mode: No QoS

MTU: Default

Zone: None

ADSL Settings

VCI: 100

VPI: 0

Multiplexing Mode: LLC

3. Define the ADSL Interface properties as explained in the table below.

Table 28.3 ADSL Interface Properties - General Tab

Option	Explanation
Interface ID	Select the number of the ADSL port on the appliance as the Interface ID. The Interface ID is automatically mapped to the ADSL port on the engine's ADSL card during the initial configuration of the engine.
QoS Mode (*Optional*)	Defines how QoS is applied to the link on this interface. If **Full QoS** or **DSCP Handling and Throttling** is selected as the QoS Mode, a QoS Policy must also be selected. If **Full QoS** is selected as the QoS Mode, the Throughput must also be defined.
QoS Policy (*Full QoS or DSCP Handling and Throttling modes only*)	The QoS Policy for the link on the interface. For more information, see Getting Started with QoS (page 782).
Throughput (*Full QoS mode only)*	Enter the throughput for the link on this interface as kilobits per second (for example, 2048). If throughput is defined for the Physical Interface to which the VLAN Interface belongs, the throughput value is automatically inherited from the Physical Interface properties. The throughput for each VLAN Interface must not be higher than the throughput for the Physical Interface to which the VLAN Interface belongs. The throughput is for uplink speed (outgoing traffic) and typically must correspond to the speed of an Internet link (such as an ADSL line), or the combined speeds of several such links when they are connected to a single Physical Interface. **Caution!** Make sure you set the interface speed correctly. When the bandwidth is set, the Firewall always scales the total amount of traffic on this interface to the bandwidth you defined. This happens even if there are no bandwidth limits or guarantees defined for any traffic.
MTU (*Optional*)	The MTU (maximum transmission unit) size on the connected link. Either enter a value between 400-65535 or select a common MTU value from the list. The default value (also the maximum standard MTU in Ethernet) is 1500. Do not set a value larger than the standard MTU unless you know that all the devices along the communications path support it.
Zone (*Optional*)	Select the network zone to which the ADSL interface belongs from the list or select **Other** to select another Zone. If the Zone is not listed, create a new Zone element through the New icon at the top of the dialog. See Defining Zone Elements (page 730) for more information.
VCI	Enter the **VCI** (Virtual Channel Identifier) value according to the configuration information provided by your ISP.
VPI	Enter the **VPI** (Virtual Path Identifier) value according to the configuration information provided by your ISP.
Multiplexing Mode	Select **LLC** (Logical Link Control) or **VC** (Virtual Circuit) according to the configuration information provided by your ISP.

4. Click **OK**. The ADSL Interface is added to the interface list.

What's Next?

- ▶ Configuring Single Firewall IP Addresses (page 439)

Related Tasks

- ▶ Configuring Advanced Interface Properties for Firewalls (page 436)
- ▶ If necessary, see Saving Initial Configuration Details (page 517) for selecting the Internet Service Provider (ISP).

Adding Wireless Interfaces for Single Firewalls

You can configure one Wireless Interface on a Single Firewall. Wireless Interfaces are only supported on specific Stonesoft appliances that have an integrated wireless network interface card. Wireless Interfaces are not supported on Firewall Clusters.

For information on which Interface ID to map with the wireless network interface on the appliance, see the relevant Appliance Installation Guide.

For settings on the **Advanced** tab, see Configuring Advanced Interface Properties for Firewalls (page 436).

▼ To define a Wireless Interface

1. In the properties dialog for the Single Firewall, switch to the **Interfaces** tab.
2. Right-click the empty space and select **New→Wireless Interface**. The Wireless Interface Properties dialog opens.

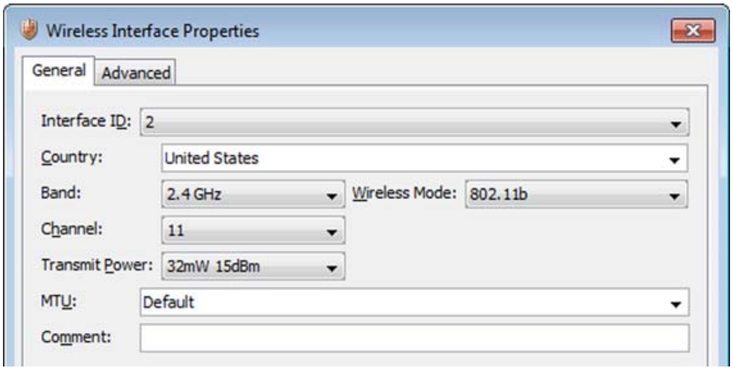

Interface ID: 2
Country: United States
Band: 2.4 GHz
Wireless Mode: 802.11b
Channel: 11
Transmit Power: 32mW 15dBm
MTU: Default
Comment:

3. Define the Wireless Interface properties as explained in the table below.

Table 28.4 Wireless Interface Properties - General Tab

Option	Explanation
Interface ID	Select the number of the wireless port on the appliance as the Interface ID. The Interface ID automatically maps to the wireless port on the engine's wireless network interface card during the initial configuration of the engine.
Country	The country where the Firewall is used as a wireless access point.

Table 28.4 Wireless Interface Properties - General Tab (Continued)

Option	Explanation
Band	The band for the wireless interface access point (2.4 GHz or 5 GHz). The Wireless Mode and Channel options depend on the band that you select.
Wireless Mode	The mode in which the wireless traffic is transmitted. Select the mode according to the capabilities of the connecting clients. The available modes depend on the selected Band. If you have selected 2.4 GHz as the Band, the options are: **802 11b**: 11 Mbit wireless-b only mode. **802 11bg**: 54 Mbit wireless-b and g modes. **802 11g**: 54 Mbit wireless-g only mode. **802 11n**: 270 Mbit wireless-n only mode. **802 11bgn**: 270 Mbit wireless-b, g, and n modes. If you have selected 5 GHz as the Band, the options are: **802 11a**: 54 Mbit wireless-a only mode. **802 11an**: 270 Mbit wireless-a and n modes. **802 11n**: 270 Mbit wireless-n only mode. **Note!** Some wireless clients do not support the 802.11n wireless mode with the WEP security mode. See Configuring Security Settings for SSID Interfaces (page 433).
Channel	The frequency for transmitting the wireless traffic. The available channels depend on the selected Band. If there are other wireless access points nearby, use channels that are as far apart as possible to avoid interference.
Transmit Power (*Optional*)	Select the maximum power of the signal for transmitting the wireless traffic. The power options are shown as milliwatts (mW) and as the power ratio in decibels of the measured power referenced to one milliwatt (dBm). The values available depend on the regulatory limits for the selected Country and the Channel for the Wireless Interface. If you are not sure what value to use, leave the default value selected.
MTU (*Optional*)	The MTU (maximum transmission unit) size on the connected link. Either enter a value between 400-65535 or select a common MTU value from the list. The default value is 1500 (the maximum standard MTU in Ethernet). Do not set a value larger than the standard MTU unless you know that all the devices along the communications path support it.

4. Click **OK**. The Wireless Interface is added to the interface list.

What's Next?

▶ Proceed to Defining SSID Interfaces for Single Firewalls (page 431).

Defining SSID Interfaces for Single Firewalls

An SSID (service set identifier) interface represents an 802.11 wireless LAN. You can define several SSID Interfaces for the Wireless Interface.

For settings on the Security tab, see Configuring Security Settings for SSID Interfaces (page 433).

For settings on the MAC Filtering tab, see Configuring MAC Filtering for SSID Interfaces (page 434).

For settings on the DHCP tab, see Activating the Internal DHCP Server on a Firewall Interface (page 511) or (for DHCP relay) Routing DHCP Messages (page 590).

▼ To define an SSID Interface

1. Right-click the Wireless Interface and select **New SSID Interface**. The SSID Interface Properties dialog opens.

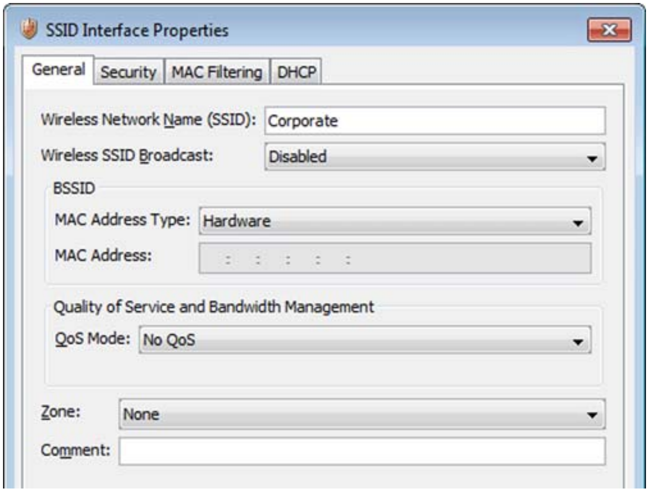

Wireless Network Name (SSID): Corporate

Wireless SSID Broadcast: Disabled

BSSID

MAC Address Type: Hardware

MAC Address: : : : : :

Quality of Service and Bandwidth Management

QoS Mode: No QoS

Zone: None

Comment:

2. Define the SSID Interface properties as explained in the table below:

Table 28.5 SSID Interface Properties - General Tab

Option		Explanation
Wireless Network Name (SSID)		The name that identifies the network to the end-users. **Note!** The name can be read by anyone within range if Wireless SSID Broadcast is enabled.
Wireless SSID Broadcast	Disabled	The end-users must enter the Wireless Network Name (SSID) to connect. Even if you disable SSID broadcast, anyone within range can discover your wireless network with detection tools widely available on the Internet.
	Enabled	End-users (and anyone else in range) can see the Wireless Network Name (SSID) in their list of available networks without further action.

Table 28.5 SSID Interface Properties - General Tab (Continued)

Option		Explanation
MAC Address Type	Hardware	The MAC address of the Firewall appliance's wireless card. This is the only MAC Address Type available when you define the first SSID interface. The first SSID interface is automatically assigned the MAC address of the wireless card.
	Custom	A custom MAC address. Enter the **MAC Address** in the field below.
QoS Mode (*Optional*)		Defines how QoS is applied to the link on this interface. If **Full QoS** or **DSCP Handling and Throttling** is selected as the QoS Mode, a QoS Policy must also be selected. If **Full QoS** is selected as the QoS Mode, the Throughput must also be defined.
QoS Policy (*Full QoS or DSCP Handling and Throttling modes only*)		The QoS Policy for the link on the interface. For more information, see Getting Started with QoS (page 782).
Throughput (*Full QoS mode only*)		Enter the throughput for the link on this interface as kilobits per second (for example, 2048). The throughput is for uplink speed (outgoing traffic) and typically must correspond to the speed of an Internet link or the combined speeds of several such links when they are connected to a single Wireless Interface. **Caution!** Make sure you set the interface speed correctly. When the bandwidth is set, the Firewall always scales the total amount of traffic on this interface to the bandwidth you defined. This happens even if there are no bandwidth limits or guarantees defined for any traffic.
Zone (*Optional*)		Select the network zone to which the wireless interface belongs from the list or select **Other** to select another Zone. If the Zone is not listed, create a new Zone element through the New icon at the top of the dialog. See Defining Zone Elements (page 730) for more information.

What's Next?

- To configure security settings for the wireless connections, proceed to Configuring Security Settings for SSID Interfaces (page 433).
- To filter wireless connections based on the clients' MAC addresses, proceed to Configuring MAC Filtering for SSID Interfaces (page 434).
- To activate or disable the internal DHCP server on the SSID interface, proceed to Activating the Internal DHCP Server on a Firewall Interface (page 511).
- To add an IPv4 address to the interface, proceed to Adding IPv4 Addresses for a Single Firewall (page 440).
- Otherwise, click **OK** close the properties of the SSID Interface and of the Firewall, and refresh the Firewall's policy to transfer the configuration changes.

Configuring Security Settings for SSID Interfaces

▼ To define security settings for SSID Interfaces

1. In the properties of the SSID Interface, switch to the **Security** tab.

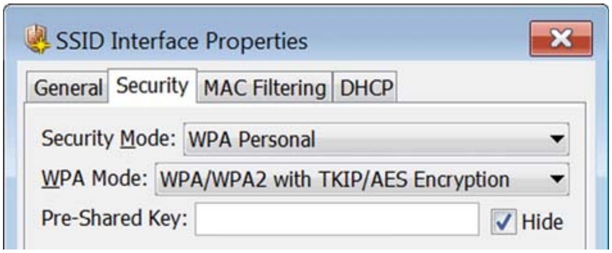

SSID Interface Properties

Security Mode: WPA Personal

WPA Mode: WPA/WPA2 with TKIP/AES Encryption

Pre-Shared Key:

☑ Hide

2. Select the **Security Mode** settings as explained in the table below.

Option	Explanation
Disabled	Wireless traffic is not encrypted. Anyone within range can freely use and intercept traffic from this wireless network. We do not recommend using this setting.
WEP Open System	After the clients have connected to the Firewall, the wireless traffic is encrypted with a 40-bit, 104-bit, or 128-bit WEP (Wired Equivalent Privacy/Wireless Encryption Protocol) key. We do not recommend this security mode. If you must use WEP for compatibility reasons, use WEP Shared Key. **Note!** Some wireless clients do not support the 802.11n wireless mode with the WEP security mode.
WEP Shared Key	The connecting clients are authenticated using WEP (Wired Equivalent Privacy/ Wireless Encryption Protocol). The wireless traffic is encrypted with a 40-bit, 104-bit, or 128-bit key. We do not recommend this security mode unless you must use WEP for compatibility reasons. **Note!** Some wireless clients do not support the 802.11n wireless mode with the WEP security mode.
WPA Personal	Wireless traffic is encrypted using the WPA or WPA2 protocol. Three encryption modes are available: either TKIP (Temporal Key Integrity Protocol) or AES Advanced Encryption Standard) or both TKIP and AES are used.
WPA Enterprise	Same as above, but RADIUS-based authentication methods provided by an external authentication server or the Authentication Server component are used to authenticate the users. This is the most secure option offered, and it is recommended if external RADIUS authentication is available.

3. Fill in the options for the selected security mode:
   - For **WEP Open System** and **WEP Shared Key**, select the **Key Length** and the **Default Key**, and enter 1 to 4 encryption keys.
   - For **WPA Personal**, select the **WPA Mode** and enter a **Pre-Shared Key** of 8 to 64 ASCII characters.
   - For **WPA Enterprise**, first select the WPA Mode and then click **Select** to choose the RADIUS Authentication Method that authenticates the users. See Integrating External Authentication Services (page 858) and Integrating Authentication Server Services

(page 861) for more information.

What's Next?

- To filter wireless connections based on the clients' MAC addresses, proceed to Configuring MAC Filtering for SSID Interfaces.
- To activate or disable internal DHCP server on the interface, proceed to Activating the Internal DHCP Server on a Firewall Interface (page 511).
- To add an IPv4 address to the interface, proceed to Adding IPv4 Addresses for a Single Firewall (page 440).
- Otherwise, close the Firewall's Properties dialog and refresh the Firewall's policy to transfer the configuration changes.

Configuring MAC Filtering for SSID Interfaces

MAC filtering allows you to restrict which clients can connect to the SSID Interfaces based on the clients' MAC address.

To configure MAC filtering

1. In the properties of the SSID Interface, switch to the **MAC Filtering** tab.

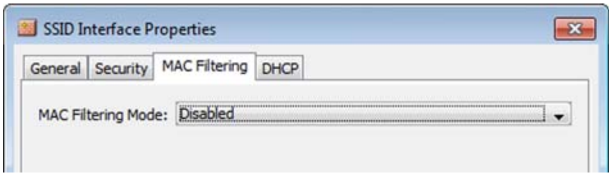

MAC Filtering Mode: Disabled

2. Select the MAC Filtering Mode and define the MAC filtering settings as explained in the table below:.

MAC Filtering Mode	Explanation
Disabled	Connecting clients are not filtered based on their MAC addresses.
Accept Unless in Denied MAC List	Clients whose MAC addresses are on the list of denied MAC addresses cannot connect. Enter each denied MAC Address and click **Add**.
Deny Unless in Allowed MAC List	Only clients whose MAC addresses are on the list of allowed MAC addresses can connect. Enter each allowed MAC Address and click **Add**.

What's Next?

- To activate or disable the internal DHCP server on the SSID interface, proceed to Activating the Internal DHCP Server on a Firewall Interface (page 511).
- To add an IPv4 address to the interface, proceed to Adding IPv4 Addresses for a Single Firewall (page 440).
- Otherwise, close the Firewall's Properties dialog and refresh the Firewall's policy to transfer the configuration changes.

Defining Tunnel Interfaces for Firewalls

Tunnel Interfaces define end-points for tunnels in the Route-Based VPN. Any traffic that is routed to a Tunnel Interface and allowed by the Firewall Access rules is sent into the tunnel. See Getting Started With IPsec VPNs (page 918) for more information about the Route-Based VPN.

You can optionally add IPv4 and/or IPv6 addresses to a Tunnel Interface. Tunnel Interfaces can only have static IP addresses. Any IP address can be added to a Tunnel Interface, even if the same IP address is used on another interface or as a loopback IP address. Adding an IP address allows you to define the source IP address of traffic sent from the engine node itself. For example, an IP address is recommended to provide a source IP address for dynamic routing daemons, for IGMP proxy, and for Protocol Independent Multicast - Sparse-Mode (PIM-SM) configuration. If no IP address is added to the Tunnel Interface, the source IP address for traffic sent from the engine node is automatically selected according to the **Bypass Default IP Address** setting in the Interface Options for the firewall. See Setting Interface Options for Firewalls (page 453).

The mapping of Tunnel Interfaces to physical network interfaces on the engine is done automatically by the engine operating system based on the routing configuration.

For settings on the **Advanced** tab, see Configuring Advanced Interface Properties for Firewalls (page 436).

▼ To define Tunnel Interfaces

1. Open the Firewall Properties and switch to the **Interfaces** tab.
2. Click **Add** and select **Tunnel Interface**. The Tunnel Interface Properties dialog opens.

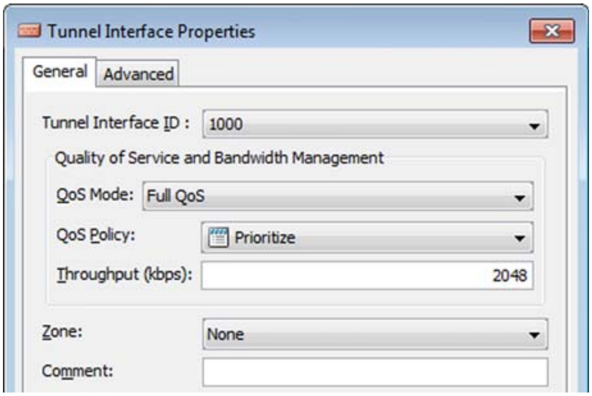

Tunnel Interface Properties

Tunnel Interface ID : 1000

Quality of Service and Bandwidth Management

QoS Mode: Full QoS

QoS Policy: Prioritize

Throughput (kbps): 2048

Zone: None

Comment:

3. Select the **Tunnel Interface ID**.
4. (*Optional*) Select the **QoS Mode** to define how QoS is applied to the interface.
5. (*Full QoS or DSCP Handling and Throttling modes only*) Select the **QoS Policy** for the link on the interface. For more information, see Getting Started with QoS (page 782).

6. (*Full QoS mode only*) Enter the **Throughtput** for the link on this interface.
   - Enter throughput as kilobits per second (for example, 2048).
   - The throughput is for uplink speed (outgoing traffic) and typically must correspond to the speed of an Internet link (such as an ADSL line), or the combined speeds of several such links when they are connected to a single physical interface).

Caution – Make sure you set the Interface speed correctly. When the bandwidth is set, the engine always scales the total amount of traffic on this interface to the bandwidth you defined. This happens even if there are no bandwidth limits or guarantees defined for any traffic.

7. Click **OK**.

What's Next?

▶ If you want to define the source IP address for traffic sent from the engine node, proceed to the next relevant section below:
   - Adding IPv4 Addresses for a Single Firewall (page 440)
   - Adding IPv4 Addresses for a Firewall Cluster (page 445).
   - Adding IPv6 Addresses for a Single Firewall (page 443)
   - Adding IPv6 Addresses for a Firewall Cluster (page 447).

▶ If you do not want to add IP addresses to the Tunnel Interface, define how the source IP address for traffic sent from the engine node is selected as explained in Setting Interface Options for Firewalls (page 453).

▶ Otherwise, close the Firewall's Properties dialog and refresh the Firewall's policy to transfer the configuration changes.

Related Tasks

▶ Defining Tunnel Interfaces for Virtual Security Engines (page 502)

Configuring Advanced Interface Properties for Firewalls

Advanced settings allow you to configure SYN Rate Limits, Log Compression, and IPv6 Router Advertisements on the Firewall's interfaces.

Log Compression is useful when the routing configuration generates a large volume of antispoofing logs or the number of Discard logs becomes high (for example, as a result of a SYN flood attack).

Router advertisements are packets that contain network layer configuration parameters. Enabling IPv6 Router Advertisements allows devices that connect to the same IPv6 network as the Firewall to acquire IP addresses automatically. This eliminates the need for a separate DHCP server. The Router Advertisement messages specify what configuration information the Firewall has available.

Note – The SYN Rate Limits and Log Compression settings in the interface properties override the general SYN Rate Limits and Log Compression settings that are defined on the Advanced tab in the Firewall or Firewall Cluster properties. See Configuring Default SYN Rate Limits (page 579) and Configuring Default Log Handling Settings (page 580).

▼ To configure advanced interface properties for Firewalls

1. In the properties dialog for the Firewall, switch to the **Interfaces** tab.
2. Right-click a Physical Interface, a VLAN, an ADSL Interface, or a Wireless Interface and select **Edit Physical Interface**, **Edit VLAN Interface**, **Edit ADSL Interface, Edit Tunnel Interface**, or **Edit Wireless Interface**. The properties dialog for the interface opens.
3. Switch to the **Advanced** tab.

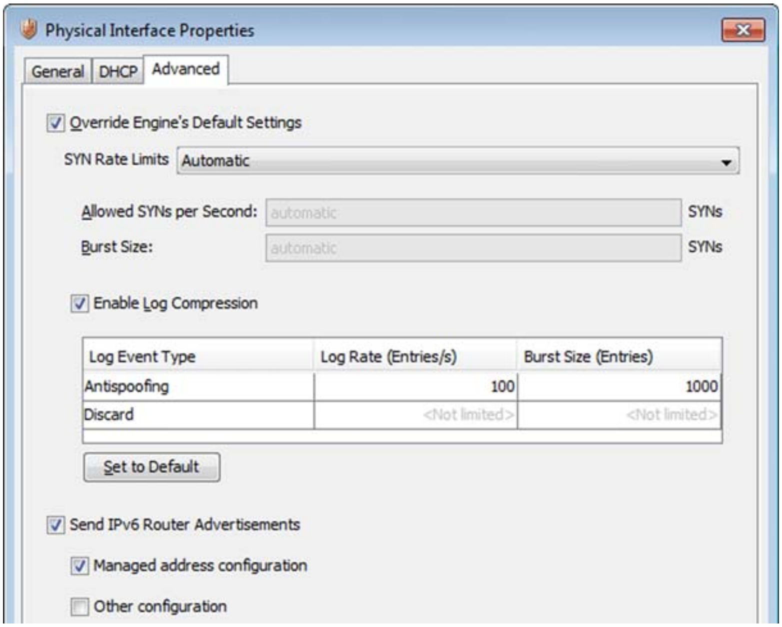

☑ Override Engine's Default Settings

SYN Rate Limits: Automatic

Allowed SYNs per Second: automatic SYNs

Burst Size: automatic SYNs

☑ Enable Log Compression

Log Event Type	Log Rate (Entries/s)	Burst Size (Entries)
Antispoofing	100	1000
Discard	<Not limited>	<Not limited>

☑ Send IPv6 Router Advertisements

☑ Managed address configuration

☐ Other configuration

4. Select **Override Engine's Default Settings**. The options for SYN Rate Limits and Log Compression are enabled.
5. (*Optional*) Define the **SYN Rate Limits**

Setting	Description
Default	The interface uses the SYN Rate Limits defined on the Advanced tab in the Firewall properties (see Configuring Default SYN Rate Limits (page 579)).
Off	Disables SYN Rate Limits on the interface.
Automatic	This is the recommended mode if you want to override the general SYN Rate Limits defined on the Advanced tab in the engine's properties. The engine automatically calculates the number of **Allowed SYNs per Second** (the number of allowed SYN packets per second) and the **Burst Size** (the number of allowed SYNs before the engine starts limiting the SYN rate) for the interface based on the engine's capacity and memory size.
Custom	Enter the desired values for **Allowed SYNs per Second** and **Burst Size**. We recommend that the Burst Size be at least one tenth of the Allowed SYNs per Second value. If the Burst Size is too small, SYN Rate Limits do not work. For example, if the value for Allowed SYNs per Second is 10000, the Burst Size should be at least 1000.

Caution – The recommended values for the SYN Flood Settings depend on your network environment. If the Custom settings are not carefully configured, the capacity of the Firewall engine may suffer or SYN Rate Limits may not work correctly.

6. (*Optional*) Enable/Disable **Log Compression** and enter the desired values for the Antispoofing entries and (optionally) for Discard entries.

Setting	Description
Log Rate (Entries/s)	The maximum number of entries per second. The default value for antispoofing entries is 100 entries/s. By default, Discard log entries are not compressed.
Burst Size (Entries)	The maximum number of matching entries in a single burst. The default value for antispoofing entries is 1000 entries. By default, Discard log entries are not compressed.

- Do not enable Log Compression if you want all the antispoofing and Discard entries to be logged as separate log entries (for example, for reporting purposes or for Firewall statistics).
- By default, each generated Antispoofing and Discard log entry is logged separately and displayed as a separate entry in the Logs view. Log Compression settings allow you to define the maximum number of separately logged entries. When the defined limit is reached, a single antispoofing log entry or Discard log entry is logged. The single entry contains information on the total number of the generated Antispoofing log entries or Discard log entries. After this, the logging returns to normal and all the generated entries are once more logged and displayed separately.

7. (*Optional*) Select **Send IPv6 Router Advertisements** and specify what configuration information is offered in the Router Advertisement messages:

Setting	Description
Managed address configuration	The Firewall offers IPv6 addresses over the Dynamic Host Configuration Protocol (DHCPv6). When this option is selected, any other available configuration information is also sent in the DHCPv6 communications.
Other configuration	The Firewall offers other configuration information, such as DNS-related information or information on other servers within the network, over DHCPv6.

8. Click **OK**.

What's Next?

▶ Close the Firewall's Properties dialog and refresh the Firewall's policy to transfer the configuration changes.

Configuring Single Firewall IP Addresses

A Single Firewall's interfaces can have the following types of IP addresses:

- A Physical Interface can have one or more static or dynamic IP addresses. A Physical Interface can have multiple dynamic IP addresses only if you add VLAN Interfaces on the Physical Interface and the VLAN Interfaces each have a dynamic IP address. Otherwise, a Physical Interface can only have a single dynamic IP address.
- A VLAN Interface and an ADSL Interface can have one or more static IP addresses or a single dynamic IP address.
- A Wireless Interface's SSID Interface can have a single static IP address.
- A Modem Interface always has a dynamic IP address that is assigned automatically by a PPP daemon.
- A Tunnel Interface can have one or more static IP addresses.

Only IPv4 addresses are supported as dynamic IP addresses.

You may need to define a *contact address* if you enter a private static address and NAT is used to translate it to a different external IP address. The external IP address must be configured as the contact address if other Management Center components need to use the external IP address to contact this Firewall or if the IP address is used as a VPN end-point. Only IPv4 addresses are used in system communications. Refer to Defining Contact IP Addresses (page 66) for more information.

What's Next?

▶ To add an IPv4 address to a Physical Interface, VLAN Interface, ADSL Interface, SSID Interface, or Tunnel Interface, proceed to Adding IPv4 Addresses for a Single Firewall (page 440).

▶ To add an IPv6 address to a Physical Interface, VLAN Interface, SSID Interface, or Tunnel Interface, proceed to Adding IPv6 Addresses for a Single Firewall (page 443).

Adding IPv4 Addresses for a Single Firewall

IPv4 addresses are supported on Physical Interfaces, VLAN Interfaces, SSID Interfaces, and Tunnel Interfaces.

> **Note – If you use VLAN Interfaces, you must add the IPv4 Addresses to the VLAN Interfaces.**

▼ To add an IPv4 address for a Single Firewall

1. In the properties dialog for the Firewall, switch to the **Interfaces** tab.
2. Right-click a Physical Interface, VLAN Interface, SSID Interface, or Tunnel Interface and select **New→IPv4 Address**, or right- click an ADSL Interface and select **New IPv4 Address**. The IP Address Properties dialog opens.

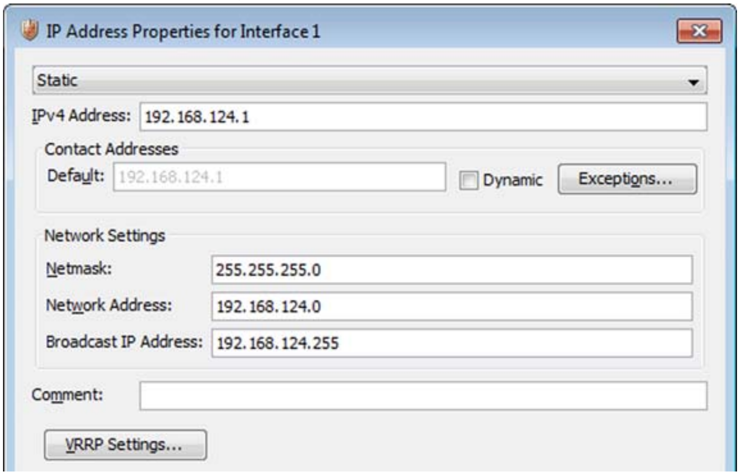

IP Address Properties for Interface 1

Static

IPv4 Address: 192.168.124.1

Contact Addresses

Default: 192.168.124.1 □ Dynamic

Network Settings

Netmask: 255.255.255.0

Network Address: 192.168.124.0

Broadcast IP Address: 192.168.124.255

Comment:

3. Configure the IP address information in one of the following ways:
   - Select **Static** and enter the **IPv4 Address**. The Network Settings are automatically entered.
   - (*Physical, VLAN, and ADSL Interfaces only*) Select **Dynamic** and select the **Dynamic index**. The Dynamic Index is a number for your own reference to identify the dynamic IP address.
4. If necessary, define the contact address information (see Defining Contact Addresses for a Single Engine or a Cluster Virtual IP Address (page 68) for detailed instructions).
   - Enter the **Default** contact address or select **Dynamic** (next to the **Default** field) if the interface has a dynamic contact address. The Default contact address is used by default whenever a component that belongs to another Location connects to this interface.
   - If components from some Locations cannot use the Default contact address, click **Exceptions** to define Location-specific contact addresses.
5. (*Optional*) Configure additional features for this interface:
   - (*Static IP address only*) If you want to use VRRP on the Physical Interface, VLAN Interface, or SSID Interface, proceed to Configuring VRRP Settings for Single Firewalls (page 441).
   - (*Dynamic IP address only*) If the interface requires PPPoE or PPPoA support, proceed to Configuring PPP Settings for Single Firewalls (page 442).
6. Click **OK**.
7. Repeat from Step 2 to add further IP addresses to the same or other interface.

What's Next?

- To add IPv6 addresses to a Physical Interface, VLAN Interface, SSID Interface, or Tunnel Interface, proceed to Adding IPv6 Addresses for a Single Firewall (page 443).
- To define Modem Interfaces, proceed to Defining Modem Interfaces for Single Firewalls (page 448).
- If you are creating a new Firewall, or if you want to change the roles the different interfaces have in the configuration, proceed to Setting Interface Options for Firewalls (page 453).
- If you are finished adding IP addresses to Tunnel Interfaces, continue the configuration of the Route-Based VPN by Defining Routing for the Route-Based VPN (page 604).
- Otherwise, close the Firewall's Properties dialog and refresh the Firewall's policy to transfer the configuration changes.

Related Tasks

- Configuring Firewall Cluster IP Addresses (page 444)
- Configuring Loopback IP Addresses for Firewalls (page 450)
- About Using a Dynamic IP Address on a Firewall Interface (page 455)
- Configuring Manual ARP Settings (page 510)
- Routing Multicast Traffic (page 593)

Configuring VRRP Settings for Single Firewalls

Virtual Router Redundancy Protocol (VRRP) allows high-availability router configurations. VRRP support is only available on Physical Interfaces, VLAN Interfaces, and SSID Interfaces of Single Firewalls. One virtual router can be configured for each Physical Interface, VLAN Interface, or SSID Interface. The virtual router can have either backup or active status. The virtual router configured for one interface does not take into account the status of possible virtual routers configured on other interfaces. VRRP is only supported for static IPv4 addresses.

▼ To set up VRRP

1. Click the **VRRP Settings** button in the IP Address properties. The VRRP Settings dialog opens.

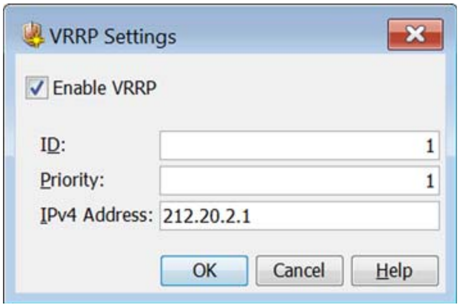

VRRP Settings

☑ Enable VRRP

ID: 1

Priority: 1

IPv4 Address: 212.20.2.1

2. Select **Enable VRRP**.
3. Enter the **ID**, **Priority**, and **IPv4 Address** according to the configuration of the virtual router.
4. Click **OK** to close the VRRP Settings dialog.

5. Click **OK** to close the IP Address Properties dialog.

What's Next?

- To add further IP addresses on the same or a different interface, repeat the steps in Adding IPv4 Addresses for a Single Firewall (page 440) or Adding IPv6 Addresses for a Single Firewall (page 443).
- To define Modem Interfaces, proceed to Defining Modem Interfaces for Single Firewalls (page 448).
- If you are adding a new Firewall, or if you want to change the roles the different interfaces have in the configuration, proceed to Setting Interface Options for Firewalls (page 453).
- Otherwise, close the Firewall's Properties dialog and refresh the Firewall's policy to transfer the configuration changes.

Configuring PPP Settings for Single Firewalls

Point-to-Point Protocol over Ethernet (PPPoE) and Point-to-Point Protocol over ATM (PPPoA) clients on the Firewall simplify the installation of a Firewall appliance when PPPoE or PPPoA are used. Activating the PPPoE or PPPoA client on the Firewall allows connecting the Firewall to an external ADSL modem without having to configure routing and NAT settings on the ADSL modem (the modem is in transparent bridge mode). You must also activate the PPPoE or PPPoA client on the Firewall if you have a specific Firewall appliance that has an ADSL port and the ISP for the ADSL connection requires PPPoE or PPPoA. The PPPoE client is supported only for dynamic IPv4 addresses on Single Firewalls.

▼ To configure the PPP client settings

1. Click the **PPP Settings** button in the IP Address properties. The PPP Settings dialog opens.

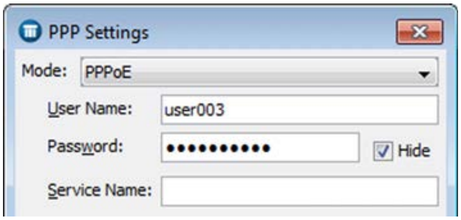

PPP Settings

Mode: PPPoE

User Name: user003

Password: ••••••••••• ☑ Hide

Service Name:

2. Select the **Mode** that the ADSL modem connected to the interface supports:
   - **PPPoE**: can be used with Physical Interfaces or ADSL Interfaces.
   - **PPPoA**: can be used with ADSL interfaces only.
3. Enter the **User Name** and **Password**, and optionally the **Service Name**, according to the local PPPoE or PPPoA configuration.

Tip – Deselect Hide if you want to view the Password in plain text in the dialog.

4. Click **OK** to close the PPP Settings dialog.
5. Click **OK** to close the IP Address Properties dialog.

What's Next?

- To add further IP addresses on a different interface, repeat the steps in Adding IPv4 Addresses for a Single Firewall.
- To define Modem Interfaces, proceed to Defining Modem Interfaces for Single Firewalls (page 448).
- If you are configuring a new Firewall, or if you want to change the roles the different interfaces have in the configuration, proceed to Setting Interface Options for Firewalls (page 453).
- Otherwise, close the Firewall's Properties dialog and refresh the Firewall's policy to transfer the configuration changes.

Adding IPv6 Addresses for a Single Firewall

IPv6 addresses are supported on Physical Interfaces, VLAN Interfaces, SSID Interfaces, and Tunnel Interfaces.

Note – If you use VLAN Interfaces, you must add the IPv4 Addresses to the VLAN Interfaces.

▼ To add an IPv6 address for a Single Firewall

1. In the properties dialog for the Firewall, switch to the **Interfaces** tab.
2. Right-click a Physical Interface, VLAN Interface, SSID Interface, or Tunnel Interface and select **New→IPv6 Address**. The IP Address Properties dialog opens.

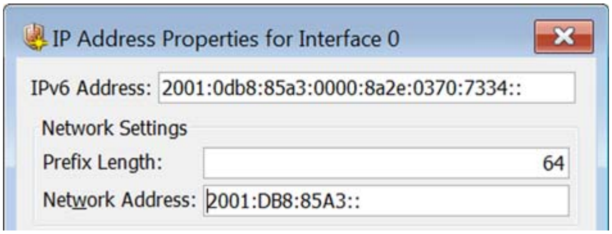

IPv6 Address: 2001:0db8:85a3:0000:8a2e:0370:7334::

Network Settings

Prefix Length: 64

Network Address: 2001:DB8:85A3::

3. Enter the **IPv6 Address**.
4. Check the automatically filled-in **Prefix Length** and adjust it if necessary by entering a value between 0-128. The Network Address is automatically generated.
5. Click **OK**.
6. Repeat from Step 2 to add more IPv6 addresses to the same or other Physical Interface, VLAN Interface, SSID Interface, or Tunnel Interface.

What's Next?

- To define Modem Interfaces, proceed to Defining Modem Interfaces for Single Firewalls (page 448).
- If you are creating a new Firewall, or if you want to change the roles the different interfaces have in the configuration, proceed to Setting Interface Options for Firewalls (page 453).
- If you are finished adding IP addresses to Tunnel Interfaces, continue the configuration of the Route-Based VPN by Defining Routing for the Route-Based VPN (page 604).
- Otherwise, close the Firewall's Properties dialog and refresh the Firewall's policy to transfer the configuration changes.

Related Tasks

- Adding IPv4 Addresses for a Single Firewall (page 440)
- Configuring Loopback IP Addresses for Firewalls (page 450)

Configuring Firewall Cluster IP Addresses

You can configure several IP addresses on each Physical Interface, VLAN Interface, or Tunnel Interface of a Firewall Cluster.

When you add an IP address to a Physical Interface, VLAN Interface, or Tunnel Interface, you must also select the purpose of the IP address. There are two types of IP addresses on Firewall Cluster interfaces:

- A *Cluster Virtual IP Address* (CVI) is a common IP address for all the nodes in the cluster. It is used for handling all of the general network traffic (traffic you want to filter with the Firewall) routed from/to the cluster.
- A *Node Dedicated IP Address* (NDI) is for handling traffic that is originating from or destined to a node in the cluster (for example, for the heartbeat connections between the nodes). Each node has a unique IP address that is used as the Node Dedicated IP Address.

You may also need to define a *contact address* if the Cluster Virtual IP Address or Node Dedicated IP Address is private and NAT is used to translate the IP address to a different external IP address. The external IP address must be configured as the contact address if other Management Center components need to use the external IP address to contact this Firewall (Node Dedicated IP Address) or if this IP address is a VPN end-point (Cluster Virtual IP Address). Refer to Defining Contact IP Addresses (page 66) for more information and configuration instructions for all components.

What's Next?

- Adding IPv4 Addresses for a Firewall Cluster (page 445)
- Adding IPv6 Addresses for a Firewall Cluster (page 447)

Adding IPv4 Addresses for a Firewall Cluster

IPv4 addresses are supported on Physical Interfaces, VLAN Interfaces, and Tunnel Interfaces.

Note – If you use VLAN Interfaces, you must add the IPv4 Addresses to the VLAN Interfaces.

▼ To add IPv4 addresses for a Firewall Cluster

1. In the properties dialog for the Firewall Cluster, switch to the **Interfaces** tab.
2. Right-click a Physical interface, VLAN Interface, or Tunnel Interface and select **New→IPv4 Address**. The IP Address Properties dialog opens.

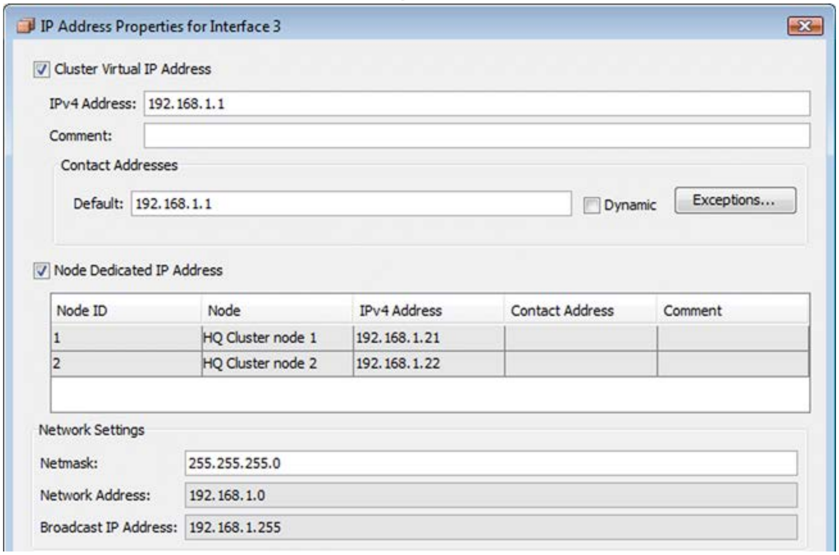

IP Address Properties for Interface 3

☑ Cluster Virtual IP Address

IPv4 Address: 192.168.1.1

Comment:

Contact Addresses

Default: 192.168.1.1 □ Dynamic Exceptions...

☑ Node Dedicated IP Address

Node ID	Node	IPv4 Address	Contact Address	Comment
1	HQ Cluster node 1	192.168.1.21		
2	HQ Cluster node 2	192.168.1.22		

Network Settings

Netmask: 255.255.255.0

Network Address: 192.168.1.0

Broadcast IP Address: 192.168.1.255

3. Select the types of IP addresses that you want to add. By default, both **Cluster Virtual IP Address** and **Node Dedicated IP Address** are selected.
   - If the interface does not receive or send traffic that the Firewall examines, there is no need to define a Cluster Virtual IP Address.
   - We recommend that you add a Node Dedicated IP Address for each (sub)network that is located behind the physical interface.
4. If you are adding a Cluster Virtual IP Address, enter the **IPv4 Address** that is used as the Cluster Virtual IP Address.
5. If you are adding a Node Dedicated IP Address for the nodes, set the **IPv4 Address** for each node by double-clicking the field.
6. Check the automatically filled-in **Netmask** and adjust it as necessary.
7. If necessary, define the contact address(es) for the Cluster Virtual IP Address (see Defining Contact Addresses for a Single Engine or a Cluster Virtual IP Address (page 68) for detailed instructions):
   - Enter the **Default** contact address or select **Dynamic** if the interface has a dynamic default contact address. The default contact address is used by default whenever a component that belongs to another Location connects to this interface.
   - If components from some Locations cannot use the Default contact address, click **Exceptions** to define Location-specific contact addresses.

8. If necessary, define the IP contact address for each node in the Node Dedicated IP Address table by double-clicking the Contact Address cell. The Exceptions dialog opens (see Defining Contact Addresses for Node Dedicated IP Addresses (page 69) for detailed instructions).
   - Enter the **Default** contact address at the top of the dialog. The Default contact address is used by default whenever a component that belongs to another Location connects to this interface.
   - If components from some Locations cannot use the Default contact address, click **Add** to define Location-specific contact addresses.
9. Click **OK**.
10. Repeat from Step 2 to add further IPv4 addresses to the same or other interface.

What's Next?

- If you are configuring a new Firewall, or if you want to change the roles the different interfaces have in the configuration, proceed to Setting Interface Options for Firewalls (page 453).
- If you are finished adding IP addresses to Tunnel Interfaces, continue the configuration of the Route-Based VPN by Defining Routing for the Route-Based VPN (page 604).
- Otherwise, close the Firewall's Properties dialog and refresh the Firewall's policy to transfer the configuration changes.

Related Tasks

- Configuring Single Firewall IP Addresses (page 439)
- Configuring Loopback IP Addresses for Firewalls (page 450)
- Configuring Manual ARP Settings (page 510)
- Routing Multicast Traffic (page 593)

Adding IPv6 Addresses for a Firewall Cluster

IPv6 addresses are supported on Physical Interfaces, VLAN Interfaces, and Tunnel Interfaces.

Note – If you use VLAN Interfaces, you must add the IPv6 Addresses to the VLAN Interfaces.

▼ To add IPv6 addresses for a Firewall Cluster

1. In the properties dialog for the Firewall, switch to the **Interfaces** tab.
2. Right-click a Physical interface, VLAN Interface, or Tunnel Interface and select **New→IPv6 Address**. The IP Address Properties dialog opens.

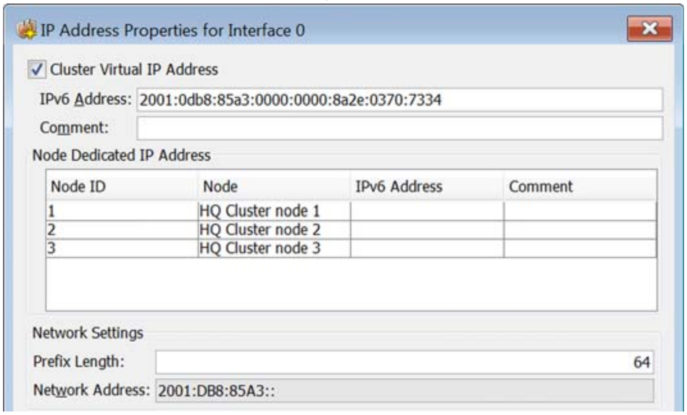

IP Address Properties for Interface 0

☑ Cluster Virtual IP Address

IPv6 Address: 2001:0db8:85a3:0000:0000:8a2e:0370:7334

Comment:

Node Dedicated IP Address

Node ID	Node	IPv6 Address	Comment
1	HQ Cluster node 1		
2	HQ Cluster node 2		
3	HQ Cluster node 3		

Network Settings

Prefix Length: 64

Network Address: 2001:DB8:85A3::

3. Select the types of IP addresses that you want to add. By default, both **Cluster Virtual IP Address** and **Node Dedicated IP Address** are selected.
   - If the interface does not receive or send traffic that the Firewall examines, there is no need to define a Cluster Virtual IP Address.
   - We recommend that you add a Node Dedicated IP Address for each (sub)network that is located behind the physical interface.
4. If you are adding a Cluster Virtual IP Address, enter the **IPv6 Address** that is used as the Cluster Virtual IP Address.
5. If you are adding a Node Dedicated IP Address for the nodes, set the **IPv6 Address** for each node by double-clicking the field.
6. *(Optional)* Modify the **Prefix Length** (0-128).
7. Check the automatically filled-in **Network Address** and adjust it as necessary.
8. Click **OK**.
9. *(Optional)* Repeat from Step 2 to add further IPv6 addresses to the same or other interface.

What's Next?

▶ If you are finished adding IP addresses to Tunnel Interfaces, continue the configuration of the Route-Based VPN by Defining Routing for the Route-Based VPN (page 604).

▶ Otherwise, close the Firewall Cluster Properties dialog and refresh the policy to transfer the configuration changes.

Related Tasks

- Configuring Single Firewall IP Addresses (page 439)
- Configuring Loopback IP Addresses for Firewalls (page 450)
- Configuring Manual ARP Settings (page 510)
- Routing Multicast Traffic (page 593)

Defining Modem Interfaces for Single Firewalls

A Modem Interface defines the settings of a 3G modem that provides a wireless outbound link for a Single Firewall. The numbering of a Modem Interface in the Management Client (the modem number) is mapped to the modem's IMEI (international mobile equipment identity) number, and each modem is assigned a unique ID when you connect the modem to the engine. You can change the mapping between the modem's IMEI number and the modem ID through the engine command line, if necessary.

▼ To define a Modem Interface

1. In the properties dialog for the Firewall, switch to the **Interfaces** tab.
2. Right-click the empty space and select **New→Modem Interface**. The Modem Interface Properties dialog opens.

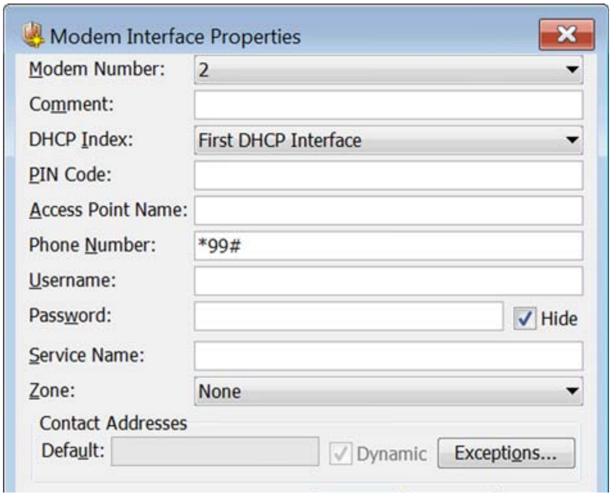

Modem Interface Properties

Modem Number: 2

Comment:

DHCP Index: First DHCP Interface

PIN Code:

Access Point Name:

Phone Number: *99#

Username:

Password: ☑ Hide

Service Name:

Zone: None

Contact Addresses

Default: ☑ Dynamic Exceptions...

3. Define the basic Modem Interface properties as explained in the table below.

Table 28.6 Modem Interface Properties

Option	Explanation
Modem Number	Select the Modem Number to represent the modem in the configuration.
DHCP Index	Select the DHCP Index. The DHCP Index is a number for your own reference to identify the DHCP interface.

Table 28.6 Modem Interface Properties (Continued)

Option	Explanation
PIN Code (*Optional*)	Enter the PIN code if it is needed for the modem's SIM card. If the PIN code is included in the configuration and you change the 3G modem's SIM card after configuring the Firewall, you must change the PIN code as described in Changing/Removing the PIN Code of a Modem Interface (page 450).
Phone Number (*Optional*)	Enter the modem's Phone Number if it differs from the default phone number.
Zone (*Optional*)	Select the network **Zone** to which the Modem interface belongs from the list or select **Other** to select another Zone. If the Zone is not listed, create a new Zone element through the New icon at the top of the dialog. See Defining Zone Elements (page 730) for more information.

4. Enter the rest of the information (**Access Point Name**, **Username** and **Password**, and **Service Name**) according to the instructions that you have received from your service provider.
5. If necessary, define the contact address information (see Defining Contact Addresses for a Single Engine or a Cluster Virtual IP Address (page 68) for detailed instructions).
   - Enter the **Default** contact address or select **Dynamic** (next to the **Default** field) if the interface has a dynamic contact address. The Default contact address is used by default whenever a component that belongs to another Location connects to this interface.
   - If components from some Locations cannot use the Default contact address, click **Exceptions** to define Location-specific contact addresses.
6. Click **OK**.

What's Next?

▶ If you are configuring a new Firewall, or if you want to change the roles the different interfaces have in the configuration, proceed to Setting Interface Options for Firewalls (page 453).

▶ Otherwise, close the Firewall's Properties dialog and adjust the routing configuration as necessary as explained in Adding Routes for Firewalls (page 589).

Related Tasks

▶ Changing/Removing the PIN Code of a Modem Interface (page 450)

▶ Configuring Manual ARP Settings (page 510)

▶ Routing Multicast Traffic (page 593)

Changing/Removing the PIN Code of a Modem Interface

If you change the SIM card of a 3G modem that is connected to a Single Firewall and the PIN code is enabled on the new SIM card, you must change the PIN code information on the related Modem Interface. You must remove the PIN code if a PIN code was enabled on the old SIM card but a PIN code is not enabled on the new SIM card.

You must also change the PIN code if you have received the message "PIN code differs from initial contact settings" at policy installation after you have finished configuring the Firewall. This message means that the PIN code in the Modem Interface properties does not match the PIN code that you have entered for the 3G modem in the engine's command line interface. If there are other problems with the PIN code, you can check through the engine's command line interface which PIN code was used in the initial configuration and change the PIN code information if necessary.

▼ To change or remove the PIN code of a modem interface

1. Run the `sg-reconfigure` tool on the command line and change the 3G modem's PIN code information in the engine's command line interface. See the *Firewall Installation Guide* for more information.
2. Open the properties of the Single Firewall in the Management Client and switch to the **Interfaces** tab.
3. Right-click the Modem Interface and select **Properties**. The Modem Interface Properties dialog opens.
4. Change the PIN code information and click **OK**.

What's Next?

▶ Close the Firewall's Properties dialog and refresh the Firewall's policy to transfer the configuration changes.

Related Tasks

▶ Defining Modem Interfaces for Single Firewalls (page 448)

Configuring Loopback IP Addresses for Firewalls

A Loopback IP address allows the firewall to communicate with itself. Loopback IP addresses are not connected to any physical interface and they do not create connectivity to any network. Loopback IP addresses only route traffic from the device to itself.

- You can add several loopback IP addresses to each Firewall.
- Any IP address that is not already used as a Cluster Virtual IP Addresses (CVI) or Node Dedicated IP Addresses (NDI) on another interface can be used as a loopback IP address.
- The same IP address can be used as a loopback IP address and as the IP address of a Tunnel Interface.
- Loopback IP addresses can be used as the Identity for Authentication Requests, the Source for Authentication Requests, and the Default IP Address for Outgoing Traffic.
- Loopback IP addresses cannot be used as Control Interfaces or Heartbeat Interfaces.

What's Next?

- ▶ Adding Loopback IP addresses for Single Firewalls (page 451)
- ▶ Adding Loopback IP Addresses for Firewall Clusters (page 452)

Adding Loopback IP addresses for Single Firewalls

▼ To add a loopback IP address for a Single Firewall

1. In the properties dialog for the Firewall, switch to the **Interfaces** tab.
2. Click **Options**. The Interface Options dialog opens.
3. Switch to the **Loopback** tab.

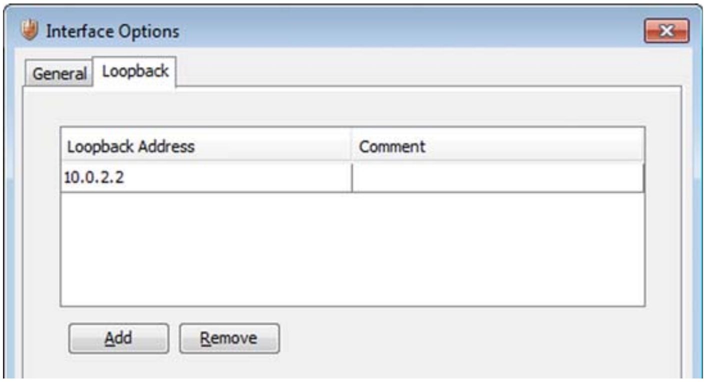

Loopback Address	Comment
10.0.2.2	

4. Click **Add**. A row is added to the table.
5. Click the **Loopback Address** cell and enter the loopback IP address.

Note – If the IP address you want to use as a loopback IP address is already used as a Cluster Virtual IP Addresses (CVI) or Node Dedicated IP Addresses (NDI) on another interface, you must remove the IP address from the interface configuration before using it as a loopback IP address.

6. Click **OK**.

What's Next?

- ▶ Close the Firewall's Properties dialog and refresh the Firewall's policy to transfer the configuration changes.

Adding Loopback IP Addresses for Firewall Clusters

You can define both Cluster Virtual IP Addresses (CVI) and Node Dedicated IP Addresses (NDI) as loopback IP addresses for Firewall Clusters. Whether to define a CVI or NDI loopback address depends on the traffic for which the loopback IP address is used:

- A CVI loopback IP address is used for loopback traffic that is sent to the whole cluster. It is shared by all the nodes in the cluster.
- An NDI loopback IP address is used for loopback traffic that is sent to a specific node in the cluster. NDI loopback IP addresses must be unique for each node.

> **Note – If the IP address you want to use as a loopback IP address is already used as a CVI or NDI on another interface, you must remove the IP address from the interface configuration before using it as a loopback IP address.**

If you use NDI Loopback IP addresses, you must define an NDI loopback IP address for all nodes.

▼ To add a loopback IP address for a Firewall Cluster

1. In the properties dialog for the Firewall Cluster, switch to the **Interfaces** tab.
2. Click **Options**. The Interface Options dialog opens.
3. Switch to the **Loopback** tab.

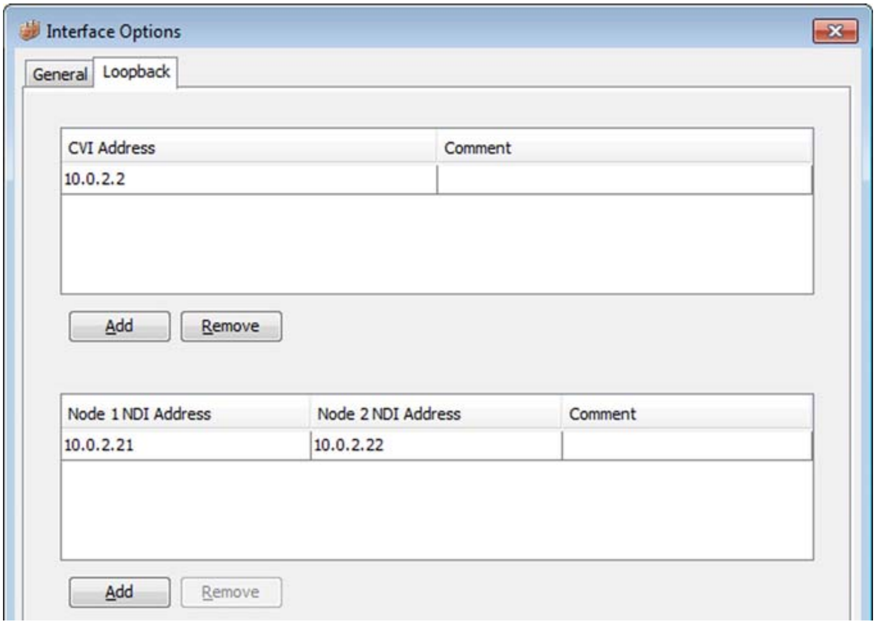

CVI Address	Comment
10.0.2.2	

Node 1 NDI Address	Node 2 NDI Address	Comment
10.0.2.21	10.0.2.22	

4. Click **Add** below the table to which you want to add a loopback IP address. A row is added to the table.
5. (*CVI only*) Click the **CVI Address** cell and enter the loopback IP address for the cluster.
6. (*NDI only*) Click the **NDI Address** cell for each node and enter the node-specific loopback IP address.
7. Click **OK**.

What's Next?

▶ Close the Firewall's Properties dialog and refresh the Firewall's policy to transfer the configuration changes.

Setting Interface Options for Firewalls

The Interface Options dialog contains the settings for selecting which IP addresses are used in particular roles in system communications (for example, in communications between the Firewall and the Management Server). Only IPv4 addresses are used in system communications.

▼ To set the interface options

1. In the properties dialog for the Firewall, switch to the **Interfaces** tab.
2. Click **Options**. The Interface Options dialog opens.

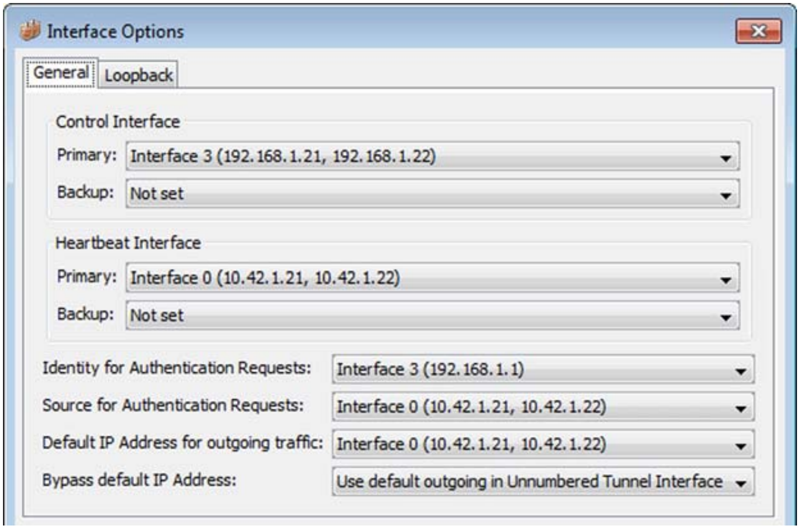

Control Interface

Primary: Interface 3 (192.168.1.21, 192.168.1.22)

Backup: Not set

Heartbeat Interface

Primary: Interface 0 (10.42.1.21, 10.42.1.22)

Backup: Not set

Identity for Authentication Requests: Interface 3 (192.168.1.1)

Source for Authentication Requests: Interface 0 (10.42.1.21, 10.42.1.22)

Default IP Address for outgoing traffic: Interface 0 (10.42.1.21, 10.42.1.22)

Bypass default IP Address: Use default outgoing in Unnumbered Tunnel Interface

3. Select the interface options as explained in the table below.

Table 28.7 Firewall Interface Options

Option	Explanation
Control Interface	Select the **Primary** Control Interface for Management Server contact. This interface is used for communications with the Management Server
	(*Optional*) Select a **Backup** Control Interface that is used if the Primary interface is not available.

Table 28.7 Firewall Interface Options (Continued)

Option	Explanation
Node-Initiated contact to Management Server (*Single Firewalls only*)	Select the Node-Initiated contact to Management Server option if the Firewall is behind a device that applies dynamic NAT to the inbound management connections or blocks them. This option is always used with a dynamic control IP address, so it is always fixed to the selected state if the IP address is dynamic. When this option is selected, the engine opens a connection to the Management Server and maintains connectivity. If the connection is not open when you command the engine through the Management Client, the command is left pending until the engine opens the connection again.
Heartbeat Interface *(Firewall Clusters only)*	Select the **Primary** Heartbeat Interface for communications between the nodes. We recommend that you use a Physical Interface, not a VLAN Interface. We strongly recommend that you do not direct any other traffic through this interface. A dedicated network helps ensure reliable and secure operation. **Caution!** Primary and Backup Heartbeat networks exchange confidential information. If dedicated networks are not possible, configure the cluster to encrypt the exchanged information. See Adjusting General Firewall Clustering Options (page 556).
	Select a **Backup** Heartbeat Interface that is used if the Primary Heartbeat Interface is unavailable. It is not mandatory to configure a backup heartbeat, but we strongly recommend it. If heartbeat traffic is not delivered, the cluster cannot operate and traffic will be disturbed. We strongly recommend that you use a dedicated interface for the backup heartbeat as well. **Caution!** Primary and Backup Heartbeat networks exchange confidential information. If dedicated networks are not possible, configure the cluster to encrypt the exchanged information. See Adjusting General Firewall Clustering Options (page 556).
Identity for Authentication Requests	The IP address of the selected interface is used when an engine contacts an external authentication server and it is also displayed (by default) to end-users in Telnet-based authentication. This option does not affect the routing of the connection with the authentication server. The IP address is used only as a parameter inside the authentication request payload to give a name to the request sender.
Source for Authentication Requests	By default, the source IP address for authentication requests is selected according to routing. If the authentication requests are sent to an external authentication server over VPN, select an interface with a Node Dedicated IP address that you want use for the authentication requests. If you use the 5.4 Authentication Server component for authenticating the users, this setting does not have any effect.

Table 28.7 Firewall Interface Options (Continued)

Option	Explanation
Default IP Address for Outgoing Traffic	This option defines the IP address that the nodes use if they have to initiate connections (system communications, ping, etc.) through an interface that has no Node Dedicated IP Address. In Firewall Clusters, you must select an interface that has an IP address defined for all nodes.
Bypass Default IP Address	Defines how the source IP address for traffic sent from the engine node is selected for Tunnel Interfaces that do not have IP addresses. **Use Link Address in Unnumbered Tunnel Interface**: The loopback IP address defined for the engine node is used as the source IP address of traffic sent from the engine node. See Configuring Loopback IP Addresses for Firewalls (page 450). This allows the loopback IP address to be used as an end-point in the Route-Based VPN. **Use Default Outgoing in Unnumbered Tunnel Interface**: the Default IP Address for Outgoing Traffic is used as the source IP address of traffic sent from the engine node.

4. Click **OK**.

What's Next?

- If this is a new Firewall, close the Firewall Properties dialog. Configure the routing as explained in Adding Routes for Firewalls (page 589), and initialize the engines as explained in Saving an Initial Configuration for Security Engines (page 515).
- If you are configuring Tunnel Interfaces for the Route-Based VPN, continue the configuration by Defining Routing for the Route-Based VPN (page 604).
- Otherwise, close the Firewall's Properties dialog and refresh the Firewall's policy to transfer the configuration changes.

About Using a Dynamic IP Address on a Firewall Interface

You can set up a Single Firewall to receive its IPv4 address through DHCP. Only dynamic IPv4 addresses are supported. Typically, the dynamic address is assigned by the ISP's DHCP service. Modem Interfaces always have dynamic IP addresses that are provided through PPPoE (see Defining Modem Interfaces for Single Firewalls (page 448)).

Since the dynamic IP address assignment includes the next-hop gateway address, the routing for interfaces with a dynamic address is defined using special dynamic Router and NetLink elements. See Adding Routes for Firewalls (page 589).

If the address the engine uses for system communications is dynamic, the engine opens a communications channel to the Management Server and the Management Server never attempts to contact the engine outside this connection. If the management traffic flows through an interface that has a dynamic address, you can adjust timeouts and other settings related to these communications. See Adjusting the Global Contact Policy for Single Engines (page 401).

Dynamic IPv4 addresses also affect policy-based VPNs: other VPN gateways cannot open VPN connections to the gateway if the address is dynamic. Instead, the gateway with the dynamic end-point must always initiate the VPN, after which the VPN can be used normally in both directions.

There are default Alias elements in the system that can be used in the Firewall's policy to represent its own dynamic addresses. For each dynamic address interface there are four Alias elements distinguished by the DHCP index number:

- **$$ DHCP Interface X.ip**: the current dynamic IP address allocated to this interface.
- **$$ DHCP Interface X.gateways**: the received IP address for the default router.
- **$$ DHCP Interface X.dns**: the received IP address of the DNS server.
- **$$ DHCP Interface X.net**: the network behind the dynamic interface.

Note – These Aliases are meant for use in the policies of the Firewall that has the dynamic IP address. They are translated to the values of the Firewall the policy is installed on, so they cannot be used in the policies of other components.

IPS Engine Interface Configuration

Prerequisites: Creating New Engine Elements / Modifying Existing Engine Elements

The interface configuration for Single IPS engines and IPS Clusters consists of the following main steps:

1. Define the required number of Physical Interfaces as explained in Defining System Communication Interfaces for IPS Engines.
2. Add the required number of VLANs as explained in Adding VLAN Interfaces for IPS Engines (page 459).
3. Add IP addresses or a traffic inspection role to the interface as explained in Configuring IP Addresses for IPS Engines (page 461) or Defining Traffic Inspection Interfaces for IPS Engines (page 463).
   - IP addresses are required for interfaces that are used for system communications.
   - Interfaces that have a traffic inspection role work transparently in the network and do not have IP addresses.
4. Select the interfaces that are used for system communications as explained in Setting Interface Options for IPS Engines (page 471).

Limitations

- Only IPv4 addresses are supported for IPS engines.

What's Next?

▶ Start by Defining System Communication Interfaces for IPS Engines (page 457).

Related Tasks

▶ Configuring Manual ARP Settings (page 510)

▶ Firewall Interface Configuration (page 421)

▶ Layer 2 Firewall Interface Configuration (page 473)

Defining System Communication Interfaces for IPS Engines

Physical Interfaces correspond to network ports on the IPS engine. By default, the numbering of the Physical Interfaces in the Management Client corresponds to the operating system interface numbering on the engine (that is, Interface ID 0 is mapped to eth0, ID 1 to eth1, etc.). However, the mapping is not fixed and you can change it through the engine command line.

Each IPS engine needs at least one interface for communicating with other system components. You can define more than one system communication interface if it is necessary in your network environment.

For settings on the **Advanced** tab, see Configuring Advanced Interface Properties for IPS Engines (page 469).

▼ To define a system communication interface for IPS engines

1. Open the properties of the IPS engine and switch to the **Interfaces** tab.
2. Right-click the empty space and select **New Physical Interface**. The Physical Interface Properties dialog opens.

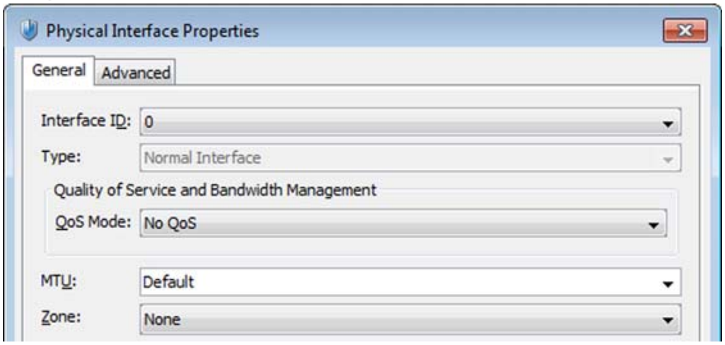

Physical Interface Properties

Interface ID: 0

Type: Normal Interface

Quality of Service and Bandwidth Management

QoS Mode: No QoS

MTU: Default

Zone: None

3. Define the Physical Interface properties as explained in the table below:

Table 28.8 Physical Interface Properties for IPS Engine System Communication

Option	Explanation
Interface ID	Select an Interface ID. The Interface ID is mapped to a physical interface during the initial configuration of the engine.
Zone (*optional*)	Select the network zone to which the Physical interface belongs from the list or select **Other** to select another Zone. If the Zone is not listed, create a new Zone element through the New icon at the top of the dialog. See Defining Zone Elements (page 730) for more information.
Type	Select **Normal Interface**.
QoS Mode (*Optional*)	Defines how QoS is applied to the link on this interface. If **Full QoS** or **DSCP Handling and Throttling** is selected as the QoS Mode, a QoS Policy must also be selected. If **Full QoS** is selected as the QoS Mode, the Throughput must also be defined.

Table 28.8 Physical Interface Properties for IPS Engine System Communication (Continued)

Option	Explanation
QoS Policy *(Full QoS or DSCP Handling and Throttling modes only)*	The QoS Policy for the link on this interface. For more information, see Getting Started with QoS (page 782).
Throughput *(Full QoS mode only)*	Enter the throughput for the link on this interface as kilobits per second (for example, 2048). The same throughput is automatically applied to any VLANs created under this Physical Interface. See Adding VLAN Interfaces for IPS Engines (page 459). The throughput is for uplink speed (outgoing traffic) and typically must correspond to the speed of an Internet link (such as an ADSL line), or the combined speeds of several such links when they are connected to a single physical interface). **Caution!** Make sure you set the Interface speed correctly. When the bandwidth is set, the Firewall always scales the total amount of traffic on this interface to the bandwidth you defined. This happens even if there are no bandwidth limits or guarantees defined for any traffic.
MTU *(Optional)*	The MTU (maximum transmission unit) size on the connected link. Either enter a value between 400-65535 or select a common MTU value from the list. The same MTU is automatically applied to any VLANs created under this Physical Interface. See Adding VLAN Interfaces for IPS Engines (page 459). The default value (also the maximum standard MTU in Ethernet) is 1500. Do not set a value larger than the standard MTU unless you know that all the devices along the communications path support it.

4. Click **OK**. The physical interface is added to the interface list.
5. Repeat from Step 2 to add more Physical Interfaces that you want to use for system communications.

What’s Next?

- If the system communications must use a particular VLAN on the directly connected network segment, proceed to Adding VLAN Interfaces for IPS Engines (page 459).
- Otherwise, proceed to Configuring IP Addresses for IPS Engines (page 461)

Related Tasks

- Configuring Advanced Interface Properties for IPS Engines (page 469)

Adding VLAN Interfaces for IPS Engines

VLANs divide a single physical network link into several virtual links. Up to 4094 VLANs can be configured for each interface.

Traffic picked up from a VLAN tagged interface can be inspected without configuring VLAN tagging on the IPS engine. However, you must configure the VLANs on the IPS engine if you want to create different traffic inspection rules for different VLANs. Even then, not all VLANs necessarily have to be specified on the IPS engine. VLANs can also optionally be used for sending the IPS engine's management and logging connections through a directly connected VLAN segment. The VLAN identifiers you configure on the IPS engine must match the switch/router configuration.

If the IPS engine encounters unknown VLANs, it may or may not inspect the traffic; this is controlled by the "Inspect Unspecified VLANs" option in the Inline and Capture Interface definitions (by default, the option is set so that all traffic is inspected).

> **Caution – Do not add any manual VLAN definitions to an interface you want to use for sending resets. Adding VLANs prevents selecting the interface as a Reset Interface and also removes the Reset Interface from any existing selections. The IPS engine automatically uses the correct VLAN when sending resets.**

▼ To define a VLAN Interface

1. Open the IPS engine element's properties and switch to the **Interfaces** tab.
2. Right-click a Physical Interface and select **New→VLAN Interface**. The VLAN Interface Properties dialog opens.
3. Define the VLAN interface properties as explained in the table below.

Table 28.9 VLAN Interface Properties

Option	Explanation
VLAN ID	Enter the VLAN ID (1-4094). The VLAN IDs you add must be the same as the VLAN IDs that are used in the switch at the other end of the VLAN trunk.
QoS Mode (*Optional*)	Defines how QoS is applied to the link on this interface. If **Full QoS** or **DSCP Handling and Throttling** is selected as the QoS Mode, a QoS Policy must also be selected. If **Full QoS** is selected as the QoS Mode, the Throughput must also be defined.
QoS Policy (*Full QoS or DSCP Handling and Throttling modes only*)	The QoS Policy for the link on the interface. For more information, see Getting Started with QoS (page 782).

Table 28.9 VLAN Interface Properties (Continued)

Option	Explanation
Throughput (*Full QoS mode only*)	Enter the throughput for the link on this interface as kilobits per second (for example, 2048). If throughput is defined for the Physical Interface to which the VLAN Interface belongs, the throughput value is automatically inherited from the Physical Interface properties. **Caution!** The throughput for each VLAN Interface must not be higher than the throughput for the Physical Interface to which the VLAN Interface belongs. The throughput is for uplink speed (outgoing traffic) and typically must correspond to the speed of an Internet link (such as an ADSL line), or the combined speeds of several such links when they are connected to a single Physical Interface. **Caution!** Make sure you set the interface speed correctly. When the bandwidth is set, the Firewall always scales the total amount of traffic on this interface to the bandwidth you defined. This happens even if there are no bandwidth limits or guarantees defined for any traffic.
MTU (*Optional*)	The MTU (maximum transmission unit) size on the connected link. Either enter a value between 400-65535 or select a common MTU value from the list. **Caution!** The MTU for each VLAN Interface must not be higher than the MTU for the Physical Interface to which the VLAN Interface belongs. The default value (also the maximum standard MTU in Ethernet) is 1500. Do not set a value larger than the standard MTU unless you know that all the devices along the communications path support it.
Zone (*Optional*)	Select the network zone to which the VLAN interface belongs from the list or select **Other** to select another Zone. If the Zone is not listed, create a new Zone element through the New icon at the top of the dialog. See Defining Zone Elements (page 730) for more information.

4. Click **OK**. The specified VLAN ID is added to the Physical Interface.
5. Repeat from Step 2 to add further VLANs (up to 4094 according to the configuration of the device at the other end of the link).

The VLAN Interface is now ready. The VLAN Interface is identified as `Interface-ID.VLAN-ID`, for example `2.100` for Interface ID 2 and VLAN ID 100.

What's Next?

- If you added a VLAN to a Physical Interface of the Normal type, continue by Configuring IP Addresses for IPS Engines (page 461)
- Otherwise, the new VLAN Interface definition is finished. Click **OK** and refresh the IPS engine's policy to activate the new interface configuration.

Related Tasks

- Configuring Advanced Interface Properties for IPS Engines (page 469)

Configuring IP Addresses for IPS Engines

Prerequisites: Defining System Communication Interfaces for IPS Engines (page 457)

An IPS engine's system communication interfaces (Normal interfaces) can have the following types of IP addresses:

- A Physical Interface can have one or more static or dynamic IP addresses. A Physical Interface can have multiple dynamic IP addresses only if you add VLAN Interfaces on the Physical Interface and the VLAN Interfaces each have a dynamic IP address. Otherwise, a Physical Interface can only have a single dynamic IP address.
- A VLAN Interface can have one or more static IP addresses or a single dynamic IP address.

You may need to define a *contact address* if you enter a private static address and NAT is used to translate it to a different external IP address. The external IP address must be configured as the contact address if other system components need to use the external IP address to contact the engine. Refer to Defining Contact IP Addresses (page 66) for more information.

What's Next?

▶ Configuring IP Addresses for Single IPS Engines

▶ Configuring IP Addresses for IPS Clusters (page 462)

Configuring IP Addresses for Single IPS Engines

Prerequisites:

▼ To add an IPv4 address for a Single IPS engine

1. In the properties dialog for the Single IPS engine, switch to the **Interfaces** tab.
2. Right-click a Physical Interface or VLAN Interface and select **New IPv4 Address**. The IP Address Properties dialog opens.

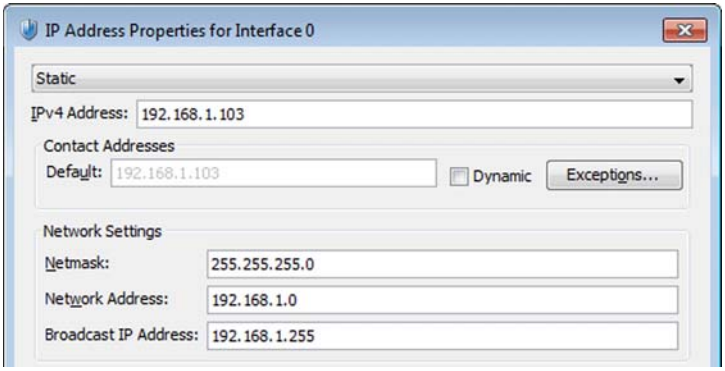

IP Address Properties for Interface 0

Static

IPv4 Address: 192.168.1.103

Contact Addresses

Default: 192.168.1.103 □ Dynamic Exceptions...

Network Settings

Netmask: 255.255.255.0

Network Address: 192.168.1.0

Broadcast IP Address: 192.168.1.255

3. Configure the Network Settings:
   - Enter the **IPv4 Address** and **Netmask** to define a static IP address.
   - Select the **Dynamic** option (top right) and the **DHCP index** if the interface gets its IP address from a DHCP server. The DHCP Index is a number for your own reference to identify the DHCP interface.

4. If necessary, define the contact address information (see Defining Contact Addresses for a Single Engine or a Cluster Virtual IP Address (page 68) for detailed instructions).
   - Enter the **Default** contact address or select **Dynamic** (next to the **Default** field) if the interface has a dynamic contact address. The Default contact address is used by default whenever a component that belongs to another Location connects to this interface.
   - If components from some Locations cannot use the Default contact address, click **Exceptions** to define Location-specific contact addresses.
5. Click **OK**.

What's Next?

- ▶ If you are defining a new IPS element, proceed to Defining Traffic Inspection Interfaces for IPS Engines (page 463).
- ▶ If you are editing an existing IPS element, proceed to Setting Interface Options for IPS Engines (page 471).

Configuring IP Addresses for IPS Clusters

▼ To add IPv4 addresses for an IPS Cluster

1. In the properties dialog for the IPS Cluster, switch to the **Interfaces** tab.
2. Right-click a Physical interface or VLAN Interface **New IPv4 Address**. The IP Address Properties dialog opens.

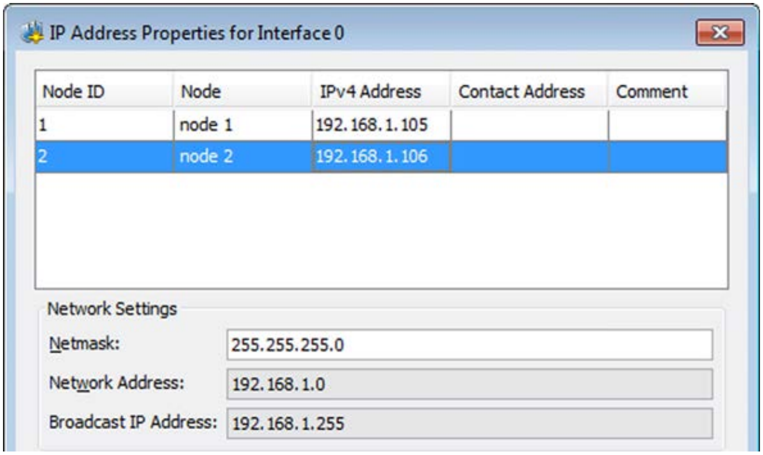

Node ID	Node	IPv4 Address	Contact Address	Comment
1	node 1	192.168.1.105		
2	node 2	192.168.1.106		

Network Settings

Netmask: 255.255.255.0

Network Address: 192.168.1.0

Broadcast IP Address: 192.168.1.255

3. Double-click the **IPv4 Address** cell in the table and enter the IP address for each node.
4. If necessary, define the contact address for each node by double-clicking the Contact Address cell. The Exceptions dialog opens (see Defining Contact Addresses for an IPS Cluster or a Layer 2 Firewall Cluster (page 70) for detailed instructions).
   - Enter the **Default** contact address at the top of the dialog. The Default contact address is used by default whenever a component that belongs to another Location connects to this interface.
   - If components from some Locations cannot use the Default contact address, click **Add** to define Location-specific contact addresses.
5. Click **OK** to close the IP Address Properties dialog.
6. Repeat the steps above as necessary to add more IP addresses to the same Physical Interface or VLAN Interface.

What's Next?

- ▶ If you are defining a new IPS element, proceed to Defining Traffic Inspection Interfaces for IPS Engines.
- ▶ If you are editing an existing IPS element, proceed to Setting Interface Options for IPS Engines (page 471).

Defining Traffic Inspection Interfaces for IPS Engines

IPS engines pick up traffic from the network for inspection. There are two ways to install the IPS engines:

- In an IDS-like configuration in which the traffic is only *captured* for inspection.
- In a full IPS configuration where the IPS engine is installed *inline*, directly on the traffic path so that traffic must always pass through the IPS engine to reach its destination. Only traffic that attempts to pass through Inline Interfaces can be actively filtered.

Connections picked up through *Capture Interfaces* can be reset through specially set up *Reset Interfaces*. Capture and *Inline Interfaces* can be defined on the same IPS engine and used simultaneously.

Logical Interface elements allow you to group interfaces together according to network segment and interface type. You can then use the Logical Interface elements as matching criteria when you edit the rules in your IPS policies.

What's Next?

- ▶ There is one predefined Logical Interface element (called default_eth) that can be used in interface configurations. If you want to create both Capture and Inline Interfaces on the same IPS engine, you must add at least one more Logical Interface: proceed to Defining Logical Interfaces for IPS Engines and Layer 2 Firewalls.
- ▶ To define a Capture Interface, you must have a Reset interface for it. See Defining Reset Interfaces for IPS Engines (page 464).
- ▶ Defining Capture Interfaces for IPS Engines (page 465).
- ▶ Defining Inline Interfaces for IPS Engines (page 466).

Defining Logical Interfaces for IPS Engines and Layer 2 Firewalls

A Logical Interface is used in IPS Policies and Layer 2 Firewall Policies and in the traffic inspection process to represent a network segment. Logical Interfaces help limit the scope of your rules. The same Logical Interface cannot be used to represent both Capture Interfaces and Inline Interfaces on the same engine. Otherwise, a Logical Interface can represent any combination of interfaces.

▼ To define a new Logical Interface element

1. Select **Configuration→Configuration→Security Engine**. The Security Engine Configuration view opens.
   - If you are currently creating a new IPS element or Layer 2 Firewall element, leave the element's Properties dialog open in the background.
2. Expand the **Other Elements** branch.

3. Right-click **Logical Interfaces** and select **New Logical Interface**. The Logical Interface Properties dialog opens.

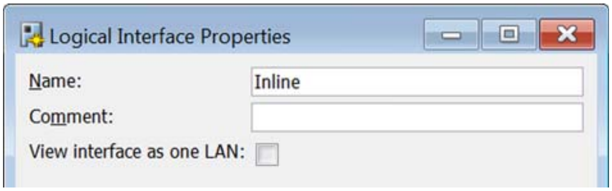

Name: Inline

Comment:

View interface as one LAN: □

4. Enter a **Name** and an optional **Comment** for the new element.
5. (*Optional*) Select **View interface as one LAN** to make the engine treat VLANs associated with the Logical Interface as a single LAN. This prevents the IPS engine or Layer 2 Firewall from seeing a single connection as multiple connections when a switch passes traffic between different VLANs or if all traffic is mirrored to the IPS engine through a SPAN port.
6. Click **OK**.

What's Next?

▶ If you are configuring a new Layer 2 Firewall element, continue by Defining Inline Interfaces for Layer 2 Firewalls (page 480).

▶ For IPS engines, proceed according to the type of interface you are configuring:

- Defining Reset Interfaces for IPS Engines
- Defining Capture Interfaces for IPS Engines (page 465)
- Defining Inline Interfaces for IPS Engines (page 466)

Defining Reset Interfaces for IPS Engines

Reset Interfaces can deliver TCP resets and ICMP "destination unreachable" messages to interrupt communications picked up through Capture Interfaces when the communications trigger a response.

The resets are sent using the source and destination addresses and MAC addresses of the communicating hosts, so an IP address is not mandatory for a Reset Interface. You can add an IP address if you also want to use this interface for system communications.

VLANs are supported for sending resets, but the correct VLAN is selected automatically and the interface you want to use as the Reset Interface must not have any manually added VLAN configuration.

You can also use a system communications interface for sending resets if the resets are routed through that interface (see Defining System Communication Interfaces for IPS Engines (page 457)) and if there are no VLANs on the interface.

▼ To add a Reset Interface

1. In the properties of the IPS engine, switch to the **Interfaces** tab.
2. Right-click the empty space and select **New Physical Interface**. The Physical Interface Properties dialog opens.
3. Select an **Interface ID**. The Interface ID is mapped to a physical interface during the initial configuration of the engine.
4. Select **Normal Interface** as the interface **Type**.
5. Click **OK**. The Physical Interface is added to the interface list.

When you set up the physical network, ensure that the Reset Interface connects to the same network(s) as the Capture interface(s).

What's Next?

▶ Set up the Capture Interfaces that use this Reset Interface as explained in Defining Capture Interfaces for IPS Engines.

Defining Capture Interfaces for IPS Engines

Capture Interfaces listen to traffic that is not routed through the IPS engine. You can have as many Capture Interfaces as there are available physical ports on the IPS engine. External equipment must be set up to mirror traffic to the Capture Interface.

For settings on the **Advanced** tab, see Configuring Advanced Interface Properties for IPS Engines (page 469).

▼ To define a Capture Interface

1. In the properties of the IPS engine element, switch to the **Interfaces** tab.
2. Right-click the empty space and select **New Physical Interface**. The Physical Interface Properties dialog opens.

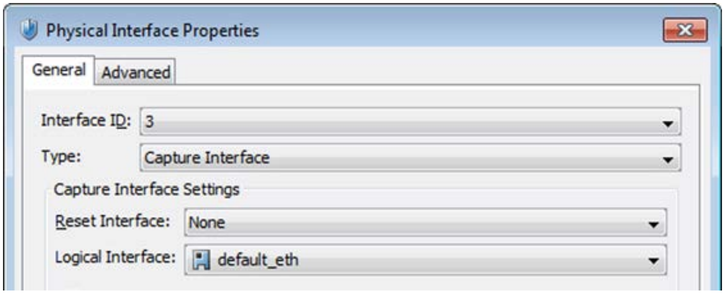

Interface ID: 3
Type: Capture Interface
Capture Interface Settings
Reset Interface: None
Logical Interface: default_eth

3. Define the interface properties as explained below.

Option	Explanation
Interface ID	Select an Interface ID. The Interface ID is mapped to a physical interface during the initial configuration of the engine.
Type	Select **Capture Interface** as the interface Type.
Reset Interface (*Optional*)	Select the **Reset Interface** to specify the interface through which TCP connection resets are sent when Reset responses are used in your IPS policy. For more information, see Defining Reset Interfaces for IPS Engines (page 464).
Logical Interface	Select the **Logical Interface**. For more information, see Defining Logical Interfaces for IPS Engines and Layer 2 Firewalls (page 463). You cannot use the same Logical Interface element for both Inline and Capture Interfaces on the same IPS engine.
Inspect Unspecified VLANs (*Optional*)	Deselect this option to make the IPS engine ignore traffic from VLANs that are not included in the IPS engine's interface configuration. We recommend that you keep this option selected if you do not have a specific reason to deselect it.

Option	Explanation
Zone *(Optional)*	Select the network zone to which the VLAN interface belongs from the list or select **Other** to select another Zone. If the Zone is not listed, create a new Zone element through the New icon at the top of the dialog. See Defining Zone Elements (page 730) for more information.

4. Click **OK** to save the interface settings.

What's Next?

- If you are creating a new element, proceed to Setting Interface Options for IPS Engines (page 471).
- If you added a new Interface to an existing element, click **OK** and refresh the IPS engine's policy to activate the new interface configuration.

Related Tasks

- Defining Inline Interfaces for IPS Engines
- Adding VLAN Interfaces for IPS Engines (page 459)
- Configuring Advanced Interface Properties for IPS Engines (page 469)

Defining Inline Interfaces for IPS Engines

Inline Interfaces allow traffic to flow through the IPS engine, so that traffic that is deemed harmful can be actively filtered out. The number of Inline Interfaces you can configure is limited by the IPS engine's license.

An Inline Interface differs from the other interfaces: it consists of two different Physical Interfaces. This way, the IPS engine can inspect the traffic coming from one interface and either stop the traffic or send it out through the other interface. The two interfaces are equal in the configuration and traffic that is allowed through is always forwarded from one interface to the other (there is no routing decision involved in this).

Note – Fail-open network cards have fixed pairs of ports. Take particular care to map these ports correctly. Otherwise, the network cards do not correctly fail open when the IPS engine is offline.

For settings on the **Advanced** tab, see Configuring Advanced Interface Properties for IPS Engines (page 469).

▼ To define an Inline Interface

1. In the properties of the IPS engine element, switch to the **Interfaces** tab.
2. Right-click the empty space and select **New Physical Interface**. The Physical Interface Properties dialog opens.

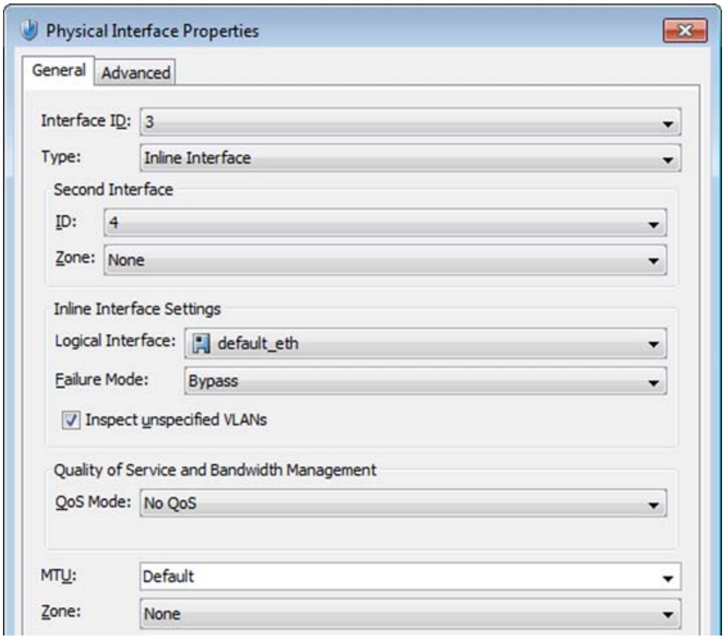

Interface ID: 3

Type: Inline Interface

Second Interface

ID: 4

Zone: None

Inline Interface Settings

Logical Interface: default_eth

Failure Mode: Bypass

☑ Inspect unspecified VLANs

Quality of Service and Bandwidth Management

QoS Mode: No QoS

MTU: Default

Zone: None

3. Define the interface properties as explained in the table below.

Table 28.10 Inline Interface Properties

Option	Explanation
Interface ID	Select an Interface ID. The Interface ID is mapped to a physical interface during the initial configuration of the engine.
Type	Select **Inline Interface** as the interface Type.
Second Interface ID	Select a Second Interface ID. The Interface ID is mapped to a physical interface during the initial configuration of the engine.
Zone (*Optional*)	Select the network **Zone** to which the interface belongs from the list or select **Other** to select another Zone. If the Zone is not listed, create a new Zone element through the New icon at the top of the dialog. See Defining Zone Elements (page 730) for more information.
Logical Interface	Select the Logical Interface. For more information, see Defining Logical Interfaces for IPS Engines and Layer 2 Firewalls (page 463). You cannot use the same Logical Interface element for both Inline and Capture Interfaces on the same IPS engine.
Failure Mode	Select how traffic to the inline interface is handled if the IPS engine goes offline. There are two options: **Bypass**: traffic is allowed through the inline interface without inspection. **Normal**: traffic is not allowed through the inline interface. **Caution!** Using the Bypass mode requires the IPS appliance to have a fail-open network interface card. If the ports that represent the pair of inline interface on the appliance cannot fail open, the policy installation fails on the IPS engine.

Table 28.10 Inline Interface Properties (Continued)

Option	Explanation
Inspect Unspecified VLANs (*Optional*)	Deselect this option to make the IPS engine ignore traffic from VLANs that are not included in the IPS engine's interface configuration. We recommend that you keep this option selected if you do not have a specific reason to deselect it.
QoS Mode (*Optional*)	Defines how QoS is applied to the link on this interface. If **Full QoS** or **DSCP Handling and Throttling** is selected as the QoS Mode, a QoS Policy must also be selected. If **Full QoS** is selected as the QoS Mode, the Throughput must also be defined.
QoS Policy (*Full QoS or DSCP Handling and Throttling modes only*)	Select the QoS Policy for the link on the interface. For more information, see Getting Started with QoS (page 782).
Throughput (*Full QoS mode only*)	Enter the throughput for the link on this interface as kilobits per second (for example, 2048). The throughput is for uplink speed (outgoing traffic) and typically must correspond to the speed of an Internet link (such as an ADSL line), or the combined speeds of several such links when they are connected to a single physical interface. **Caution!** Make sure you set the Interface speed correctly. When the bandwidth is set, the engine always scales the total amount of traffic on this interface to the bandwidth you defined. This happens even if there are no bandwidth limits or guarantees defined for any traffic
MTU (*Optional*)	The MTU (maximum transmission unit) size on the connected link. Either enter a value between 400-65535 or select a common MTU value from the list. The default value (also the maximum standard MTU in Ethernet) is 1500. Do not set a value larger than the standard MTU unless you know that all the devices along the communications path support it.
Zone (*Optional*)	Select the network **Zone** to which the interface belongs from the list or select **Other** to select another Zone. If the Zone is not listed, create a new Zone element through the New icon at the top of the dialog. See Defining Zone Elements (page 730) for more information.

4. Click **OK** to save the interface settings.

What's Next?

- Adding VLAN Interfaces for IPS Engines (page 459)
- If you are creating a new element, proceed to Setting Interface Options for IPS Engines (page 471).
- If you added a new Interface to an existing element, click **OK** and refresh the IPS engine's policy to activate the new interface configuration.

Related Tasks

- Configuring Advanced Interface Properties for IPS Engines

Configuring Advanced Interface Properties for IPS Engines

Advanced settings allow you to configure SYN Rate Limits and Log Compression on the IPS engine's interfaces. Log Compression is useful when the number of Discard logs becomes high (for example, as a result of a SYN flood attack).

Note – The SYN Rate Limits and Log Compression settings in the interface properties override the IPS engine's general SYN Rate Limits and Log Compression settings that are defined on the Advanced tab in the Single IPS and IPS Cluster properties. See Configuring Default SYN Rate Limits (page 579) and Configuring Default Log Handling Settings (page 580).

▼ To configure advanced interface properties for IPS engines

1. In the properties dialog for the IPS engine, switch to the **Interfaces** tab.
2. Right-click a Physical Interface or a VLAN and select **Edit Physical Interface** or **Edit VLAN Interface**. The properties dialog for the interface opens.
3. Switch to the **Advanced** tab.

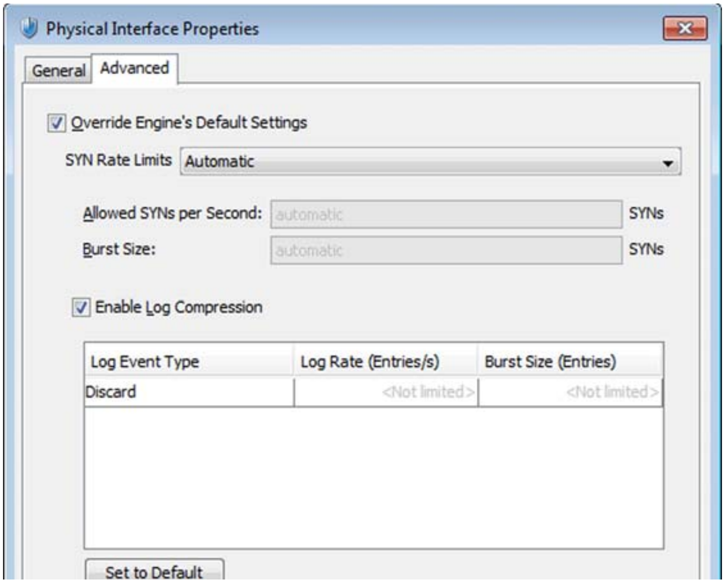

☑ Override Engine's Default Settings

SYN Rate Limits: Automatic

Allowed SYNs per Second: automatic SYNs

Burst Size: automatic SYNs

☑ Enable Log Compression

Log Event Type	Log Rate (Entries/s)	Burst Size (Entries)
Discard	<Not limited>	<Not limited>

4. Select **Override Engine's Default Settings**. The options for SYN Rate Limits and Log Compression are enabled.

5. (*Optional*) Define the **SYN Rate Limits**

Table 28.11 SYN Rate Limits

Setting	Description
Default	The interface uses the SYN Rate Limits defined on the Advanced tab in the engine properties (see Configuring Default SYN Rate Limits (page 579).
Off	Disables SYN Rate Limits on the interface.
Automatic	This is the recommended mode if you want to override the general SYN Rate Limits defined on the Advanced tab in the engine's properties. The engine automatically calculates the number of **Allowed SYNs per Second** (the number of allowed SYN packets per second) and the **Burst Size** (the number of allowed SYNs before the engine starts limiting the SYN rate) for the interface based on the engine's capacity and memory size.
Custom	Enter the desired values for **Allowed SYNs per Second** and **Burst Size**. We recommend that the Burst Size be at least one tenth of the Allowed SYNs per Second value. If the Burst Size is too small, SYN Rate Limits do not work. For example, if the value for Allowed SYNs per Second is 10000, the Burst Size should be at least 1000.

Caution – The recommended values for the SYN Flood Settings depend on your network environment. If the Custom settings are not carefully configured, the capacity of the engine may suffer or SYN Rate Limits may not work correctly.

6. (*Optional*) Enable/Disable **Log Compression** and enter the desired values for the Discard entries.

Setting	Description
Log Rate (Entries/s)	The maximum number of entries per second. By default, Discard log entries are not compressed.
Burst Size (Entries)	The maximum number of matching entries in a single burst. The default value for antispoofing entries is 1000 entries. By default, Discard log entries are not compressed.

- Do not enable Log Compression if you want all the Discard entries to be logged as separate log entries (for example, for reporting purposes or for engine statistics).
- By default, each generated Discard log entry is logged separately and displayed as a separate entry in the Logs view. Log Compression settings allow you to define the maximum number of separately logged entries. When the defined limit is reached, a single Discard log entry is logged. The single entry contains information on the total number of the generated Discard log entries. After this, the logging returns to normal and all the generated entries are once more logged and displayed separately.

7. Click **OK**.

What's Next?

- Close the IPS engine's Properties dialog and refresh the engine's policy to transfer the configuration changes.

Setting Interface Options for IPS Engines

Interface options allow you to select which interfaces are used for which types of system communications.

▼ To set the interface options

1. Open the IPS engine's properties and switch to the **Interfaces** tab.
2. Click the **Options** button below the interface table. The Interface Options dialog opens.

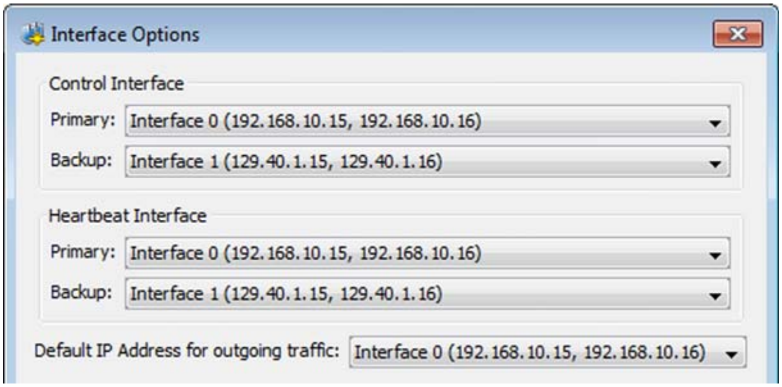

Interface Options

Control Interface

Primary: Interface 0 (192.168.10.15, 192.168.10.16)

Backup: Interface 1 (129.40.1.15, 129.40.1.16)

Heartbeat Interface

Primary: Interface 0 (192.168.10.15, 192.168.10.16)

Backup: Interface 1 (129.40.1.15, 129.40.1.16)

Default IP Address for outgoing traffic: Interface 0 (192.168.10.15, 192.168.10.16)

3. Select the interface options as explained in the table below.

Table 28.12 IPS Engine Interface Options

Option	Explanation
Control Interface	Select the **Primary** Control Interface for Management Server contact. This interface is used for communications with the Management Server
	(*Optional*) Select a **Backup** Control Interface that is used if the Primary interface is not available.
Node-Initiated contact to Management Server (*Single IPS engines only*)	Select the Node-Initiated contact to Management Server option if the engine is behind a device that applies dynamic NAT to the inbound management connections or blocks them. This option is always used with a dynamic control IP address, so it is always fixed to the selected state if the IP address is dynamic. When this option is selected, the engine opens a connection to the Management Server and maintains connectivity. If the connection is not open when you command the engine through the Management Client, the command is left pending until the engine opens the connection again.

Table 28.12 IPS Engine Interface Options (Continued)

Option	Explanation
Heartbeat Interface *(IPS Clusters only)*	Select the **Primary** Heartbeat Interface for communications between the nodes in the IPS cluster. This must not be a VLAN Interface. A dedicated network (without other traffic) is strongly recommended for security and reliability of heartbeat communications. In an IPS cluster with only two nodes, a crossover cable without additional intermediary devices is recommended for the heartbeat link.
	Select a **Backup** Heartbeat Interface that is used if the Primary Heartbeat Interface is unavailable. It is not mandatory to configure a backup heartbeat, but we strongly recommend it. If heartbeat traffic is not delivered, the cluster cannot operate and traffic will be disturbed. We strongly recommend that you use a dedicated interface for the backup heartbeat as well.
Default IP Address for Outgoing Traffic (*IPS Clusters only*)	Select the IP address that the nodes use if they have to initiate connections (system communications, ping, etc.) through an interface that has no Node Dedicated IP Address. You must select an interface that has an IP address defined for all nodes.

4. Click **OK**.

What's Next?

- If you are creating a new element, configure the routing as explained in Adding Routes for IPS Engines and Layer 2 Firewalls (page 605), and connect the engine to the Management Server as explained in Saving an Initial Configuration for Security Engines (page 515).
- Otherwise, click **OK** and refresh the IPS engine's policy to activate the new interface configuration.

Layer 2 Firewall Interface Configuration

Prerequisites: Creating and Modifying Engine Elements /Network Interface Configuration

The interface configuration for Single Layer 2 Firewalls and Layer 2 Firewall Clusters consists of the following main steps:

1. Define the required number of Physical Interfaces as explained in Defining System Communication Interfaces for Layer Firewalls.
2. Add the required number of VLANs as explained in Configuring VLAN Interfaces for Layer 2 Firewalls (page 475).
3. Add IP addresses or a traffic inspection role to the interface as explained in Configuring IP Addresses for Layer 2 Firewalls (page 477) or Defining Traffic Inspection Interfaces for Layer 2 Firewalls (page 479).
   - IP addresses are required for interfaces that are used for system communications.
   - Interfaces that have a traffic inspection role work transparently in the network and do not have IP addresses.
4. Select the interfaces that are used for system communications as explained in Setting Interface Options for Layer 2 Firewalls (page 484).

Limitations

Only IPv4 addresses are supported for Layer 2 Firewalls.

What's Next?

▶ Start by Defining System Communication Interfaces for Layer Firewalls.

Related Tasks

▶ Firewall Interface Configuration (page 421)

▶ IPS Engine Interface Configuration (page 456)

▶ Configuring Manual ARP Settings (page 510)

Defining System Communication Interfaces for Layer 2 Firewalls

A *Physical Interface* element corresponds to a network port on the Layer 2 Firewall. By default, the numbering of the Physical Interfaces in the Management Client corresponds to the operating system interface numbering on the engine (that is, Interface ID 0 is mapped to eth0, ID 1 to eth1, etc.). However, the mapping is not fixed and you can change it through the engine command line.

Each Layer 2 Firewall needs at least one interface for communicating with other system components. You can define more than one system communication interface if it is necessary in your network environment.

▼ To define a system communication interface for Layer 2 Firewalls

1. Open the properties of the Layer 2 Firewall and switch to the **Interfaces** tab.

2. Right-click the empty space and select **New Physical Interface**. The Physical Interface Properties dialog opens.

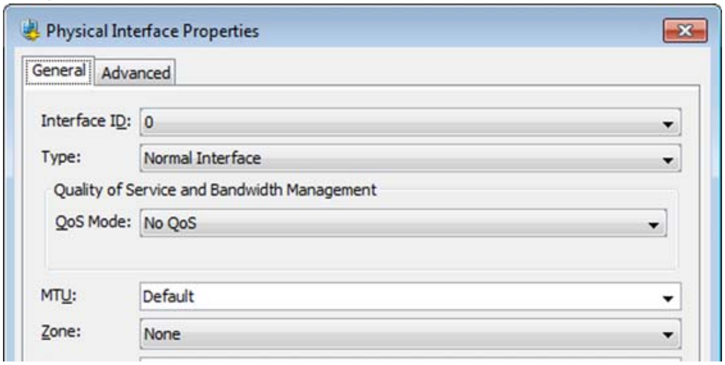

Physical Interface Properties

Interface ID: 0

Type: Normal Interface

Quality of Service and Bandwidth Management

QoS Mode: No QoS

MTU: Default

Zone: None

3. Define the Physical Interface properties as explained in the table below:

Table 28.13 Physical Interface Properties for Layer 2 Firewall System Communication

Option	Explanation
Interface ID	Select an Interface ID. The Interface ID is mapped to a physical interface during the initial configuration of the engine.
Type	Select **Normal Interface**.
QoS Mode (*Optional*)	Defines how QoS is applied to the link on this interface. If **Full QoS** or **DSCP Handling and Throttling** is selected as the QoS Mode, a QoS Policy must also be selected. If **Full QoS** is selected as the QoS Mode, the Throughput must also be defined.
QoS Policy (*Full QoS or DSCP Handling and Throttling modes only*)	The QoS Policy for the link on this interface. For more information, see Getting Started with QoS (page 782).
Throughput (*Full QoS mode only)*	Enter the throughput for the link on this interface as kilobits per second (for example, 2048). The same throughput is automatically applied to any VLANs created under this Physical Interface. See Configuring VLAN Interfaces for Layer 2 Firewalls (page 475). The throughput is for uplink speed (outgoing traffic) and typically must correspond to the speed of an Internet link (such as an ADSL line), or the combined speeds of several such links when they are connected to a single physical interface). **Caution!** Make sure you set the Interface speed correctly. When the bandwidth is set, the Firewall always scales the total amount of traffic on this interface to the bandwidth you defined. This happens even if there are no bandwidth limits or guarantees defined for any traffic.

Table 28.13 Physical Interface Properties for Layer 2 Firewall System Communication (Continued)

Option	Explanation
MTU *(Optional)*	The MTU (maximum transmission unit) size on the connected link. Either enter a value between 400-65535 or select a common MTU value from the list. The same MTU is automatically applied to any VLANs created under this Physical Interface. See Configuring VLAN Interfaces for Layer 2 Firewalls (page 475). The default value (also the maximum standard MTU in Ethernet) is 1500. Do not set a value larger than the standard MTU unless you know that all the devices along the communications path support it.
Zone (*optional*)	Select the network zone to which the Physical interface belongs from the list or select **Other** to select another Zone. If the Zone is not listed, create a new Zone element through the New icon at the top of the dialog. See Defining Zone Elements (page 730) for more information.

4. Click **OK**. The physical interface is added to the interface list.
5. Repeat from Step 2 to add more Physical Interfaces that you want to use for system communications.

What's Next?

- If the system communications must use a particular VLAN on the directly connected network segment, proceed to Configuring VLAN Interfaces for Layer 2 Firewalls.
- Configuring IP Addresses for Layer 2 Firewalls (page 477)

Related Tasks

- Configuring Advanced Interface Properties for Layer 2 Firewalls (page 482)

Configuring VLAN Interfaces for Layer 2 Firewalls

VLANs divide a single physical network link into several virtual links. VLANs can be defined for both single and clustered Layer 2 Firewalls. The maximum number of VLANs for a single Physical Interface or Inline Interface is 4094. The VLANs must also be defined in the configuration of the switch/router to which the interface is connected.

Traffic picked up from a VLAN tagged interface can be inspected without configuring VLAN tagging on the Layer 2 Firewall. However, you must configure VLANs on the Layer 2 Firewall if you want to create different traffic inspection rules for different VLANs. Even then, not all VLANs necessarily have to be specified on the Layer 2 Firewall. VLANs can optionally also be used for sending the Layer 2 Firewall's management and logging connections through a directly connected VLAN segment.

If the Layer 2 Firewall encounters unknown VLANs, it may or may not inspect the traffic; this is controlled by the "Inspect Unspecified VLANs" option in the Inline Interface definitions (by default, the option is set so that all traffic is inspected). See Defining Traffic Inspection Interfaces for Layer 2 Firewalls for more information.

▼ To add a VLAN Interface for a single or clustered Layer 2 Firewall

1. In the properties dialog for the Layer 2 Firewall, switch to the **Interfaces** tab.
2. Right-click a Physical Interface or Inline Interface and select **New→VLAN Interface**. The VLAN Interface Properties dialog opens.
3. Define the VLAN interface properties as explained in the table below.

Table 28.14 Layer 2 Firewall VLAN Interface Properties

Option	Explanation
VLAN ID	Enter the VLAN ID (1-4094). The VLAN IDs you add must be the same as the VLAN IDs that are used in the switch at the other end of the VLAN trunk.
QoS Mode (*Optional*)	Defines how QoS is applied to the link on this interface. If **Full QoS** or **DSCP Handling and Throttling** is selected as the QoS Mode, a QoS Policy must also be selected. If **Full QoS** is selected as the QoS Mode, the Throughput must also be defined.
QoS Policy (*Full QoS or DSCP Handling and Throttling modes only*)	The QoS Policy for the link on the interface. For more information, see Getting Started with QoS (page 782).
Throughput (*Full QoS mode only*)	The throughput for the link on this interface as kilobits per second (for example, 2048). If throughput is defined for the Physical Interface to which the VLAN Interface belongs, the throughput value is automatically inherited from the Physical Interface properties. **Caution!** The throughput for each VLAN Interface must not be higher than the throughput for the Physical Interface to which the VLAN Interface belongs. The throughput is for uplink speed (outgoing traffic) and typically must correspond to the speed of an Internet link (such as an ADSL line), or the combined speeds of several such links when they are connected to a single Physical Interface. **Caution!** Make sure you set the interface speed correctly. When the bandwidth is set, the Firewall always scales the total amount of traffic on this interface to the bandwidth you defined. This happens even if there are no bandwidth limits or guarantees defined for any traffic.
MTU (*Optional*)	The MTU (maximum transmission unit) size on the connected link. Either enter a value between 400-65535 or select a common MTU value from the list. **Caution!** The MTU for each VLAN Interface must not be higher than the MTU for the Physical Interface to which the VLAN Interface belongs. The default value (also the maximum standard MTU in Ethernet) is 1500. Do not set a value larger than the standard MTU unless you know that all the devices along the communications path support it.
Zone (*Optional*)	Select the network zone to which the VLAN interface belongs from the list or select **Other** to select another Zone. If the Zone is not listed, create a new Zone element through the New icon at the top of the dialog. See Defining Zone Elements (page 730) for more information.

4. Click **OK**. The specified VLAN ID is added to the Physical Interface.

5. (*Optional*) Repeat from Step 2 to add further VLANs (up to 4094 according to the configuration of the device at the other end of the link).

What's Next?

- If you added a VLAN to a Normal interface, proceed to Configuring IP Addresses for Layer 2 Firewalls.
- If you added a VLAN to an Inline Interface, the new Inline Interface definition is finished. Click **OK** and refresh the IPS engine's policy to activate the new interface configuration.

Related Tasks

- Configuring Advanced Interface Properties for Layer 2 Firewalls (page 482)

Configuring IP Addresses for Layer 2 Firewalls

A Layer 2 Firewall's system communication interfaces (Normal interfaces) can have the following types of IP addresses:

- A Physical Interface can have one or more static or dynamic IP addresses. A Physical Interface can have multiple dynamic IP addresses only if you add VLAN Interfaces on the Physical Interface and the VLAN Interfaces have dynamic IP addresses. Otherwise, a Physical Interface can only have a single dynamic IP address.
- A VLAN Interface can have one or more static IP addresses or a single dynamic IP address.

You may need to define a *contact address* if you enter a private static address and NAT is used to translate it to a different external IP address. The external IP address must be configured as the contact address if other system components need to use the external IP address to contact the engine. Refer to Defining Contact IP Addresses (page 66) for more information.

What's Next?

- Configuring IP Addresses for Single Layer 2 Firewalls (page 478)
- Configuring IP Addresses for Layer 2 Firewall Clusters (page 479)

Configuring IP Addresses for Single Layer 2 Firewalls

▼ To add an IPv4 address for a Single Layer 2 Firewall

1. In the properties dialog for the Layer 2 Firewall, switch to the **Interfaces** tab.
2. Right-click a Physical Interface or VLAN Interface and select **New IPv4 Address**. The IP Address Properties dialog opens.

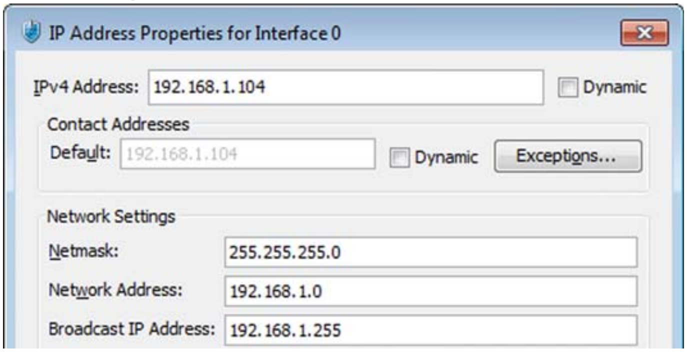

IP Address Properties for Interface 0

IPv4 Address: 192.168.1.104 □ Dynamic

Contact Addresses

Default: 192.168.1.104 □ Dynamic Exceptions...

Network Settings

Netmask: 255.255.255.0

Network Address: 192.168.1.0

Broadcast IP Address: 192.168.1.255

3. Configure the IP address information:
   - Enter the **IPv4 Address** and **Netmask** to define a static IP address.
   - Select the **Dynamic** option (top right) and the **DHCP index** if the interface gets its IP address from a DHCP server. The DHCP Index is a number for your own reference to identify the DHCP interface.
4. If necessary, define the contact address information (see Defining Contact Addresses for a Single Engine or a Cluster Virtual IP Address (page 68) for detailed instructions).
   - Enter the **Default** contact address or select **Dynamic** (next to the **Default** field) if the interface has a dynamic contact address. The Default contact address is used by default whenever a component that belongs to another Location connects to this interface.
   - If components from some Locations cannot use the Default contact address, click **Exceptions** to define Location-specific contact addresses.
5. Click **OK**.

What's Next?

▶ If you are defining a new Layer 2 Firewall element, proceed to Defining Traffic Inspection Interfaces for Layer 2 Firewalls (page 479).

▶ If you are editing an existing Layer 2 Firewall element and do not want to define traffic inspection interfaces, proceed to Setting Interface Options for Layer 2 Firewalls (page 484).

Configuring IP Addresses for Layer 2 Firewall Clusters

▼ To add IPv4 addresses for a Layer 2 Firewall Cluster

1. In the properties dialog for the Layer 2 Firewall Cluster, switch to the **Interfaces** tab.
2. Right-click a Physical interface or VLAN Interface and select **New IPv4 Address**. The IP Address Properties dialog opens.

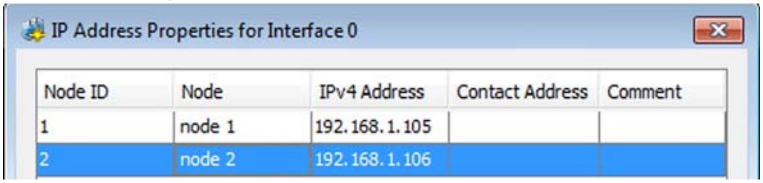

IP Address Properties for Interface 0

Node ID	Node	IPv4 Address	Contact Address	Comment
1	node 1	192.168.1.105		
2	node 2	192.168.1.106		

3. Double-click the **IPv4 Address** cell in the table and enter the IP address for each node.
4. If necessary, define the contact address for each node by double-clicking the Contact Address cell. The Exceptions dialog opens (see Defining Contact Addresses for an IPS Cluster or a Layer 2 Firewall Cluster (page 70) for detailed instructions).
   - Enter the **Default** contact address at the top of the dialog. The Default contact address is used by default whenever a component that belongs to another Location connects to this interface.
   - If components from some Locations cannot use the Default contact address, click **Add** to define Location-specific contact addresses.
5. Click **OK** to close the IP Address Properties dialog.
6. Repeat the steps above as necessary to add more IP addresses to the same Physical Interface or VLAN Interface.

What's Next?

- ▶ If you are defining a new Layer 2 Firewall Cluster element, proceed to Defining Traffic Inspection Interfaces for Layer 2 Firewalls (page 479).
- ▶ If you are editing an existing Layer 2 Firewall Cluster element and do not want to define traffic inspection interfaces, proceed to Setting Interface Options for Layer 2 Firewalls (page 484).

Defining Traffic Inspection Interfaces for Layer 2 Firewalls

Layer 2 Firewalls inspect network traffic. Layer 2 Firewalls are installed *inline*, directly on the traffic path so that traffic must always pass through the Layer 2 Firewall engine to reach its destination. Only traffic that attempts to pass through *Inline Interfaces* can be actively filtered.

Logical Interface elements allow you to group interfaces together according to network segment. You can then use the Logical Interface elements as matching criteria when you edit the rules in your Layer 2 Firewall policies.

What's Next?

- ▶ There is one predefined Logical Interface element (called default_eth) that can be used in interface configurations. If you need to create more Logical Interfaces, configure the Logical Interfces as described above in Defining Logical Interfaces for IPS Engines and Layer 2 Firewalls (page 463).
- ▶ Defining Inline Interfaces for Layer 2 Firewalls (page 480).

Defining Inline Interfaces for Layer 2 Firewalls

Inline Interfaces allow traffic to flow through the Layer 2 Firewall, so that traffic that is deemed harmful can be actively filtered out. The number of Inline Interfaces you can configure is limited by the Layer 2 Firewall's license.

An Inline Interface consists of two different Physical Interfaces. This way, the Layer 2 Firewall can inspect the traffic coming from one interface and either stop the traffic or send it out through the other interface. The two interfaces are equal in the configuration and traffic that is allowed through is always forwarded from one interface to the other (there is no routing decision involved in this).

▼ To define an Inline Interface

1. In the properties of the Layer 2 Firewall element, switch to the **Interfaces** tab.
2. Right-click the empty space and select **New Physical Interface**. The Physical Interface Properties dialog opens.

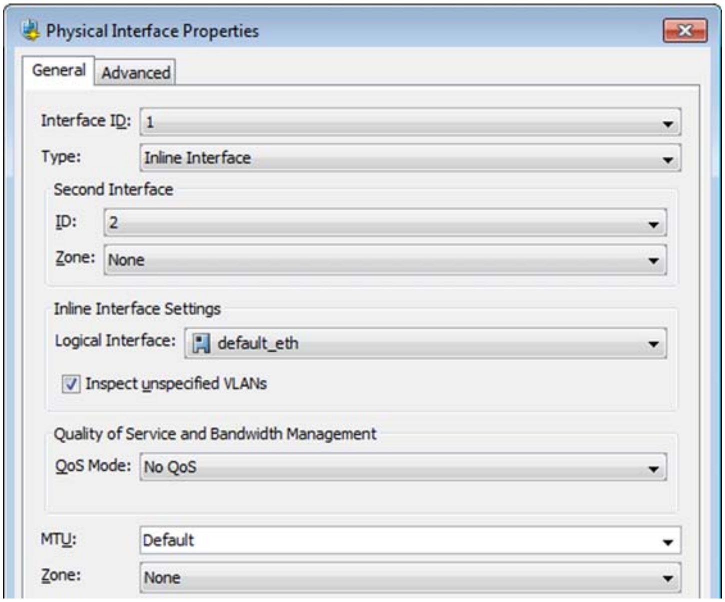

Interface ID: 1
Type: Inline Interface
Second Interface
ID: 2
Zone: None
Inline Interface Settings
Logical Interface: default_eth
☑ Inspect unspecified VLANs
Quality of Service and Bandwidth Management
QoS Mode: No QoS
MTU: Default
Zone: None

3. Define the interface properties as explained in the table below.

Table 28.15 Inline Interface Properties

Option	Explanation
Interface ID	Select an Interface ID. The Interface ID is mapped to a physical interface during the initial configuration of the engine.
Type	Select **Inline Interface** as the interface Type.
Second Interface ID	Select a Second Interface ID. The Interface ID is mapped to a physical interface during the initial configuration of the engine.

Table 28.15 Inline Interface Properties (Continued)

Option	Explanation
Zone (*Optional*)	Select the network **Zone** to which the interface belongs from the list or select **Other** to select another Zone. If the Zone is not listed, create a new Zone element through the New icon at the top of the dialog. See Defining Zone Elements (page 730) for more information.
Logical Interface	Select the Logical Interface. For more information, see Defining Logical Interfaces for IPS Engines and Layer 2 Firewalls (page 463).
Inspect Unspecified VLANs (*Optional*)	Deselect this option to make the Layer 2 Firewall ignore traffic from VLANs that are not included in the Layer 2 Firewall's interface configuration. We recommend that you keep this option selected if you do not have a specific reason to deselect it.
QoS Mode (*Optional*)	Defines how QoS is applied to the link on this interface. If **Full QoS** or **DSCP Handling and Throttling** is selected as the QoS Mode, a QoS Policy must also be selected. If **Full QoS** is selected as the QoS Mode, the Throughput must also be defined.
QoS Policy (*Full QoS or DSCP Handling and Throttling modes only*)	Select the QoS Policy for the link on the interface. For more information, see Getting Started with QoS (page 782).
Throughput (*Full QoS mode only*)	Enter the Throughput for the link on this interface as kilobits per second (for example, 2048). The throughput is for uplink speed (outgoing traffic) and typically must correspond to the speed of an Internet link (such as an ADSL line), or the combined speeds of several such links when they are connected to a single physical interface. **Caution!** Make sure you set the Interface speed correctly. When the bandwidth is set, the engine always scales the total amount of traffic on this interface to the bandwidth you defined. This happens even if there are no bandwidth limits or guarantees defined for any traffic
Zone (*Optional*)	Select the network **Zone** to which the interface belongs from the list or select **Other** to select another Zone. If the Zone is not listed, create a new Zone element through the New icon at the top of the dialog. See Defining Zone Elements (page 730) for more information.
MTU (*Optional*)	The MTU (maximum transmission unit) size on the connected link. Either enter a value between 400-65535 or select a common MTU value from the list. The default value (also the maximum standard MTU in Ethernet) is 1500. Do not set a value larger than the standard MTU unless you know that all the devices along the communications path support it.

4. Click **OK** to save the interface settings.

What's Next?

- ▶ If you want to define a VLAN Interface for the Inline Interface, add the VLAN Interface as described above in Configuring VLAN Interfaces for Layer 2 Firewalls (page 475).
- ▶ If you are creating a new Layer 2 Firewall element, proceed to Setting Interface Options for IPS Engines (page 471).
- ▶ If you added a new Interface to an existing element, click **OK** and refresh the IPS engine's policy to activate the new interface configuration.

Configuring Advanced Interface Properties for Layer 2 Firewalls

Advanced settings allow you to configure SYN Rate Limits and Log Compression on the Layer 2 Firewall's interfaces. Log Compression is useful when the number of Discard logs becomes high (for example, as a result of a SYN flood attack).

Note – The SYN Rate Limits and Log Compression settings in the interface properties override the Layer 2 Firewall's general SYN Rate Limits and Log Compression settings that are defined on the Advanced tab in the Single Layer 2 Firewall and Layer 2 Firewall Cluster properties. See Configuring Default SYN Rate Limits (page 579) and Configuring Default Log Handling Settings (page 580).

▼ To configure advanced interface properties for Layer 2 Firewalls

1. In the properties dialog for the Layer 2 Firewall, switch to the **Interfaces** tab.
2. Right-click a Physical Interface or VLAN and select **Edit Physical Interface** or **Edit VLAN Interface**. The properties dialog for the interface opens.
3. Switch to the **Advanced** tab.

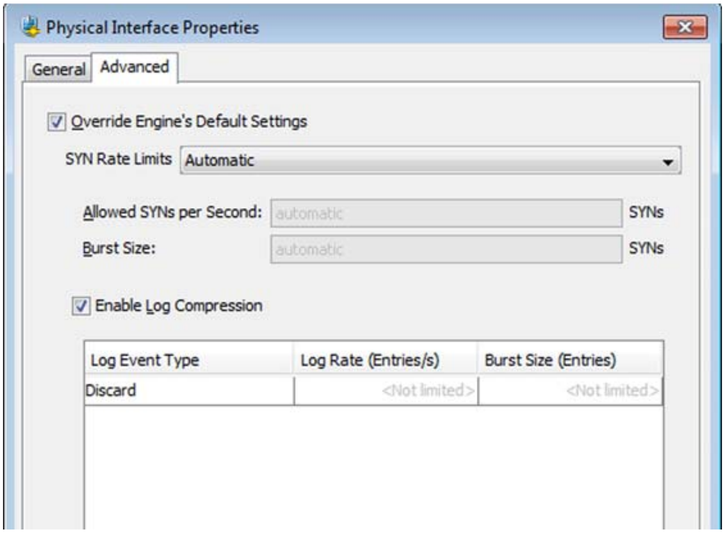

☑ Override Engine's Default Settings

SYN Rate Limits: Automatic

Allowed SYNs per Second: automatic SYNs

Burst Size: automatic SYNs

☑ Enable Log Compression

Log Event Type	Log Rate (Entries/s)	Burst Size (Entries)
Discard	<Not limited>	<Not limited>

4. Select **Override Engine's Default Settings**. The options for SYN Rate Limits and Log Compression are enabled.

5. (*Optional*) Define the **SYN Rate Limits**

Table 28.16 Layer 2 Firewall SYN Rate Limits Settings

Setting	Description
Default	The interface uses the SYN Rate Limits defined on the Advanced tab in the engine properties. See Configuring Default SYN Rate Limits (page 579).
Off	Disables SYN Rate Limits on the interface.
Automatic	This is the recommended mode if you want to override the general SYN Rate Limits defined on the Advanced tab in the engine's properties. The engine automatically calculates the number of **Allowed SYNs per Second** (the number of allowed SYN packets per second) and the **Burst Size** (the number of allowed SYNs before the engine starts limiting the SYN rate) for the interface based on the engine's capacity and memory size.
Custom	Enter the desired values for **Allowed SYNs per Second** and **Burst Size**. We recommend that the Burst Size be at least one tenth of the Allowed SYNs per Second value. If the Burst Size is too small, SYN Rate Limits do not work. For example, if the value for Allowed SYNs per Second is 10000, the Burst Size should be at least 1000.

Caution – The recommended values for the SYN Flood Settings depend on your network environment. If the Custom settings are not carefully configured, the capacity of the engine may suffer or SYN Rate Limits may not work correctly.

6. (*Optional*) Enable/Disable **Log Compression** and enter the desired values for the Discard entries.

Table 28.17 Log Compression Settings

Setting	Description
Log Rate (Entries/s)	The maximum number of entries per second. By default, Discard log entries are not compressed.
Burst Size (Entries)	The maximum number of matching entries in a single burst. The default value for antispoofing entries is 1000 entries. By default, Discard log entries are not compressed.

- Do not enable Log Compression if you want all the Discard entries to be logged as separate log entries (for example, for reporting purposes or for engine statistics).
- By default, each generated Discard log entry is logged separately and displayed as a separate entry in the Logs view. Log Compression settings allow you to define the maximum number of separately logged entries. When the defined limit is reached, a single Discard log entry is logged. The single entry contains information on the total number of the generated Discard log entries. After this, the logging returns to normal and all the generated entries are once more logged and displayed separately.

7. Click **OK**.

What's Next?

- Close the Layer 2 Firewall's Properties dialog and refresh the engine's policy to transfer the configuration changes.

Setting Interface Options for Layer 2 Firewalls

Interface options allow you to select which interfaces are used for which types of system communications.

▼ To set the interface options

1. Open the Layer 2 Firewall element's properties and switch to the **Interfaces** tab.
2. Click the **Options** button below the interface table. The Interface Options dialog opens.

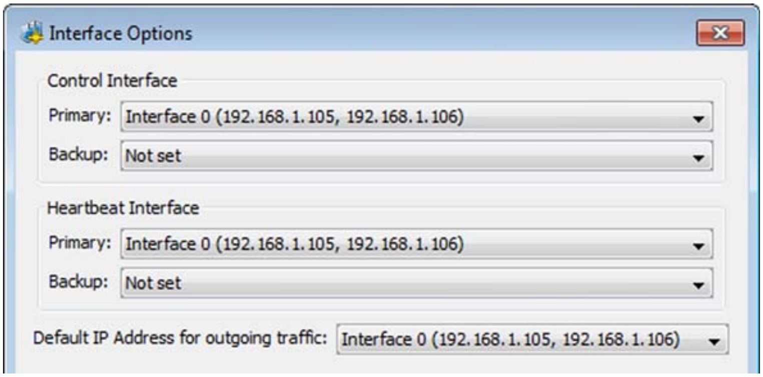

Control Interface

Primary: Interface 0 (192.168.1.105, 192.168.1.106)

Backup: Not set

Heartbeat Interface

Primary: Interface 0 (192.168.1.105, 192.168.1.106)

Backup: Not set

Default IP Address for outgoing traffic: Interface 0 (192.168.1.105, 192.168.1.106)

3. Select the interface options as explained in the table below.

Table 28.18 Layer 2 Firewall Interface Options

Option	Explanation
Control Interface	Select the **Primary** Control Interface for Management Server contact. This interface is used for communications with the Management Server
	(*Optional*) Select a **Backup** Control Interface that is used if the Primary interface is not available.
Node-Initiated contact to Management Server (*Single Layer 2 Firewalls only*)	Select the Node-Initiated contact to Management Server option if the engine is behind a device that applies dynamic NAT to the inbound management connections or blocks them. This option is always used with a dynamic control IP address, so it is always fixed to the selected state if the IP address is dynamic. When this option is selected, the engine opens a connection to the Management Server and maintains connectivity. If the connection is not open when you command the engine through the Management Client, the command is left pending until the engine opens the connection again.

Table 28.18 Layer 2 Firewall Interface Options (Continued)

Option	Explanation
Heartbeat Interface *(Layer 2 Firewall Clusters only)*	Select the **Primary** Heartbeat Interface for communications between the nodes in the IPS cluster. This must not be a VLAN Interface. A dedicated network (without other traffic) is strongly recommended for security and reliability of heartbeat communications. In a Layer 2 Firewall Cluster with only two nodes, a crossover cable without additional intermediary devices is recommended for the heartbeat link.
	Select a **Backup** Heartbeat Interface that is used if the Primary Heartbeat Interface is unavailable. It is not mandatory to configure a backup heartbeat, but we strongly recommend it. If heartbeat traffic is not delivered, the cluster cannot operate and traffic will be disturbed. We strongly recommend that you use a dedicated interface for the backup heartbeat as well.
Default IP Address for Outgoing Traffic (*Layer 2 Firewall Clusters only*)	This option defines the IP address that the nodes use if they have to initiate connections (system communications, ping, etc.) through an interface that has no Node Dedicated IP Address. You must select an interface that has an IP address defined for all nodes.

4. Click **OK**.

What's Next?

- If you are creating a new element, configure the routing as explained in Adding Routes for IPS Engines and Layer 2 Firewalls (page 605), and connect the engine to the Management Server as explained in Saving an Initial Configuration for Security Engines (page 515).
- Otherwise, click **OK** and refresh the Layer 2 Firewall's policy to activate the new interface configuration.

Master Engine Interface Configuration

Prerequisites: Creating New Master Engine Elements / Editing Master Engine Properties

The interface configuration for Master Engines consists of the following main steps:

1. Add the required number of network connections as explained in Defining Physical Interfaces for Master Engines (page 486).
   - If you want to use the interface to host a Virtual Security Engine, associate a Virtual Resource with the interface.
2. (*Optional)* Add the required number of VLANs as explained in Adding VLAN Interfaces for Master Engines (page 490).
   - If you want to use the interface to host a Virtual Security Engine, associate a Virtual Resource with the interface.
3. Configure the IP address settings as explained in Configuring Master Engine IP Addresses (page 493).
4. Select the interfaces that are used for system communications as explained in Setting Interface Options for Master Engines (page 497).

Related Tasks

- Firewall Interface Configuration (page 421)
- IPS Engine Interface Configuration (page 456)
- Layer 2 Firewall Interface Configuration (page 473)
- Virtual Security Engine Interface Configuration (page 498)
- Configuring Manual ARP Settings (page 510)

Defining Physical Interfaces for Master Engines

Physical Interfaces correspond to network ports on the Master Engine. By default, the numbering of the Physical Interfaces in the Management Client corresponds to the operating system interface numbering on the engine (Interface ID 0 is mapped to eth0, ID 1 to eth1, etc.). However, the mapping is not fixed and you can change it through the engine command line. See the relevant *Appliance Installation Guide* for details on which Interface IDs to map with which network ports.

For settings on the **Advanced** tab, see Configuring Advanced Interface Properties for Master Engines (page 495).

What's Next?

- To define Physical Interfaces for the Master Engine's own communications, proceed to Defining Physical Interfaces for Master Engine Communications (page 487).
- To define Physical Interfaces that are used by the Virtual Security Engines hosted on the Master Engine, proceed to Defining Physical Interfaces for Hosted Virtual Engine Communications (page 489).

Defining Physical Interfaces for Master Engine Communications

▼ To define physical interface for Master Engine communications

1. In the properties dialog for the Master Engine, switch to the **Interfaces** tab.
2. Right-click the empty space and select **New Physical Interface**. The Physical Interface Properties dialog opens.
3. Define the Physical Interface properties as explained in the table below.

Table 28.19 Physical Interface Properties for Master Engine Communications

Options		Explanation
Interface ID		The Interface ID automatically maps to a physical interface of the same number during the initial configuration of the engine, but the mapping can be changed as necessary through the engine's command line interface.
Type	Normal Interface	Corresponds to a single network interface on the Master Engine appliance.
	Aggregated Link in High-Availability Mode	Represents two interfaces on the Master Engine appliance. Only the first interface in the aggregated link is actively used. The second interface becomes active only if the first interface fails. If you configure an Aggregated Link in High-Availability mode, connect the first interface to one switch and the second interface to another switch.
	Aggregated Link in Load-Balancing Mode	Represents two interfaces on the Master Engine appliance. Both interfaces in the aggregated link are actively used and connections are automatically balanced between the two interfaces. Link aggregation in the load-balancing mode is implemented based on the IEEE 802.3ad Link Aggregation standard. If you configure an Aggregated Link in Load-Balancing Mode, connect both interfaces to a single switch. Make sure that the switch supports the Link Aggregation Control Protocol (LACP) and that LACP is configured on the switch.
Second Interface ID (*Only if interface type is Aggregated Link*)		The second interface in the aggregated link.
Virtual Resource		Do not select a Virtual Resource.
Cluster MAC Address		The MAC address for the Master Engine. Do not use the MAC address of any actual network card on any of the nodes.
QoS Mode (*Optional*)		Defines how QoS is applied to the link on this interface. If **Full QoS** or **DSCP Handling and Throttling** is selected as the QoS Mode, a QoS Policy must also be selected. If **Full QoS** is selected as the QoS Mode, the Throughput must also be defined.

Table 28.19 Physical Interface Properties for Master Engine Communications (Continued)

Options	Explanation
QoS Policy (*Full QoS or DSCP Handling and Throttling modes only*)	The QoS Policy for the link on this interface. For more information, see Getting Started with QoS (page 782).
Throughput (*Full QoS mode only*)	Enter the throughput for the link on this interface as kilobits per second (for example, 2048). The same throughput is automatically applied to any VLANs created under this Physical Interface. See Adding VLAN Interfaces for Master Engines (page 490). The throughput is for uplink speed (outgoing traffic) and typically must correspond to the speed of an Internet link (such as an ADSL line), or the combined speeds of several such links when they are connected to a single physical interface). **Caution!** Make sure you set the Interface speed correctly. When the bandwidth is set, the Firewall always scales the total amount of traffic on this interface to the bandwidth you defined. This happens even if there are no bandwidth limits or guarantees defined for any traffic.
MTU (*Optional*)	The MTU (maximum transmission unit) size on the connected link. Either enter a value between 400-65535 or select a common MTU value from the list. The default value (also the maximum standard MTU in Ethernet) is 1500. Do not set a value larger than the standard MTU unless you know that all the devices along the communications path support it.
Zone (*Optional*)	Select the network zone to which the Physical interface belongs from the list or select **Other** to select another Zone. If the Zone is not listed, create a new Zone element through the New icon at the top of the dialog. See Defining Zone Elements (page 730) for more information.

4. Click **OK**. The Physical Interface is added to the interface list.
5. Repeat from Step 2 to add more Physical Interfaces.

What’s Next?

- If you want to use VLANs on this Physical Interface, add the VLANs before adding IP addresses. Proceed to Adding VLAN Interfaces for Master Engines.
- Otherwise, add IP addresses directly to the Physical Interfaces as instructed in Configuring Master Engine IP Addresses (page 493).

Defining Physical Interfaces for Hosted Virtual Engine Communications

▼ To define physical interface for hosted Virtual Engine communications

1. In the properties dialog for the Master Engine, switch to the **Interfaces** tab.
2. Right-click the empty space and select **New Physical Interface**. The Physical Interface Properties dialog opens.
3. Define the Physical Interface properties as explained in the table below.

Table 28.20 Physical Interface Properties for Hosted Virtual Engine Communications

Options		Explanation
Interface ID		The Interface ID automatically maps to a physical interface of the same number during the initial configuration of the engine, but the mapping can be changed as necessary through the engine's command line interface.
Type	Normal Interface	Corresponds to a single network interface on the Master Engine appliance.
	Aggregated Link in High-Availability Mode	Represents two interfaces on the Master Engine appliance. Only the first interface in the aggregated link is actively used. The second interface becomes active only if the first interface fails. If you configure an Aggregated Link in High-Availability mode, connect the first interface to one switch and the second interface to another switch.
	Aggregated Link in Load-Balancing Mode	Represents two interfaces on the Master Engine appliance. Both interfaces in the aggregated link are actively used and connections are automatically balanced between the two interfaces. Link aggregation in the load-balancing mode is implemented based on the IEEE 802.3ad Link Aggregation standard. If you configure an Aggregated Link in Load-Balancing Mode, connect both interfaces to a single switch. Make sure that the switch supports the Link Aggregation Control Protocol (LACP) and that LACP is configured on the switch.
Second Interface ID (*Only if interface type is Aggregated Link*)		The second interface in the aggregated link.
Virtual Resource		The Virtual Resource associated with the interface. Select the same Virtual Resource in the properties of the Virtual Firewall element to add the Virtual Security Engine to the Master Engine. See Creating New Virtual Security Engines (page 407). Only one Virtual Resource can be selected for each Physical Interface. If you want to add multiple Virtual Resources, add VLAN interfaces to the Physical Interface and select the Virtual Resource in the VLAN interface properties as explained in Adding VLAN Interfaces for Master Engines (page 490).

Table 28.20 Physical Interface Properties for Hosted Virtual Engine Communications (Continued)

Options	Explanation
Allow VLAN Definition in Virtual Engine	Select this option to allow VLAN interfaces to be added to the automatically created Physical Interfaces in the Virtual Security Engine that is associated with this interface.
Virtual Engine Interface ID	Select the Interface ID of the Physical Interface in the Virtual Security Engine that is associated with this interface.
Cluster MAC Address	The MAC address for the Master Engine. Do not use the MAC address of any actual network card on any of the nodes.
Throughput (kbps) (*Optional*)	Enter the maximum throughput for Virtual Engines that use this interface as kilobits per second (for example, 2048). The same throughput is automatically applied to any VLANs created under this Physical Interface. See Adding VLAN Interfaces for Master Engines (page 490).
MTU (*Optional*)	The MTU (maximum transmission unit) size for Virtual Engines that use this interface. Either enter a value between 400-65535 or select a common MTU value from the list. The default value (also the maximum standard MTU in Ethernet) is 1500. Do not set a value larger than the standard MTU unless you know that all the devices along the communications path support it.

4. Click **OK**. The Physical Interface is added to the interface list.
5. Repeat from Step 2 to add more Physical Interfaces.

What's Next?

- If you want to use VLANs, add the VLANs before adding IP addresses. Proceed to Adding VLAN Interfaces for Master Engines.
- Otherwise, add IP addresses directly to the Physical Interfaces as instructed in Configuring Master Engine IP Addresses (page 493).

Adding VLAN Interfaces for Master Engines

VLANs divide a single physical network link into several virtual links. The maximum number of VLANs for a single Physical Interface is 4094. The VLANs must also be defined in the configuration of the switch/router to which the interface is connected.

What's Next?

- To define VLAN Interfaces for the Master Engine's own communications, proceed to Adding VLAN Interfaces for Master Engine Communications
- To define VLAN Interfaces that are used by the Virtual Security Engines hosted on the Master Engine, proceed to Adding VLAN Interfaces for Hosted Virtual Engine Communications (page 492).

Adding VLAN Interfaces for Master Engine Communications

For settings on the **Advanced** tab, see Configuring Advanced Interface Properties for Master Engines (page 495).

▼ To add a VLAN Interface for a Master Engine

1. In the properties dialog for the Master Engine, switch to the **Interfaces** tab.
2. Right-click a Physical Interface and select **New→VLAN Interface**. The VLAN Interface Properties dialog opens.
3. Define the VLAN interface properties as explained in the table below.

Table 28.21 VLAN Interface Properties - General Tab

Option	Explanation
VLAN ID	Enter the VLAN ID (1-4094). The VLAN IDs you add must be the same as the VLAN IDs that are used in the switch at the other end of the VLAN trunk. Each VLAN interface is identified as Interface-ID.VLAN-ID, for example 2.100 for Interface ID 2 and VLAN ID 100.
Virtual Resource	Do not select a Virtual Resource.
QoS Mode (*Optional*)	Defines how QoS is applied to the link on this interface. If **Full QoS** or **DSCP Handling and Throttling** is selected as the QoS Mode, a QoS Policy must also be selected. If **Full QoS** is selected as the QoS Mode, the Throughput must also be defined.
QoS Policy (*Full QoS or DSCP Handling and Throttling modes only*)	The QoS Policy for the link on the interface. For more information, see Getting Started with QoS (page 782).
Throughput (*Full QoS mode only)*	Enter the throughput for the link on this interface as kilobits per second (for example, 2048). If throughput is defined for the Physical Interface to which the VLAN Interface belongs, the throughput value is automatically inherited from the Physical Interface properties. **Caution!** The throughput for each VLAN Interface must not be higher than the throughput for the Physical Interface to which the VLAN Interface belongs. The throughput is for uplink speed (outgoing traffic) and typically must correspond to the speed of an Internet link (such as an ADSL line), or the combined speeds of several such links when they are connected to a single Physical Interface. **Caution!** Make sure you set the interface speed correctly. When the bandwidth is set, the Firewall always scales the total amount of traffic on this interface to the bandwidth you defined. This happens even if there are no bandwidth limits or guarantees defined for any traffic.

Table 28.21 VLAN Interface Properties - General Tab (Continued)

Option	Explanation
MTU *(Optional)*	The MTU (maximum transmission unit) size on the connected link. Either enter a value between 400-65535 or select a common MTU value from the list. **Caution!** The MTU for each VLAN Interface must not be higher than the MTU for the Physical Interface to which the VLAN Interface belongs. The default value (also the maximum standard MTU in Ethernet) is 1500. Do not set a value larger than the standard MTU unless you know that all the devices along the communications path support it.
Zone *(Optional)*	Select the network zone to which the VLAN interface belongs from the list or select **Other** to select another Zone. If the Zone is not listed, create a new Zone element through the New icon at the top of the dialog. See Defining Zone Elements (page 730) for more information.

4. Click **OK**. The specified VLAN ID is added to the Physical Interface.
5. Repeat from Step 2 to add further VLANs on the same or other Physical Interfaces.

What's Next?

▶ Add IP addresses to the VLAN Interfaces as instructed in Configuring Master Engine IP Addresses (page 493).

Adding VLAN Interfaces for Hosted Virtual Engine Communications

▼ To add a VLAN Interface for a hosted Virtual Engine

1. In the properties dialog for the Master Engine, switch to the **Interfaces** tab.
2. Right-click a Physical Interface and select **New→VLAN Interface**. The VLAN Interface Properties dialog opens.
3. Define the VLAN interface properties as explained in the table below.

Table 28.22 VLAN Interface Properties - General Tab

Option	Explanation
VLAN ID	Enter the VLAN ID (1-4094). The VLAN IDs you add must be the same as the VLAN IDs that are used in the switch at the other end of the VLAN trunk. Each VLAN interface is identified as Interface-ID.VLAN-ID, for example 2.100 for Interface ID 2 and VLAN ID 100.
Virtual Resource	The Virtual Resource associated with the interface. Select the same Virtual Resource in the properties of the Virtual Firewall element to add the Virtual Security Engine to the Master Engine. See Creating New Virtual Security Engines (page 407). Only one Virtual Resource can be selected for each VLAN interface.

Table 28.22 VLAN Interface Properties - General Tab (Continued)

Option	Explanation
Throughput (*Optional*)	The maximum throughput for the Virtual Engines that use this VLAN interface. Enter the throughput as kilobits per second (for example, 2048). If throughput is defined for the Physical Interface to which the VLAN Interface belongs, the throughput value is automatically inherited from the Physical Interface properties. **Caution!** The throughput for each VLAN Interface must not be higher than the throughput for the Physical Interface to which the VLAN Interface belongs. The throughput is for uplink speed (outgoing traffic) and typically must correspond to the speed of an Internet link (such as an ADSL line), or the combined speeds of several such links when they are connected to a single Physical Interface. **Caution!** Make sure you set the interface speed correctly. When the bandwidth is set, the Firewall always scales the total amount of traffic on this interface to the bandwidth you defined. This happens even if there are no bandwidth limits or guarantees defined for any traffic.
MTU (*Optional*)	The MTU (maximum transmission unit) size for Virtual Engines that use this interface. Either enter a value between 400-65535 or select a common MTU value from the list. **Caution!** The MTU for each VLAN Interface must not be higher than the MTU for the Physical Interface to which the VLAN Interface belongs. The default value (also the maximum standard MTU in Ethernet) is 1500. Do not set a value larger than the standard MTU unless you know that all the devices along the communications path support it.

4. Click **OK**. The specified VLAN ID is added to the Physical Interface.
5. Repeat from Step 2 to add further VLANs on the same or other Physical Interfaces.

What's Next?

▶ Add IP addresses to the VLAN Interfaces as instructed in Configuring Master Engine IP Addresses.

Configuring Master Engine IP Addresses

You can configure several IP addresses on each Physical Interface or VLAN Interface that does not have a Virtual Resource associated with it.

What's Next?

▶ Adding IPv4 Addresses for a Master Engine (page 494)

▶ Adding IPv6 Addresses for a Master Engine (page 494)

Adding IPv4 Addresses for a Master Engine

▼ To add IPv4 addresses for a Master Engine

1. In the properties dialog for the Master Engine, switch to the **Interfaces** tab.
2. Right-click a Physical interface or VLAN Interface and select **New→IPv4 Address**. The IP Address Properties dialog opens.

> **Note – If you use VLAN Interfaces, you must add the IPv4 Addresses to the VLAN Interfaces.**

3. Enter the **IPv4 Address** for each node.
4. If necessary, double-click the **Contact Address** field and define the contact address(es). See Defining Contact IP Addresses (page 66) for more information.
   - Enter the **Default** contact address. The default contact address is used by default whenever a component that belongs to another Location connects to this interface.
   - If components from some Locations cannot use the Default contact address, click **Exceptions** to define Location-specific contact addresses.
5. Check the automatically filled-in **Netmask** and adjust it as necessary.
6. Click **OK**.
7. Repeat from Step 2 to add further IPv4 addresses to the same or other interface.

What's Next?

- ▶ If you are configuring a new Master Engine, or if you want to change the roles the different interfaces have in the configuration, proceed to Setting Interface Options for Master Engines (page 497).
- ▶ Otherwise, close the Master Engine Properties dialog and refresh the policy to transfer the configuration changes.

Adding IPv6 Addresses for a Master Engine

▼ To add IPv6 addresses for a Master Engine

1. In the properties dialog for the Master Engine, switch to the **Interfaces** tab.
2. Right-click a Physical interface or VLAN Interface and select **New→IPv6 Address**. The IP Address Properties dialog opens.

> **Note – If you use VLAN Interfaces, you must add the IPv6 Addresses to the VLAN Interfaces.**

3. Enter the **IPv6 Address** for each node.
4. Check the automatically filled-in **Prefix Length** and adjust it if necessary by entering a value between 0-128. The Network Address is automatically generated.
5. Click **OK**.
6. *(Optional)* Repeat from Step 2 to add further IPv6 addresses to the same or other interface.

What's Next?

- If you are configuring a new Master Engine, or if you want to change the roles the different interfaces have in the configuration, proceed to Setting Interface Options for Master Engines (page 497).
- Otherwise, close the Master Engine Properties dialog and refresh the policy to transfer the configuration changes.

Configuring Advanced Interface Properties for Master Engines

Advanced settings allow you to configure SYN Rate Limits and Log Compression for interfaces that are used for the Master Engine's own communications. Advanced settings are not available for interfaces that are associated with a Virtual Resource.

By default, each generated Antispoofing and Discard log entry is logged separately and displayed as a separate entry in the Logs view. Log Compression settings allow you to define the maximum number of separately logged entries. When the defined limit is reached, a single antispoofing log entry or Discard log entry is logged. The single entry contains information on the total number of the generated Antispoofing log entries or Discard log entries. After this, the logging returns to normal and all the generated entries are once more logged and displayed separately. Log Compression is useful when the routing configuration generates a large volume of antispoofing logs or the number of Discard logs becomes high (for example, as a result of a SYN flood attack).

Note – The SYN Rate Limits and Log Compression settings in the interface properties override the general SYN Rate Limits and Log Compression settings that are defined on the Advanced tab in the Master Engine properties. See Adjusting Master Engine Traffic Handling Parameters (page 569).

▼ To configure advanced interface properties for Master Engines

1. In the properties dialog for the Master Engine, switch to the **Interfaces** tab.
2. Right-click a Physical Interface or a VLAN Interface and select **Edit Physical Interface**, or **Edit VLAN Interface**. The properties dialog for the interface opens.
3. Switch to the **Advanced** tab.
4. Select **Override Engine's Default Settings**. The options for SYN Rate Limits and Log Compression are enabled.
5. (*Optional*) Define the **SYN Rate Limits**

Setting	Description
Default	The interface uses the SYN Rate Limits defined on the Advanced tab in the Master Engine properties.
Off	Disables SYN Rate Limits on the interface.

Setting	Description
Automatic	This is the recommended mode if you want to override the general SYN Rate Limits defined on the Advanced tab in the engine’s properties. The engine automatically calculates the number of **Allowed SYNs per Second** (the number of allowed SYN packets per second) and the **Burst Size** (the number of allowed SYNs before the engine starts limiting the SYN rate) for the interface based on the engine’s capacity and memory size.
Custom	Enter the desired values for **Allowed SYNs per Second** and **Burst Size**. We recommend that the Burst Size be at least one tenth of the Allowed SYNs per Second value. If the Burst Size is too small, SYN Rate Limits do not work. For example, if the value for Allowed SYNs per Second is 10000, the Burst Size should be at least 1000.

Caution – The recommended values for the SYN Flood Settings depend on your network environment. If the Custom settings are not carefully configured, the capacity of the Master Engine may suffer or SYN Rate Limits may not work correctly.

6. (*Optional*) Enable/Disable **Log Compression** and enter the desired values for the Antispoofing entries and (optionally) for Discard entries.

Setting	Description
Log Rate (Entries/s)	The maximum number of entries per second. The default value for antispoofing entries is 100 entries/s. By default, Discard log entries are not compressed.
Burst Size (Entries)	The maximum number of matching entries in a single burst. The default value for antispoofing entries is 1000 entries. By default, Discard log entries are not compressed.

Note – Do not enable Log Compression if you want all the antispoofing and Discard entries to be logged as separate log entries (for example, for reporting purposes or for statistics).

7. Click **OK**.

What’s Next?

▶ Close the Master Engine Properties dialog and refresh the policy to transfer the configuration changes.

Setting Interface Options for Master Engines

The Interface Options dialog contains the settings for selecting which IP addresses are used in particular roles in system communications (for example, in communications between the Master Engine and the Management Server). Only IPv4 addresses are used in system communications.

▼ To set the interface options for a Master Engine

1. In the properties dialog for the Master Engine, switch to the **Interfaces** tab.
2. Click **Options**. The Interface Options dialog opens.
3. Select the interface options as explained in the table below.

Table 28.23 Firewall Interface Options

Option	Explanation
Control Interface	Select the **Primary** Control Interface for Management Server contact. This interface is used for communications with the Management Server
	(*Optional*) Select a **Backup** Control Interface that is used if the Primary interface is not available.
Heartbeat Interface	Select the **Primary** Heartbeat Interface for communications between the nodes. We recommend that you use a Physical Interface, not a VLAN Interface. We strongly recommend that you do not direct any other traffic through this interface. A dedicated network helps ensure reliable and secure operation. **Caution!** Primary and Backup Heartbeat networks exchange confidential information. If dedicated networks are not possible, configure the cluster to encrypt the exchanged information. See Adjusting Master Engine Clustering Options (page 571).
	Select a **Backup** Heartbeat Interface that is used if the Primary Heartbeat Interface is unavailable. It is not mandatory to configure a backup heartbeat, but we strongly recommend it. If heartbeat traffic is not delivered, the cluster cannot operate and traffic will be disturbed. We strongly recommend that you use a dedicated interface for the backup heartbeat as well. **Caution!** Primary and Backup Heartbeat networks exchange confidential information. If dedicated networks are not possible, configure the cluster to encrypt the exchanged information. See Adjusting Master Engine Clustering Options (page 571).
Default IP Address for Outgoing Traffic	This option defines the IP address that the nodes use if they have to initiate connections (system communications, ping, etc.) through an interface that has no IP Address. You must select an interface that has an IP address defined for all nodes.

4. Click **OK**.

What’s Next?

- If you are configuring a new Master Engine as part of the configuration process for Virtual Security Engines, continue by Creating New Virtual Security Engines (page 407).
- Otherwise, close the Master Engine Properties dialog and refresh the policy to transfer the configuration changes.

Virtual Security Engine Interface Configuration

Prerequisites: Master Engine Interface Configuration

The interface configuration for Virtual Security Engines consists of the following main steps:

1. Modify the automatically-created Physical Interfaces as explained in Modifying Physical Interfaces for Virtual Security Engines (page 499).
2. (*Optional)* Add the required number of VLANs as explained in Adding VLAN Interfaces for Virtual Security Engines.
3. (*Optional)* Define Tunnel Interfaces for the Route-Based VPN as explained in Defining Tunnel Interfaces for Virtual Security Engines (page 502)
4. Configure the IP address settings as explained in Configuring Virtual Security Engine IP Addresses (page 504).
5. (*Optional*) Define Loopback IP addresses to allow the firewall to communicate with itself as explained in Adding Loopback IP Addresses for Virtual Firewalls (page 505).
6. Select the interfaces that are used for system communications as explained in Setting Interface Options for Virtual Security Engines (page 509).

Related Tasks

- Firewall Interface Configuration (page 421)
- IPS Engine Interface Configuration (page 456)
- Layer 2 Firewall Interface Configuration (page 473)
- Master Engine Interface Configuration (page 486)
- Configuring Manual ARP Settings (page 510)
- Routing Multicast Traffic (page 593)

Modifying Physical Interfaces for Virtual Security Engines

Physical Interfaces for Virtual Security Engines represent interfaces allocated to the Virtual Security Engine in the Master Engine. When you select the Virtual Resource for the Virtual Security Engine, Physical Interfaces are automatically created based on the interface configuration in the Master Engine properties. The number of Physical Interfaces depends on the number of interfaces allocated to the Virtual Security Engine in the Master Engine. You can modify the automatically-created Physical Interfaces.

For settings on the **DHCP** tab, see Routing DHCP Messages (page 590).

For settings on the **Advanced** tab, see Configuring Advanced Interface Properties for Virtual Security Engines (page 506).

▼ To modify physical interfaces for Virtual Security Engines

1. In the properties dialog for the Virtual Firewall, switch to the **Interfaces** tab.
2. Right-click a Physical Interface and select **Edit Physical Interface**. The Physical Interface Properties dialog opens.
3. Define the Physical Interface properties as explained in the table below.

Table 28.24 Physical Interface Properties - General Tab

Options	Explanation
Interface ID	The Interface ID automatically maps to a physical interface of the same number during the initial configuration of the engine, but the mapping can be changed as necessary through the engine’s command line interface.
QoS Mode (*Optional*)	Defines how QoS is applied to the link on this interface. If **Full QoS** or **DSCP Handling and Throttling** is selected as the QoS Mode, a QoS Policy must also be selected. If **Full QoS** is selected as the QoS Mode, the Throughput must also be defined.
QoS Policy (*Full QoS or DSCP Handling and Throttling modes only*)	The QoS Policy for the link on this interface. For more information, see Getting Started with QoS (page 782).

Table 28.24 Physical Interface Properties - General Tab (Continued)

Options	Explanation
Throughput (*Full QoS mode only*)	The Throughput for the interface may be limited by settings defined in the properties of the Master Engine. Enter the throughput as kilobits per second (for example, 2048). The same throughput is automatically applied to any VLANs created under this Physical Interface. See Adding VLAN Interfaces for Virtual Security Engines (page 501). **Caution!** The throughput for the Virtual Firewall Physical Interface must not be higher than the throughput for the Master Engine interface that hosts the Virtual Firewall. Contact the administrator of the Master Engine before changing this setting. The throughput is for uplink speed (outgoing traffic) and typically must correspond to the speed of an Internet link (such as an ADSL line), or the combined speeds of several such links when they are connected to a single physical interface). **Caution!** Make sure you set the Interface speed correctly. When the bandwidth is set, the Firewall always scales the total amount of traffic on this interface to the bandwidth you defined. This happens even if there are no bandwidth limits or guarantees defined for any traffic.
Zone (*Optional*)	Select the network zone to which the Physical interface belongs from the list or select **Other** to select another Zone. If the Zone is not listed, create a new Zone element through the New icon at the top of the dialog. See Defining Zone Elements (page 730) for more information.

4. Click **OK**. The Physical Interface is added to the interface list.
5. Repeat from Step 2 to add more Physical Interfaces.

What's Next?

- If you want to use VLANs, add the VLANs before adding IP addresses. Proceed to Adding VLAN Interfaces for Virtual Security Engines (page 501).
- If you want to define Tunnel Interfaces for the Route-Based VPN, continue by Defining Tunnel Interfaces for Virtual Security Engines (page 502).
- Otherwise, add IP addresses directly to the Physical Interfaces as instructed in Configuring Virtual Security Engine IP Addresses (page 504).

Adding VLAN Interfaces for Virtual Security Engines

VLAN Interfaces can only be added for Virtual Firewalls if the creation of VLAN Interfaces for Virtual Firewalls is enabled in the Master Engine Properties. VLANs divide a single physical network link into several virtual links. The maximum number of VLANs for a single Physical Interface is 4094. The VLANs must also be defined in the configuration of the switch/router to which the interface is connected.

Note – You cannot add VLAN Interfaces on top of other VLAN Interfaces. Depending on the configuration of the Master Engine that hosts the Virtual Security Engine, you may not be able to create valid VLAN Interfaces for the Virtual Security Engine. Contact the administrator who configured the Master Engine.

For settings on the **DHCP** tab, see Routing DHCP Messages (page 590).

For settings on the **Advanced** tab, see Configuring Advanced Interface Properties for Virtual Security Engines (page 506).

▼ To add a VLAN Interface for a Virtual Security Engine

1. In the properties dialog for the Virtual Firewall, switch to the **Interfaces** tab.
2. Right-click a Physical Interface and select **New→VLAN Interface**. The VLAN Interface Properties dialog opens.
3. Define the VLAN interface properties as explained in the table below.

Table 28.25 VLAN Interface Properties - General Tab

Option	Explanation
VLAN ID	Enter the VLAN ID (1-4094). The VLAN IDs you add must be the same as the VLAN IDs that are used in the switch at the other end of the VLAN trunk. Each VLAN interface is identified as Interface-ID.VLAN-ID, for example 2.100 for Interface ID 2 and VLAN ID 100.
QoS Mode (*Optional*)	Defines how QoS is applied to the link on this interface. If **Full QoS** or **DSCP Handling and Throttling** is selected as the QoS Mode, a QoS Policy must also be selected. If **Full QoS** is selected as the QoS Mode, the Throughput must also be defined.
QoS Policy (*Full QoS or DSCP Handling and Throttling modes only*)	The QoS Policy for the link on the interface. For more information, see Getting Started with QoS (page 782).

Table 28.25 VLAN Interface Properties - General Tab (Continued)

Option	Explanation
Throughput (*Full QoS mode only*)	The Throughput for the interface may be limited by settings defined in the properties of the Master Engine. Enter the throughput as kilobits per second (for example, 2048). If throughput is defined for the Physical Interface to which the VLAN Interface belongs, the throughput value is automatically inherited from the Physical Interface properties. **Caution!** The throughput for the Virtual Firewall Physical Interface must not be higher than the throughput for the Master Engine interface that hosts the Virtual Firewall. Contact the administrator of the Master Engine before changing this setting. The throughput is for uplink speed (outgoing traffic) and typically must correspond to the speed of an Internet link (such as an ADSL line), or the combined speeds of several such links when they are connected to a single Physical Interface. **Caution!** Make sure you set the interface speed correctly. When the bandwidth is set, the Firewall always scales the total amount of traffic on this interface to the bandwidth you defined. This happens even if there are no bandwidth limits or guarantees defined for any traffic.
Zone (*Optional*)	Select the network zone to which the VLAN interface belongs from the list or select **Other** to select another Zone. If the Zone is not listed, create a new Zone element through the New icon at the top of the dialog. See Defining Zone Elements (page 730) for more information.

4. Click **OK**. The specified VLAN ID is added to the Physical Interface.
5. Repeat from Step 2 to add further VLANs on the same or other Physical Interfaces.

What's Next?

- If you want to define Tunnel Intefaces for the Route-Based VPN, continue by Defining Tunnel Interfaces for Virtual Security Engines.
- Otherwise, add IP addresses to the VLAN Interfaces as instructed in Configuring Virtual Security Engine IP Addresses (page 504).

Defining Tunnel Interfaces for Virtual Security Engines

Tunnel Interfaces define end-points for tunnels in the Route-Based VPN. Any traffic that is routed to a Tunnel Interface and allowed by the Firewall Access rules is sent into the tunnel. See Getting Started With IPsec VPNs (page 918) for more information about the Route-Based VPN.

You can optionally add IPv4 and/or IPv6 addresses to a Tunnel Interface. Tunnel Interfaces can only have static IP addresses. Any IP address can be added to a Tunnel Interface, even if the same IP address is used on another interface or as a loopback IP address. Adding an IP address allows you to define the source IP address of traffic sent from the engine node itself. For example, an IP address is recommended to provide a source IP address for dynamic routing daemons, for IGMP proxy, and for Protocol Independent Multicast - Sparse-Mode (PIM-SM) configuration. If no IP address is added to the Tunnel Interface, the source IP address for traffic sent from the engine node is automatically selected according to the **Bypass Default IP Address** setting in the Interface Options for the firewall. See Setting Interface Options for Virtual Security Engines (page 509).

The mapping of Tunnel Interfaces to physical network interfaces on the engine is done automatically by the engine operating system based on the routing configuration.

For settings on the **Advanced** tab, see Configuring Advanced Interface Properties for Virtual Security Engines (page 506).

▼ To define Tunnel Interfaces

1. Open the Virtual Firewall Properties and switch to the **Interfaces** tab.
2. Click **Add** and select **Tunnel Interface**. The Tunnel Interface Properties dialog opens.

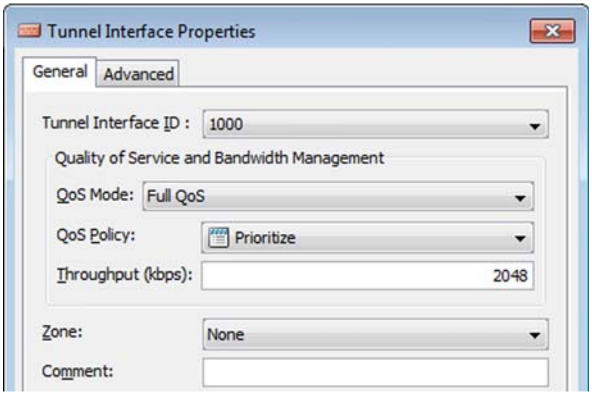

Tunnel Interface Properties
General
Advanced
Quality of Service and Bandwidth Management

Tunnel Interface ID : 1000

QoS Mode: Full QoS

QoS Policy: Prioritize

Throughput (kbps): 2048

Zone: None

Comment:

3. Select the **Tunnel Interface ID**.
4. (*Optional*) Select the **QoS Mode** to define how QoS is applied to the interface.
5. (*Full QoS or DSCP Handling and Throttling modes only*) Select the **QoS Policy** for the link on the interface. For more information, see Getting Started with QoS (page 782).
6. (*Full QoS mode only*) Enter the **Throughtput** for the link on this interface.
   - Enter throughput as kilobits per second (for example, 2048).

Caution – The throughput for the Virtual Firewall Physical Interface must not be higher than the throughput for the Master Engine interface that hosts the Virtual Firewall. Contact the administrator of the Master Engine before changing this setting.

   - The throughput is for uplink speed (outgoing traffic) and typically must correspond to the speed of an Internet link (such as an ADSL line), or the combined speeds of several such links when they are connected to a single physical interface).

Caution – Make sure you set the Interface speed correctly. When the bandwidth is set, the engine always scales the total amount of traffic on this interface to the bandwidth you defined. This happens even if there are no bandwidth limits or guarantees defined for any traffic.

7. Click **OK**.

What's Next?

▶ If you are configuring a new Virtual Firewall, continue by Configuring Virtual Security Engine IP Addresses (page 504).

▶ Otherwise, close the Firewall's Properties dialog and refresh the Firewall's policy to transfer the configuration changes.

Related Tasks

- ▶ Defining Tunnel Interfaces for Firewalls (page 435)

Configuring Virtual Security Engine IP Addresses

A Virtual Security Engine's interfaces can have the following types of IP addresses:

- A Physical Interface can have one or more static IP addresses.
- A Tunnel Interface can have one or more static IP addresses.

You may need to define a *contact address* if you enter a private static address and NAT is used to translate it to a different external IP address. The external IP address must be configured as the contact address if other Management Center components need to use the external IP address to contact this engine or if the IP address is used as a VPN end-point. Only IPv4 addresses are used in system communications. Refer to Defining Contact IP Addresses (page 66) for more information.

What's Next?

- ▶ Adding IPv4 Addresses for Virtual Security Engines.
- ▶ Adding IPv6 Addresses for Virtual Security Engines (page 505).

Adding IPv4 Addresses for Virtual Security Engines

▼ To add an IPv4 address for a Virtual Security Engine

1. In the properties dialog for the Virtual Firewall, switch to the **Interfaces** tab.
2. Right-click a Physical Interface, VLAN Interface, or Tunnel Interface and select **New→IPv4 Address**. The IP Address Properties dialog opens.

Note – If you use VLAN Interfaces, you must add the IPv4 Addresses to the VLAN Interfaces.

3. Enter the **IPv4 Address**.
4. If necessary, define the Contact Address information (see Defining Contact IP Addresses (page 66) for detailed instructions).
   - Enter the **Default** that is used by when a component that belongs to another Location connects to this interface.
   - If components from some Locations cannot use the Default contact address, click **Exceptions** to define Location-specific contact addresses.
5. Check the automatically filled-in **Netmask** and adjust it as necessary.
6. Click **OK**.
7. Repeat from Step 2 to add further IP addresses to the same or other interface.

What's Next?

- To add IPv6 addresses, proceed to Adding IPv6 Addresses for Virtual Security Engines.
- If you are creating a new Virtual Firewall, or if you want to change the roles the different interfaces have in the configuration, proceed to Setting Interface Options for Virtual Security Engines (page 509).
- Otherwise, close the Virtual Firewall's Properties dialog and refresh the Virtual Firewall's policy to transfer the configuration changes.

Adding IPv6 Addresses for Virtual Security Engines

▼ To add IPv6 addresses for a Virtual Security Engine

1. In the properties dialog for the Virtual Security Engine, switch to the **Interfaces** tab.
2. Right-click a Physical interface or VLAN Interface and select **New→IPv6 Address**. The IP Address Properties dialog opens.

Note – If you use VLAN Interfaces, you must add the IPv6 Addresses to the VLAN Interfaces.

3. Enter the **IPv6 Address** for each node.
4. Check the automatically filled-in **Prefix Length** and adjust it if necessary by entering a value between 0-128. The Network Address is automatically generated.
5. Click **OK**.
6. *(Optional)* Repeat from Step 2 to add further IPv6 addresses to the same or other interface.

What's Next?

- If you are creating a new Virtual Firewall, or if you want to change the roles the different interfaces have in the configuration, proceed to Setting Interface Options for Virtual Security Engines (page 509).
- Otherwise, close the Virtual Firewall's Properties dialog and refresh the Virtual Firewall's policy to transfer the configuration changes.

Adding Loopback IP Addresses for Virtual Firewalls

A Loopback IP address allows the Virtual Firewall to communicate with itself. Loopback IP addresses do not create connectivity to any network. Loopback IP addresses only route traffic from the Virtual Firewall to itself.

- You can add several loopback IP addresses to each Virtual Firewall.
- Any IP address that is not already used on another Physical or VLAN interface in the same Virtual Firewall can be used as a loopback IP address.
- The same IP address can be used as a loopback IP address and as the IP address of a Tunnel Interface.
- Loopback IP addresses can be used as the Identity for Authentication Requests, the Source for Authentication Requests, and the Default IP Address for Outgoing Traffic.

▼ To add a loopback IP address for a Virtual Firewall

1. In the properties dialog for the Virtual Firewall, switch to the **Interfaces** tab.
2. Click **Options**. The Interface Options dialog opens.
3. Switch to the **Loopback** tab.

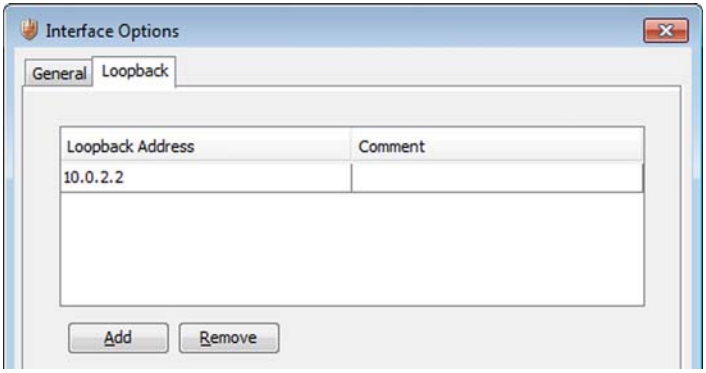

Loopback Address	Comment
10.0.2.2	

4. Click **Add**. A row is added to the table.
5. Click the **Loopback Address** cell and enter the loopback IP address.

Note – If the IP address you want to use as a loopback IP address is already used on another Physical or VLAN Interface, you must remove the IP address from the interface configuration before using it as a loopback IP address.

6. Click **OK**.

What's Next?

▶ Close the Virtual Firewall's Properties dialog and refresh the policy to transfer the configuration changes.

Configuring Advanced Interface Properties for Virtual Security Engines

Advanced settings allow you to configure SYN Rate Limits, Log Compression, and IPv6 Router Advertisements on the Virtual Firewall's interfaces.

By default, each generated Antispoofing and Discard log entry is logged separately and displayed as a separate entry in the Logs view. Log Compression settings allow you to define the maximum number of separately logged entries. When the defined limit is reached, a single antispoofing log entry or Discard log entry is logged. The single entry contains information on the total number of the generated Antispoofing log entries or Discard log entries. After this, the logging returns to normal and all the generated entries are once more logged and displayed separately. Log Compression is useful when the routing configuration generates a large volume of antispoofing logs or the number of Discard logs becomes high (for example, as a result of a SYN flood attack).

Router advertisements are packets that contain network layer configuration parameters. Enabling IPv6 Router Advertisements allows devices that connect to the same IPv6 network as the Virtual Firewall to acquire IP addresses automatically. This eliminates the need for a separate DHCP server. The Router Advertisement messages specify what configuration information the Virtual Firewall has available.

Note – The SYN Rate Limits and Log Compression settings in the interface properties override the general SYN Rate Limits and Log Compression settings that are defined on the Advanced tab in the Firewall or Firewall Cluster properties. See Configuring Default SYN Rate Limits (page 579) and Configuring Default Log Handling Settings (page 580).

▼ To configure advanced interface properties for Virtual Firewalls

1. In the properties dialog for the Virtual Firewall, switch to the **Interfaces** tab.
2. Right-click a Physical Interface or a VLAN Interface and select **Edit Physical Interface** or **Edit VLAN Interface**. The properties dialog for the interface opens.
3. Switch to the **Advanced** tab.

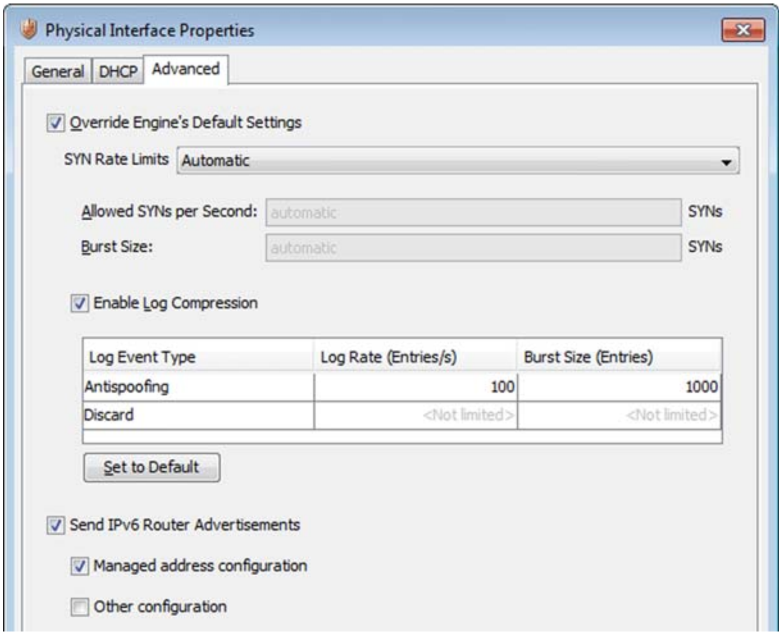

☑ Override Engine's Default Settings

SYN Rate Limits: Automatic

Allowed SYNs per Second: automatic SYNs

Burst Size: automatic SYNs

☑ Enable Log Compression

Log Event Type	Log Rate (Entries/s)	Burst Size (Entries)
Antispoofing	100	1000
Discard	<Not limited>	<Not limited>

☑ Send IPv6 Router Advertisements

☑ Managed address configuration

□ Other configuration

4. Select **Override Engine's Default Settings**. The options for SYN Rate Limits, Log Compression, and IPv6 Router Advertisements are enabled.
5. (*Optional*) Define the **SYN Rate Limits**

Setting	Description
Default	The interface uses the SYN Rate Limits defined on the Advanced tab in the Virtual Firewall properties (see Adjusting Virtual Security Engine Traffic Handling Parameters (page 575).
Off	Disables SYN Rate Limits on the interface.

Setting	Description
Automatic	This is the recommended mode if you want to override the general SYN Rate Limits defined on the Advanced tab in the engine's properties. The engine automatically calculates the number of **Allowed SYNs per Second** (the number of allowed SYN packets per second) and the **Burst Size** (the number of allowed SYNs before the engine starts limiting the SYN rate) for the interface based on the engine's capacity and memory size.
Custom	Enter the desired values for **Allowed SYNs per Second** and **Burst Size**. We recommend that the Burst Size be at least one tenth of the Allowed SYNs per Second value. If the Burst Size is too small, SYN Rate Limits do not work. For example, if the value for Allowed SYNs per Second is 10000, the Burst Size should be at least 1000.

Caution – The recommended values for the SYN Flood Settings depend on your network environment. If the Custom settings are not carefully configured, the capacity of the engine may suffer or SYN Rate Limits may not work correctly.

6. (*Optional*) Enable/Disable **Log Compression** and enter the desired values for the Antispoofing entries and (optionally) for Discard entries.

Setting	Description
Log Rate (Entries/s)	The maximum number of entries per second. The default value for antispoofing entries is 100 entries/s. By default, Discard log entries are not compressed.
Burst Size (Entries)	The maximum number of matching entries in a single burst. The default value for antispoofing entries is 1000 entries. By default, Discard log entries are not compressed.

Note – Do not enable Log Compression if you want all the antispoofing and Discard entries to be logged as separate log entries (for example, for reporting purposes).

7. (*Optional*) Select **Send IPv6 Router Advertisements** and specify what configuration information is offered in the Router Advertisement messages:

Setting	Description
Managed address configuration	The Virtual Firewall offers IPv6 addresses over the Dynamic Host Configuration Protocol (DHCPv6). When this option is selected, any other available configuration information is also sent in the DHCPv6 communications.
Other configuration	The Virtual Firewall offers other configuration information, such as DNS-related information or information on other servers within the network, over DHCPv6.

8. Click **OK**.

What's Next?

- ▶ Close the Virtual Firewall's Properties dialog and refresh the policy to transfer the configuration changes.

Setting Interface Options for Virtual Security Engines

The Interface Options dialog contains the settings for selecting which IP addresses are used in particular roles. All communication between Virtual Security Engines and the SMC is proxied by the Master Engine. Virtual Security Engines do not have any interfaces for system communication.

▼ To set the interface options

1. In the properties dialog for the Virtual Firewall, switch to the **Interfaces** tab.
2. Click **Options**. The Interface Options dialog opens.
3. Select the interface options as explained in the table below.

Table 28.26 Virtual Security Engine Interface Options

Option	Explanation
Identity for Authentication Requests	The IP address of the selected interface is used when an engine contacts an external authentication server and it is also displayed (by default) to end-users in Telnet-based authentication. This option does not affect the routing of the connection with the authentication server. The IP address is used only as a parameter inside the authentication request payload to give a name to the request sender.
Source for Authentication Requests	By default, the source IP address for authentication requests is selected according to routing. If the authentication requests are sent to an external authentication server over VPN, select an interface with a Node Dedicated IP address that you want use for the authentication requests. If you use the 5.4 Authentication Server component for authenticating the users, this setting does not have any effect.
Default IP Address for Outgoing Traffic	This option defines the IP address that the nodes use if they have to initiate connections (system communications, ping, etc.) through an interface that has no Node Dedicated IP Address. In Firewall Clusters, you must select an interface that has an IP address defined for all nodes.
Bypass Default IP Address	Defines how the source IP address for traffic sent from the engine node is selected for Tunnel Interfaces that do not have IP addresses. **Use Link Address in Unnumbered Tunnel Interface**: The loopback IP address defined for the engine node is used as the source IP address of traffic sent from the engine node. See Adding Loopback IP Addresses for Virtual Firewalls (page 505). This allows the loopback IP address to be used as an end-point in the Route-Based VPN. **Use Default Outgoing in Unnumbered Tunnel Interface**: the Default IP Address for Outgoing Traffic is used as the source IP address of traffic sent from the engine node.

4. Click **OK**.

What's Next?

- Close the Virtual Firewall's Properties dialog and refresh the Virtual Firewall's policy to transfer the configuration changes.

Configuring Manual ARP Settings

Prerequisites: See Creating and Modifying Engine Elements and Network Interface Configuration for overviews

ARP (Address Resolution Protocol) entries are normally managed automatically based on the routing configuration. It is not necessary to add manual ARP entries on security engines unless there are ARP-related problems (such as devices that do not respond to gratuitous ARP requests, or that impose a significant delay on such operations).

You can add and remove manual ARP entries in the Properties dialog for the engine elements. The manual ARP entries are generated by the engines regardless of the policy installed on the nodes. Both static and proxy ARP entries are supported on Firewalls. Only static ARP entries defined with IPv4 addresses are supported on IPS engines and Layer 2 Firewalls.

▼ To add or suppress an ARP entry

1. Right-click the engine and select **Properties**. The Properties dialog opens.
2. Switch to the **Interfaces** tab.
3. Click the **ARP Entries** button below the interface table. The ARP Entry Properties dialog opens.

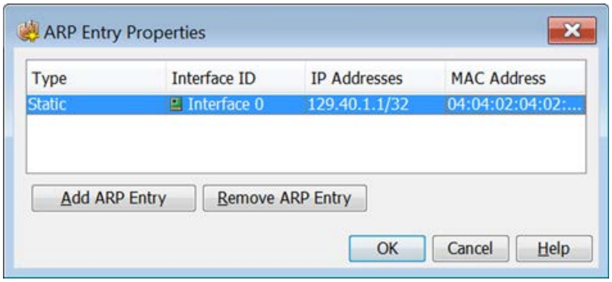

Type	Interface ID	IP Addresses	MAC Address
Static	Interface 0	129.40.1.1/32	04:04:02:04:02:...

4. Modify the ARP entries:
   - To suppress an entry, select the entry and click **Remove ARP Entry**.
   - To add an entry, click **Add ARP Entry**. A new entry is added in the table.
5. Click the **Type** cell for the new entry and select the type of ARP entry you want to add:
   - A **Static** ARP entry gives the engine a permanent reference to an IP address/MAC address pair. All entries are of this type on IPS engines and Layer 2 Firewalls.
   - A **Proxy** ARP entry gives a Firewall engine a reference to an IP address/MAC address pair that the Firewall should perform proxy ARP for. Proxy ARP is possible only for hosts located in networks directly connected to the Firewall.
6. Select the **Interface ID** for the interface on which you want to apply this ARP entry.
7. Double-click the **IP Address** cell and enter the IP Address. Only IPv4 addresses are supported on IPS engines and Layer 2 Firewalls.
8. Double-click the **MAC Address** cell and enter the MAC Address.
9. Click **OK** to close the ARP Entry Properties dialog.

10. Click **OK** to close the engine properties.

What's Next?

▶ Refresh the policy to transfer the new configuration to the engines.

Activating the Internal DHCP Server on a Firewall Interface

Prerequisites: See Getting Started with Interface Configuration

Single Firewalls and Firewall Clusters have an integrated DHCP server that can be set up independently on several Physical Interfaces and VLANs. When VLANs are configured, the DHCP server must be set up separately for each VLAN. Only IPv4 addresses are supported. To use this feature, the Firewall interface must have at least one IPv4 address.

Note – You can use the internal DHCP server to provide IP addresses to the VPN client Virtual Adapter only if you use Single Firewalls as VPN gateways.

▼ To activate the internal DHCP server on a Firewall interface

1. Right-click the Single Firewall or Firewall Cluster element and select **Properties**. The Properties dialog for the Firewall opens.
2. Switch to the **Interfaces** tab.
3. Right-click a Physical Interface, VLAN Interface, or SSID Interface, and select **Edit Physical Interface**, **Edit VLAN Interface**, or **Edit SSID Interface**. The Properties dialog opens.
4. Switch to the **DHCP** tab.

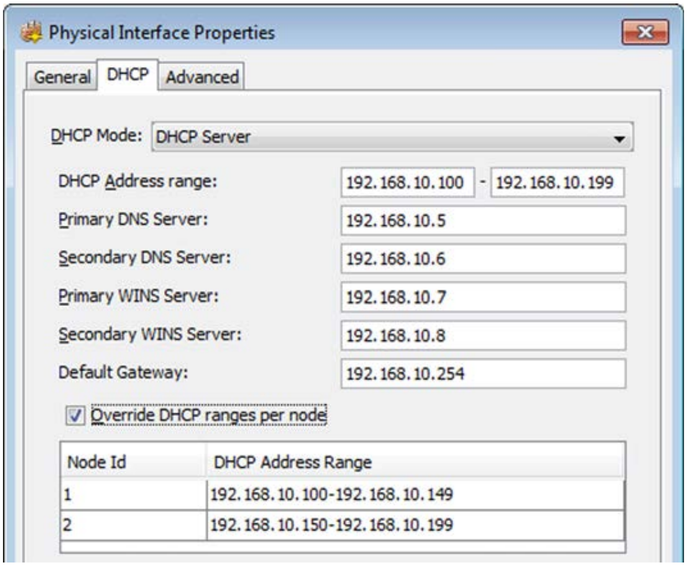

DHCP Mode: DHCP Server

DHCP Address range: 192.168.10.100 - 192.168.10.199

Primary DNS Server: 192.168.10.5

Secondary DNS Server: 192.168.10.6

Primary WINS Server: 192.168.10.7

Secondary WINS Server: 192.168.10.8

Default Gateway: 192.168.10.254

☑ Override DHCP ranges per node

Node Id	DHCP Address Range
1	192.168.10.100-192.168.10.149
2	192.168.10.150-192.168.10.199

5. Select **DHCP Server** as the **DHCP Mode**. The DHCP Server options appear.

6. Enter the beginning and the end addresses of the **DHCP Address Range** that the Firewall assigns to clients. On Firewall Clusters, the DHCP Address range is automatically divided between the nodes.

Note – The DHCP address range must be in the same network space defined for the Physical Interface. The DHCP address range must not contain the Firewall's NDI/CVI addresses or broadcast IP addresses of networks behind the Firewall.

7. (*Optional*) Enter the **Primary DNS Server** and **Secondary DNS Server** IP address that clients use to resolve domain names.
8. (*Optional*) Enter the **Primary WINS Server** and **Secondary WINS Server** IP address that clients use to resolve NetBIOS computer names.
9. Enter the IP address of the **Default Gateway** through which traffic from clients is routed.
10. (*Optional, clusters only*) Select **Override DHCP Ranges per Node** and enter the **DHCP Address Range** for each node.

Caution – Enter unique ranges for each node. Overlapping ranges can cause IP address duplication.

11. Click **OK** and repeat the steps for all interfaces on which you want to enable the DHCP server.
12. Allow the DHCP-related traffic in the Firewall's IPv4 Access rules:
    - Allow DHCP broadcast messages. Use the **BOOTPS (UDP)** Service element.
    - Allow messages from the DHCP address range defined in Step 6 to the IP address(es) that the Firewall uses towards the clients. Use the **BOOTPS (UDP)** Service element.
    - Allow messages from the IP address(es) that the Firewall uses towards the clients to the DHCP address range defined in Step 6. Use the **BOOTPC (UDP)** Service element.

What's Next?

- Refresh the policy to transfer the new configuration to the engines.

Related Tasks

- To relay clients' DHCP messages to/from an external DHCP server, see Routing DHCP Messages (page 590).
- Editing Access Rules (page 672)

Chapter 29

Connecting Engines to the Management Center

To maintain the security of your system, the engines establish an authenticated and encrypted connection with Log Servers and Management Servers.

The following sections are included:

Getting Started with Connecting Engines to the SMC

Prerequisites: See Getting Started with Engine Elements

What Connecting Engines to the SMC Does

When you connect the engines to the SMC, the engines make initial contact with the Management Server and receive a certificate. The certificate allows the engine to authenticate itself to other components in all further communications. When components contact each other, they check if the other component's certificate is signed by the same Internal Certificate Authority as their own certificate. The certificate authority runs on the Management Server, but is separate from the Management Server itself. The initial contact procedure is secured using a one-time password.

If the engines in your system are Stonesoft appliances, you can connect them to the SMC using the plug-and-play configuration method. In plug-and-play configuration, you upload the initial configuration to the Stonesoft Installation Server. When the engines are powered on with all the cables connected, they download the initial configuration from the Installation Server. After this, the engines automatically install the initial configuration and make initial contact with the Management Server. You can also specify a predefined policy to be installed on the engines when they make initial contact with the Management Server. For information on how to view the status of the configuration process, see Viewing Appliance Configuration Status (page 98).

If the certificate of the engine is lost or expires, the initial contact procedure must be repeated to reconnect the engine to the other components.

Limitations

- The plug-and-play configuration method is only available for Stonesoft appliances. You must have a valid proof-of-serial (POS) code for each appliance you want to configure using the plug-and-play configuration method.
- Virtual Security Engines do not communicate directly with the SMC. All communication between Virtual Security Engines and the SMC is proxied by the Master Engine.

What Should I Know Before I Begin?

- Engine certificates expire in three years from the date that they are issued. If the automatic certificate renewal option is active for the engine, the certificate is renewed automatically before it expires.
- The internal certificate authority that signs the engine certificates is valid for ten years. The internal certificate authority is automatically renewed six months before the expiration date and new certificates signed by the new internal certificate authority are automatically created for the engines. If the automatic certificate renewal fails, you must again make initial contact with the Management Server so that the engine receives a new certificate.
- When a new Internal Certificate Authority is created, its initial status is Ready to Use and it is not yet Active. A new Internal Certificate Authority in a Ready to Use state only signs Management Server certificates. Certificates for other system components are signed by the Internal Certificate Authority that is currently used by the Management Server. In an environment with multiple Management Servers, the new Internal Certificate Authority reaches Active status when all of the Management Servers are using the new Internal Certificate Authority.

Configuration Overview

1. Save an initial configuration for the engine on the Management Server. This triggers the creation of a one-time password that the engine can use to log in to the Management Server.
2. Make the engine contact the Management Server.
   - For Firewall, IPS, and Layer 2 Firewall engines, see Reconfiguring Basic Engine Settings (page 233).
   - For SSL VPN gateways, see Connecting SSL VPN Gateways to the SMC (page 518).
3. (*Not applicable to SSL VPN gateways*) Install a policy on the engine to transfer the full working configuration from the Management Server to the engine.
   - For Firewall engines, IPS engines, Layer 2 Firewall engines, and Master Engines see Installing Policies (page 652).

What's Next?

▶ Saving an Initial Configuration for Security Engines

▶ Connecting SSL VPN Gateways to the SMC (page 518)

Saving an Initial Configuration for Security Engines

Prerequisites: See Getting Started with Engine Elements

This operation allows you to establish a management connection for new Firewall engines, IPS engines, Layer 2 Firewall engines, and Master Engines for the first time. It also allows you to reconnect previously configured engines that have lost the connection due to a missing or expired certificate or because the internal certificate authority that signs the engine certificates has been renewed and the engines have not yet received a new certificate signed by the new internal certificate authority.

If you are installing a new engine or want to replace the engine's previous working configuration, this operation also allows you to save relevant parts of the configuration on a USB memory stick and import it during the engine installation.

What's Next?

▶ Start by Creating One-Time Passwords (page 516).

Creating One-Time Passwords

The one-time password that is created is specific to each engine. Keep track of which engine has which password. If you mix them up or lose them, you can repeat the procedure and create new initial configurations for those engines.

▼ To create a one-time password

↪ Save the initial configuration information:

- For an individual engine: right-click the bottom-level node element and select **Save Initial Configuration**.
- For all engines in a cluster: right-click the top-level cluster element and select **Configuration→Save Initial Configuration**.

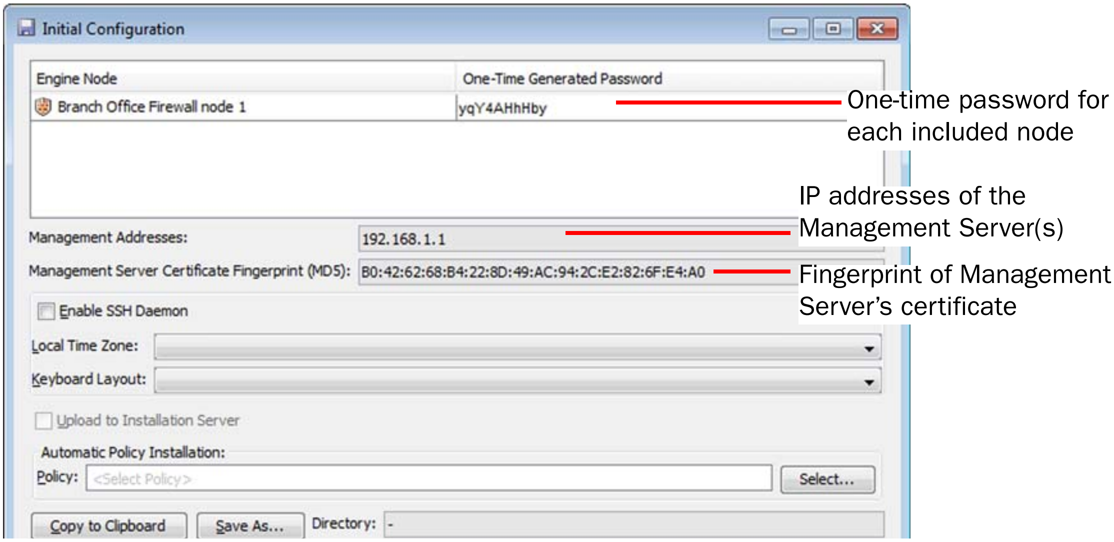

Initial Configuration

Engine Node	One-Time Generated Password
Branch Office Firewall node 1	yqY4AHhHby

Management Addresses: 192.168.1.1

Management Server Certificate Fingerprint (MD5): B0:42:62:68:B4:22:8D:49:AC:94:2C:E2:82:6F:E4:A0

□ Enable SSH Daemon

Local Time Zone:

Keyboard Layout:

□ Upload to Installation Server

Automatic Policy Installation:

Policy: <Select Policy>

Directory: -

One-time password for each included node

IP addresses of the Management Server(s)

Fingerprint of Management Server's certificate

The **One-Time Generated Password** and the **Management SSL Fingerprint** are displayed. The one-time password is mandatory and is used to authenticate the engine to the Management Server. The fingerprint can also be optionally used to make the engine verify the identity of the Management Server.

> **Note – If there is a Firewall between the engine and the Management Server, you must allow the connection in the Firewall's IPv4 Access rules. If there is a NAT device between the engine and the Management Server, you must also configure NAT rules for the connection. Otherwise, the engine cannot contact the Management Server.**

What's Next?

- ▶ If you want to manually enter details in the Engine Configuration Wizard or if the engine already has the correct configuration, take note of the one-time password and use it for initial contact from the engine (you can copy the password through the right-click menu).
- ▶ Otherwise, continue by Saving Initial Configuration Details (page 517).

Saving Initial Configuration Details

You can save the initial configuration details by uploading them to the Installation Server or saving them on a USB stick. Uploading the initial configuration details to the Installation Server is the simplest configuration method. However, this method can only be used with Stonesoft appliances that support plug-and-play configuration.

Saving the initial configuration details on a USB stick allows automatic configuration by booting the appliance with the USB stick inserted. Alternatively, you can import the configuration details from a USB stick in the manual command line Configuration Wizard.

Caution – Because the initial configuration files include the one-time password for establishing the trust relationship between the Management Server and the engine, the information must be handled securely when saving the initial configuration details on a USB stick.

▼ To save initial configuration details

1. (*Optional*) Select **Enable SSH Daemon** to allow remote access to the engine command line. SSH may be helpful for remote troubleshooting.
   - You can enable and disable remote command line access to the engine at any time after management contact is established through the right-click menu of the engine. We recommend that you disable SSH access whenever it is not needed and that you make sure your Access rules allow SSH access to the engines from the administrators’ IP addresses only.
   - The Firewall Template, IPS Template, and Layer 2 Template policies do not allow these connections, but the temporary policy activated right after the engine’s initial configuration (in force until you install the working policy) allows SSH access from the Management Server’s IP address. Alternatively, you can upload a working policy to be installed on the engine after it has contacted the Management Server.

Caution – If you enable SSH, set the password for command line access after the initial configuration either through the Management Client or by logging in to the command line. When the password is not set, anyone with SSH access to the engine can set the password.

2. (*Optional*) Select the **Local Time Zone** and the **Keyboard Layout** for use on the command line. The time zone setting is only for displaying the time on the local console; the engines always use UTC (GMT) time internally. The clock is automatically synchronized to match the Management Server’s time.
3. *(Optional)* If you are using the plug-and-play configuration method, select **Upload to Installation Server** to upload the initial configuration automatically to the Stonesoft Installation Server.
4. *(Optional)* If you already have a policy you want to use for the engine, click **Select** and select the appropriate policy for **Automatic Policy Installation**.
5. *(Optional)* If you are using a USB stick for the initial configuration, save the configuration information:
   - Click **Save As** to save the initial configuration file.
   - Select **Copy to Clipboard** or write down the one-time password to manually enter the password during the engine configuration.
6. Click **OK**.

What's Next?

- If you selected to upload the initial configuration to the Installation Server, connect the cables and power on the appliance. The appliance contacts the Installation Server and downloads the initial configuration.
- To configure the appliance automatically using a USB stick, power on the engine with a USB stick inserted.
- Otherwise, power on the engine and import the configuration to the command line Configuration Wizard. See Reconfiguring Basic Engine Settings (page 233) for detailed instructions.

Related Tasks

- Viewing Appliance Configuration Status (page 98)

Connecting SSL VPN Gateways to the SMC

Prerequisites: Creating a New SSL VPN Gateway Element

To connect an SSL VPN Gateway to the SMC

1. Right-click the SSL VPN Gateway element and select **Configuration→Save Initial Configuration**. The Initial Configuration dialog opens and the **Management SSL Fingerprint** and the **One-Time Generated Password** (one for each node in a mirrored pair) are displayed.
   - The one-time password is mandatory and is used to authenticate the engine to the Management Server.
   - The fingerprint can also be optionally used to make the engine verify the identity of the Management Server.
2. Take note of the one-time password (you can also copy the password by right-clicking it and selecting **Copy**).
3. Log in to the SSL VPN Web Console and use the one-time password to perform initial contact through **System→Initial Contact**. In the case of a mirrored pair, do this for both nodes (using different passwords).

Note – If there is a Firewall between the SSL VPN gateway and the Management Server, you must allow the connection in the Firewall's IPv4 Access rules. If there is a NAT device between the SSL VPN gateway and the Management Server, you must also configure NAT rules for the connection. Otherwise, the SSL VPN gateway cannot contact the Management Server.

The SSL VPN gateway is now configured for monitoring through the Management Client and you can view its status right away. However, you must configure the SSL VPN appliance separately to send syslog messages to the Log Server to be able to view logs from the SSL VPN appliance.

CHAPTER 30

CONFIGURING THE ENGINE TESTER

The tester runs various checks on the engines and initiates responses based on the success or failure of these tests.

The following sections are included:

Getting Started with the Engine Tester

Prerequisites: Creating New Engine Elements / Modifying Existing Engine Elements

The engines can be configured with periodical self-tests to ensure proper functioning of each Firewall engine, IPS engine, Layer 2 Firewall engine, and Master Engine.

What the Tester Does

The tester performs checks at certain intervals depending on the state of the engine (online/offline). Depending on the result, the basic actions the tester can perform are turning the engine online/offline and/or sending alerts and/or SNMP traps.

Limitations

- The engine tester also runs internal system tests that cannot be edited or disabled. These are meant for recognizing certain configuration problems and internal error conditions, such as nodes in the cluster having different policies.
- Tests cannot be run on Virtual Security Engines

Configuration Overview

1. Configure the global tester settings that are common for all tests. See Specifying Global Engine Tester Settings (page 521).
2. Configure the test entries (settings for running individual tests). See Adding Engine Tests (page 522).

What's Next?

▶ Proceed to Specifying Global Engine Tester Settings (page 521).

Related Tasks

▶ Removing Engine Tests (page 527)

▶ Disabling/Enabling Configured Engine Tests (page 528)

Specifying Global Engine Tester Settings

Prerequisites: None

The global settings of the tester have default values that you can override to meet your needs.

▼ To configure the global tester settings

1. Right-click the Firewall, IPS engine, Layer 2 Firewall, or Master Engine element, and select **Properties**. The Properties dialog opens.

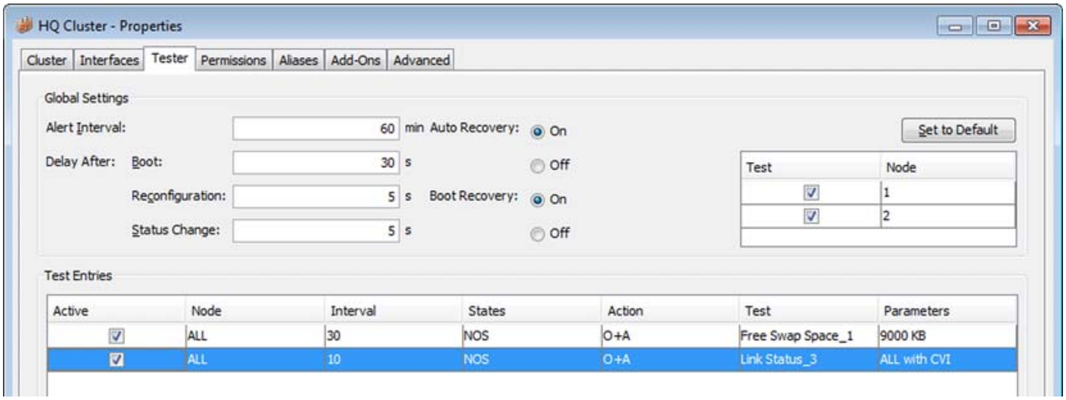

Global Settings

Alert Interval: 60 min Auto Recovery: ● On ○ Off

Delay After: Boot: 30 s

Reconfiguration: 5 s Boot Recovery: ● On ○ Off

Status Change: 5 s

Test	Node
☑	1
☑	2

Test Entries

Active	Node	Interval	States	Action	Test	Parameters
☑	ALL	30	NOS	O+A	Free Swap Space_1	9000 KB
☑	ALL	10	NOS	O+A	Link Status_3	ALL with CVI

2. Switch to the **Tester** tab. The global settings of the tester and the currently configured test entries are displayed.
3. Configure the settings as explained below:

Setting	Configuration
Alert Interval	Enter the time in minutes the system waits before sending a new alert when the same test keeps failing repeatedly. The default value is 60 minutes. **Note!** If the interval is too short, the alerts can potentially overload the system or the alert recipient.
Delay After	Enter the time in seconds that the engine waits before it resumes running the tests after the listed events. The delays prevent false test failures that can occur due to variations in how quickly different processes and subsystems can start and stop. The default values are event-specific: **Boot**: the default is 30 seconds (maximum: 1800 s) **Reconfiguration**: the default is 5 seconds (maximum: 1800 s) **Status Change**: the default is 5 seconds (maximum: 1800 s)
Auto Recovery (*Optional*)	Select **On** if you want the engine to automatically go back online when a previously failed test completes successfully. **Note!** Make sure to run the test in both online and offline states if you activate this option.
Boot Recovery (*Optional*)	Select **On** if you want the engine to automatically go back online after a reboot (or events such as a power failure or system crash) if all offline tests report a success.

What's Next?

▶ Configure the test entry settings as explained in Adding Engine Tests (page 522).

Adding Engine Tests

Prerequisites: Specifying Global Engine Tester Settings

Some settings are nearly identical for all tests. These common settings are explained below. Some tests have test-specific settings in addition to the common settings, which are explained separately for each test (see listing below the steps).

Test failures can be notified to administrators as Alerts or as SNMP traps (requires that SNMP is configured as explained in Getting Started with SNMP Configuration (page 582)). Additionally, a test can switch nodes offline/online based on the result.

▼ To specify the common settings

1. Open the properties dialog of the engine and switch to the **Tester** tab.
2. Click the **Add** button under the Test Entries table. The Test Entry Settings dialog opens.

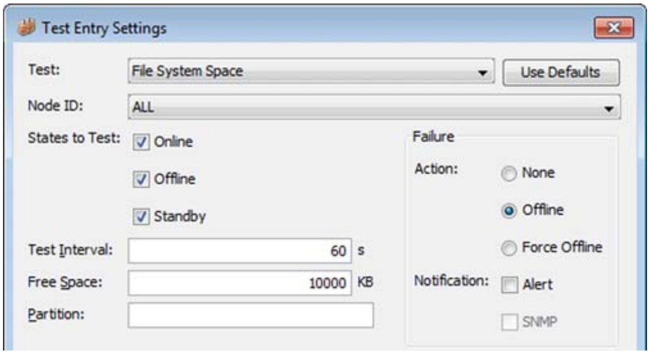

Test: File System Space

Node ID: ALL

States to Test: ☑ Online ☑ Offline ☑ Standby

Test Interval: 60 s

Free Space: 10000 KB

Partition:

Failure

Action: ○ None ● Offline ○ Force Offline

Notification: □ Alert □ SNMP

3. Select the type of **Test** you want to configure. The following tests are available:

Table 30.1 Tests

Test	Description
External	Runs a custom script stored on the engine. If the script returns the code zero (0), the test is considered successful, otherwise the test is considered failed.
File System Space	Tests the free disk space on a hard disk partition. **Note!** Do not use this test on appliances that use a flash card instead of a hard disk.
Policy	Checks whether a new policy is activated on the engine. It is mainly intended for sending SNMP notifications.
Free Swap Space	Checks the available swap space on the hard disk. **Note!** Do not use this test on appliances that use a flash card instead of a hard disk.

Table 30.1 Tests (Continued)

Test	Description
Link Status	Checks whether a network port reports the link as up or down. - The link of a Modem Interface (*Single Firewalls only*) is up if a modem reports that a call is in progress. - The link of an ADSL interface (*Single Firewalls only*) is up when the ADSL modem is connected to a Digital Subscriber Line Access Multiplexer (DSLM), and the ADSL modem reports that the connection is working. - The link of a Wireless Interface (*Single Firewalls only*) is up when the related firewall engine has been configured and is working. - The status of Aggregated Links (*Firewalls only*) depends on the link mode. An Aggregated Link in High-Availability Mode is up if one of the interfaces that belong to the Aggregated Link is up. An Aggregated Link in Load-Balancing Mode is up if both the interfaces that belong to the Aggregated Link are up.
Multiping	Sends out a series of ping requests to determine whether there is connectivity through a network link. Only IPv4 addresses are supported as target addresses for the Multiping test.
Inline Pair Link Speed (*Security Engines with inline interfaces only*)	Checks whether the network settings (speed/duplex) match on the two ports that form the inline pair and can force both ports to use the same settings.

4. Configure the other settings as explained below:

Table 30.2 Test Settings

Setting	Configuration
Node ID (*clusters only*)	Select whether to run the test on **ALL** nodes or only on a specific node.
States to Test	Select one or more engine states in which to run the test: - **Online**: test is run when node is online. - **Offline**: test is run when node is offline. - **Standby** (*clusters only*): test is run when node is in the Standby state.
Test Interval	Enter the time in seconds to specify how often the test is run. The minimum interval is 4 seconds and the maximum is 86,400 seconds (one day). **Note!** Running a test too frequently can increase overhead.

Table 30.2 Test Settings (Continued)

Setting		Configuration
Failure	Action	Select the action the tester takes if a test fails: - **None**: no state transition is performed. - **Offline**: the test switches the node offline. However, the tester does not turn offline the last active node in a cluster or nodes in the Locked Online state. - **Force Offline**: the node is forced to an offline state, even if it is the last node online in a cluster or if it is in the Locked Online state. Use in cases in which a complete cut in traffic is a better option than a partially working Firewall. - **Force Speed** (*Inline Pair Link Speed test only*): if the test finds that the pair of inline ports have different speed/duplex settings, the speed/duplex on the higher speed/full duplex link is set to match the lower speed/half duplex setting on the other port. Correct and reliable operation requires identical settings.
	Notification	Select how to notify administrators that a test has failed. - **Alert**: test failure triggers an alert. - **SNMP**: an SNMP trap is sent. To use SNMP notification, SNMP must be configured for this engine. See Getting Started with SNMP Configuration (page 582).

What's Next?

▶ Proceed to the next revelant section for the type of test you are configuring:

- Configuring Additional Settings for the External Test.
- Configuring Additional Settings for the File System Space Test (page 525).
- Configuring Additional Settings for the Free Swap Space Test (page 525).
- Configuring Additional Settings for the Link Status Test (page 525).
- Configuring Additional Settings for the Multiping Test (page 526).

▶ Otherwise, click **OK** and refresh the policy to transfer the new configuration to the engine(s).

Configuring Additional Test-Specific Settings

Configuring Additional Settings for the External Test

▼ To define settings for the External test

1. Enter the **Retry Count** to configure the number of times the Tester tries to execute the test.
   - We recommend always setting the retry count to more than 1 to avoid creating overly sensitive tests that burden the system unnecessarily.
2. Enter **Test Timeout** in milliseconds.
   - If the test being run does not return a response in the specified time, the test is considered to have failed.
   - Avoid overly short timeout values. We recommend a timeout of 500 - 1000 ms depending on the external test script.
3. Enter the script path in the **Command Line** field.

Example `/usr/local/bin/connectivity.sh`

4. Click **OK**.

> **Note – The external script must return an exit code of 0 (zero) if it succeeds. Any non-zero return value is considered a failure.**

What's Next?

- ▶ Refresh the policy to transfer the new configuration to the engine(s).

Configuring Additional Settings for the File System Space Test

▼ To define settings for the File System Space test

1. Enter the minimum amount of **Free Space** in kilobytes. When the amount of free space drops below this amount, the engine executes the chosen action.
2. Specify the partition in the **Partition** field.
3. Click **OK**.

What's Next?

- ▶ Refresh the policy to transfer the new configuration to the engine(s).

Configuring Additional Settings for the Free Swap Space Test

▼ To define settings for the Free Swap Space test

1. Enter the minimum amount of **Free Space** in kilobytes. When the amount of free space drops below this amount, the engine executes the chosen action.
2. Click **OK**.

What's Next?

- ▶ Refresh the policy to transfer the new configuration to the engine(s).

Configuring Additional Settings for the Link Status Test

▼ To define settings for the Link Status test

1. Select the **Interface** on which the test is run:
   - **ALL** Physical Interfaces, Modem Interfaces (*Single Firewalls only*), ADSL Interfaces (*Single Firewalls only*), and Wireless Interfaces (*Single Firewalls only*).
   - (*Clusters only*) **ALL with CVI**.
   - A single Physical Interface, a Modem Interface (*Single Firewalls only*), an ADSL Interface (*single firewalls only*), or a Wireless Interface (*Single Firewalls only*).

> **Note – (*Firewalls only*) Only the first interface that belongs to an Aggregated Link is shown in the list of interfaces. However, the Link Status test checks the status of both interfaces in the Aggregated Link.**

2. Click **OK**.

What’s Next?

▶ Refresh the policy to transfer the new configuration to the engine(s).

Configuring Additional Settings for the Multiping Test

▼ To define settings for the Multiping test

1. Select the **Source Address** for the test:
   - **DEFAULT**: no source address is forced on the test. The source IP address is selected automatically based on the standard routing rules.
     On a cluster, if the Physical Interface that routes the ping packet out has an NDI (Node Dedicated IP Address), this address is used as the source address. Otherwise, the NDI selected as Default IP for Outgoing Traffic is used.
   - A single Physical Interface, a VLAN Interface, a Modem Interface (*Single Firewalls only*), an ADSL Interface (*single firewalls only*), or an SSID Interface (*Single Firewalls only*). If the **Node ID** selection is **ALL**, each node uses the IP address of the selected interface as the source IP address for the test. If a single node is selected in **Node ID** on a cluster, the source address and the multiping test itself apply to that node only.
2. Enter the **Retry Count** to configure the number of times the Tester tries to execute the test. We recommend always setting the retry count to more than 1 to avoid creating overly sensitive tests that burden the system unnecessarily.
3. Enter **Test Timeout** in milliseconds. If the test being run does not return a response in the specified time, the test is considered to have failed. Avoid overly short timeout values.
4. Specify the **Target Addresses** of ICMP echo requests. Only IPv4 addresses are supported.
   - You can add or remove target addresses with the corresponding buttons.
   - The test is considered failed if none of the target addresses responds.
5. Click **OK**.

What’s Next?

▶ Refresh the policy to transfer the new configuration to the engine(s).

Checking Configured Tests

Prerequisites: Adding Engine Tests

The test entries that you have configured with the tester are displayed in the Test Entries table in the engine properties. You can deactivate, activate, and remove tests directly in the Test Entry table.

▼ To check configured tests

1. Right-click the Firewall, IPS engine, Layer 2 Firewall, or Master Engine element and select **Properties**. The Properties dialog opens.
2. Switch to the **Tester** tab. All configured test entries are displayed in the Test Entries table. The selected states and actions for the tests are indicated with letters as summarized in the table below.

Table 30.3 Test Entry Display of Selected Actions and States

Types	Table Display	Description
States to test	N	Online
	O	Offline
	S	Standby
Action on failure	A	Alert
	O+A	Offline + Alert
	FO+A	Force Offline + Alert
	A+S	Alert + SNMP
	O+A+S	Offline + Alert + SNMP
	FO+A+S	Force Offline + Alert + SNMP

What's Next?

- If you want to remove a test permanently, proceed to Removing Engine Tests.
- If you only want to activate or reactivate a test, proceed to Disabling/Enabling Configured Engine Tests (page 528).
- Otherwise, click **OK** to close the engine properties.

Removing Engine Tests

Prerequisites: None

Removing a test permanently removes the test settings from the selected engine. If you want to stop a test from running without removing it, see Disabling/Enabling Configured Engine Tests (page 528).

▼ To remove a test

1. Right-click the Firewall, IPS engine, Layer 2 Firewall, or Master Engine element and select **Properties**. The Properties dialog opens.
2. Switch to the **Tester** tab.
3. Select the test entry you want to remove in the **Test Entries** list.
4. Click **Remove**. The test is permanently removed from the selected engine.
5. Click **OK**.

What's Next?

- Refresh the policy to transfer the new configuration to the engine(s).

Disabling/Enabling Configured Engine Tests

Prerequisites: There are one or more enabled/disabled user-configurable tests configured

All test entries that you configure are automatically activated and available to the tester after a policy refresh or install. You can decide to disable any individual test in the Test Entries table or all of the tests that are applied on a specified node.

What's Next?

▶ Disabling/Enabling Individual Engine Tests

▶ Disabling/Enabling All Custom Engine Tests

Disabling/Enabling Individual Engine Tests

▼ To disable or enable an individual test

1. Right-click the Firewall, IPS engine, Layer 2 Firewall, or Master Engine element and select **Properties**. The Properties dialog opens.
2. Switch to the **Tester** tab.
3. Find the test you want to disable/enable in the Test Entries table and deselect/select the option in the **Active** column to disable/enable the test.
4. Click **OK**.

What's Next?

▶ Refresh the policy to transfer the new configuration to the engine(s).

Disabling/Enabling All Custom Engine Tests

▼ To disable or enable all customized tests on a specific node

1. Right-click the Firewall, IPS engine, Layer 2 Firewall, or Master Engine element and select **Properties**. The Properties dialog opens.
2. Switch to the **Tester** tab.
3. In the Global Settings section, deselect/select the option in the **Test** column of the node(s) for which you want to disable/enable testing.
4. Click **OK**.

What's Next?

▶ Refresh the policy to transfer the new configuration to the engine(s).

CHAPTER 31

ENGINE PERMISSIONS

You can set permissions to control the administration of the engines.

The following sections are included:

Getting Started with Engine Permissions

Prerequisites: Your administrator account must have editing privileges to the engine element

What Engine Permission Control Does

Engine permissions control can be used in two ways:

- To prevent some administrators from editing and viewing an engine's properties to prevent unauthorized modifications and protect confidential details.
- To restrict the policies that can be installed on the engine to prevent service outages caused by the wrong policy being installed on the engine by accident.

What Do I Need to Know Before I Begin?

- Permissions for Master Engines and Virtual Security Engines are configured separately. Otherwise, engine permissions are configured for Master Engines and Virtual Security Engines in the same way as for other types of engines.
- See Getting Started with Administrator Accounts (page 248) for more information about access control in the Stonesoft Management Center.

Configuration Overview

1. Define which administrators are allowed to edit and view the element. See Defining Administrator Permissions on Engines.
2. Define which policies can be installed on the engine. See Selecting Permitted Policies for Engines (page 531).

Defining Administrator Permissions on Engines

Prerequisites: Your administrator account must have editing privileges to the engine element

You can either add an Access Control List or an individual Administrator-Administrator Role pair as permitted on the engine. The rights that the Access Control List grants to the administrators are defined in the properties of the administrator accounts (defined with Administrator elements).

An administrator with a restricted account can refresh or upload a policy on an engine only if both the engine and the policy have the administrator on the list of permitted administrators. The engines may not accept all policies. This can be adjusted as explained in Selecting Permitted Policies for Engines (page 531).

▼ To add or modify administrator permissions on an engine

1. Right-click the engine or policy and select **Properties**. Switch to the **Permissions** tab.
2. To add the engine to an Access Control List, click **Add** for the **Access Control Lists** and select the Access Control List(s) to which you want the engine to belong.
   - If you want to create a new Access Control List, click the **New** icon in the Select Element dialog and proceed as explained in Defining Access Control Lists (page 252).
3. To add a permission, click **Add Permission** under Permissions. A new row appears on the administrator list.
4. Click the **Administrator** cell and select the Administrator.
5. Right-click the **Administrator Role** cell and select **Edit Administrator Role**. The Select Element(s) dialog opens.

6. Select a role and click **Add**.
7. Click **OK** to close the Select Element(s) dialog.
8. If necessary, repeat from Step 3 to define other administrator permissions.
9. Click **OK**.

Note – Changes to administrator permissions are immediately distributed and taken into account in all related elements.

Related Tasks

- Defining Administrator Roles and Access Control Lists (page 249)

Selecting Permitted Policies for Engines

Prerequisites: Your administrator account must have editing privileges to the engine element

By default, any policy can be installed on any component as long as the policy is of appropriate type for the type of engine. To prevent accidental installations of the wrong policy, you can select the allowed policy for each engine in the engine element's properties. The policy selection is enforced regardless of the administrator privileges that the installing administrator has.

▼ To grant a policy to be installed on an engine

1. Right-click the correct engine element and select **Properties**. The properties dialog for the element opens.
2. Switch to the **Permissions** tab.
3. In the Policies section at the bottom, select the Policy:
   - To allow the installation of any policy, click **Set to Any**.
   - Otherwise, click **Add**. The Select Element dialog opens.
4. Select the correct Policy or Template Policy.
   - If you select a Template Policy, any policy based on the template can be installed.
5. Click **OK**.

CHAPTER 32

ALIAS TRANSLATIONS FOR ENGINES

Aliases can be used to represent other network elements in configurations. The value an Alias takes in a configuration can be different on each engine where the Alias is used.

The following sections are included:

Getting Started with Alias Translations

Prerequisites: Creating New Engine Elements / Modifying Existing Engine Elements

What Aliases Do

You can use Alias elements in your policies to represent IP addresses. Aliases differ from Group elements in that they do not represent all the elements at once. The Alias elements do not contain any IP address information themselves. The values that the Aliases receive when they are encountered in a policy depend on the translation value you set for each Alias in the engine elements' properties. This way, the same policy can be used on several engines, and the IP address information is filled in correctly according to the translation values for each engine. Alias elements are especially useful in policy templates and sub-policies.

What Do I Need to Know Before I Begin?

Aliases are configured for Security Engines. The Security Engines use the translated values in the data they send to the Log Server.

If you use Master Engines and Virtual Security Engines, you must configure Aliases separately for the Master Engines and for the Virtual Security Engines.

There are some predefined Aliases provided by the system. Some system Aliases cannot be modified. The names of the unmodifiable system Aliases start with two $$ symbols. The unmodifiable system Aliases receive their translation values automatically based on the engine's configuration. You can also create new Aliases. User-modifiable and user-created Aliases start with one $ symbol. User-modifiable default Aliases do not receive their translation values automatically. For a listing of predefined Aliases available in the system, see Predefined Aliases (page 1175).

Defining Alias Translation Values

Prerequisites: Creating New Engine Elements / Modifying Existing Engine Elements

Alias translation values are set for each engine. You can edit the translation values through the engine properties or in the properties of the Alias element itself.

Adding Alias Translation Values

▼ To add elements to Alias translation values

1. Right-click an engine element and select **Properties**. The Properties dialog opens.
2. Switch to the **Aliases** tab. The table with Aliases and their values opens.
3. Select the Alias whose value you want to edit.
4. Define the value for the Alias:
   - If you want the Alias to match any IP address on this engine, right-click the **Value** cell for the Alias and select **Set to Any**. Skip to Step 8.
   - Otherwise, right-click the **Value** cell and select **Edit Value**. The Alias Value Properties dialog opens.
5. Browse in the Resources panel for an element you want to use as the translation value on this engine.
   - If the element does not exist yet, you can create a new one through the right-click menu for the correct element type.

6. Select the correct element(s) and click **Add**. The element(s) are added to the Alias Value panel.
7. Click **OK** to close the Alias Value Properties dialog.
8. Click **OK** to close the engine properties.

Removing Alias Translation Values

If there are no translation values for a particular Alias for an engine, the Translation Value is **None**. If Aliases with the value **None** are used alone as matching criteria in a rule in a Firewall, IPS, or Layer 2 Firewall Policy, the rule never matches any traffic.

▼ To remove Alias translation values

1. Right-click an engine element and select **Properties**. The Properties dialog opens.
2. Switch to the **Aliases** tab. The table with Aliases and their values opens.
3. Right-click the **Value** cell of the Alias whose value you want to modify and select **Edit Value**. The Alias Value Properties dialog opens.
4. Select the element(s) to be removed in the Alias Value panel and click **Remove**.
5. Click **OK** to close the Alias Value Properties dialog.
6. Click **OK** to close the engine properties.

CHAPTER 33

ADD-ON FEATURES

This section explains the settings on the Add-Ons tab in the properties of Firewall, IPS, Layer 2 Firewall, and Virtual Firewall elements.

The following sections are included:

Getting Started with Add-On Features

Prerequisites: Creating and Modifying Engine Elements

There are several optional add-ons that you can enable on the Add-Ons tab in Firewall, IPS, Layer 2 Firewall, and Virtual Firewall properties. Most add-ons are separately licensed features.

Table 33.1 Add-Ons

Add-On	Supported On	See
TLS Inspection	Firewalls, IPS engines, Layer 2 Firewalls, Virtual Firewalls	Getting Started with TLS inspection (page 794)
User Agent	Firewalls, IPS engines, Layer 2 Firewalls, Virtual Firewalls	Enabling Access Control by User (page 838)
User Authentication	Firewalls, Virtual Firewalls	Enabling Browser-Based User Authentication (page 874)
Anti-virus	Firewalls	Configuring Anti-Virus Settings (page 540)
Anti-spam	Firewalls	Configuring Anti-Spam Settings (page 541)

Add-ons are not supported on Master Engines.

What’s Next?

▶ Continue by Editing Add-On Settings.

Editing Add-On Settings

Prerequisites: Getting Started with Add-On Features

▼ To edit add-ons

1. Right-click an engine element and select **Properties**. The element's Properties dialog opens.

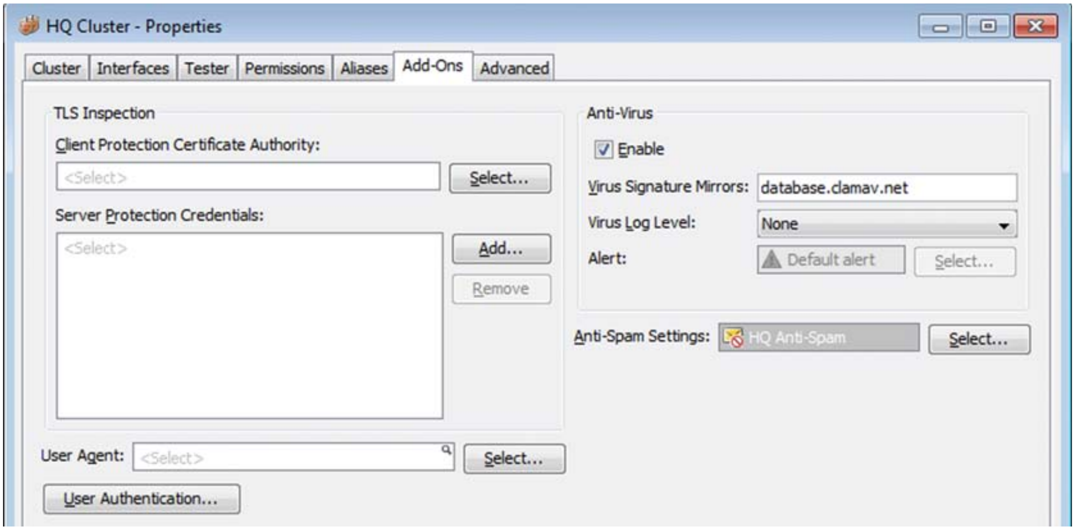

HQ Cluster - Properties

TLS Inspection

Client Protection Certificate Authority: <Select>

Server Protection Credentials: <Select>

User Agent: <Select>

Anti-Virus

☑ Enable

Virus Signature Mirrors: database.clamav.net

Virus Log Level: None

Alert: Default alert

Anti-Spam Settings: HQ Anti-Spam

2. Switch to the **Add-Ons** tab and adjust the settings as explained in the sections listed below.
   - Getting Started with TLS inspection (page 794)
   - Enabling Access Control by User (page 838)
   - (*Firewalls and Virtual Firewalls Only*) Enabling Browser-Based User Authentication (page 874)
   - (*Firewalls Only*) Configuring Anti-Virus Settings
   - (*Firewalls Only*) Configuring Anti-Spam Settings

Configuring Anti-Virus Settings

Prerequisites: Creating a New Single Firewall Element / Creating a New Firewall Cluster Element

Anti-virus is available for Firewalls as a separately licensed feature on selected platforms. Only IPv4 traffic can be inspected for viruses. The IPv4 Access rules in the Firewall Policy determine which traffic is inspected for viruses (see Editing Access Rules (page 672)).

▼ To configure anti-virus settings

1. Select **Enable** in the Anti-Virus section.
2. (*Optional*) Enter the URL of the **Virus Signature Mirrors** that you want to use. The engines contact the mirrors to update their virus signatures. Separate multiple addresses with commas.
   - You must have at least one IP address defined in the **DNS IP Addresses** list on the Single Node or Cluster tab of the Firewall Properties.
   - Engines query the DNS servers to resolve the virus signature mirror URLs.
3. Select the **Virus Log Level**. The log levels are the same as in Access rules. See Defining Access Rule Logging Options (page 684).
4. (*Optional*) If you selected Alert as the Virus Log Level, click **Select** and select the Alert.
5. Click **OK**.

What's Next?

▶ If you have not yet defined which traffic is selected for anti-virus inspection, modify the IPv4 Access rules in the Firewall Policy. Proceed to Editing Access Rules (page 672).

▶ Otherwise, refresh the firewall's policy to transfer the changes.

Configuring Anti-Spam Settings

Prerequisites: Creating a New Single Firewall Element / Creating a New Firewall Cluster Element

Anti-spam is available for Firewalls as a separately licensed feature on selected platforms. The feature allows the Firewall to filter incoming spam e-mails. Only IPv4 addresses are supported in spam filtering. The IPv4 Access rules in the Firewall Policy determine which traffic is filtered for spam (see Editing Access Rules (page 672)).

It is not recommended to use spam filtering between your own mail servers. Modify the IPV4 Access rules to exclude traffic between your mail servers from spam filtering. For more information on editing Access Rules, see Editing Access Rules (page 672).

▼ To configure anti-spam settings

1. Click **Select** in the Anti-Spam section. The Select Element dialog opens.
2. Select the Anti-Spam Settings element from the list, or create a new Anti-Spam Settings element through the New icon at the top of the dialog.

What's Next?

- ▶ If you created a new Anti-Spam Settings element, continue by Defining General Anti-Spam Settings.
- ▶ Otherwise, click **OK** to close the Firewall Properties. Modify the IPv4 Access rules in the Firewall Policy as instructed in Editing Access Rules (page 672) to define which traffic is filtered for spam.

Defining General Anti-Spam Settings

General anti-spam settings include antispoofing and anti-relay settings, honeypot address settings, and settings for SPF (Sender Policy Framework) and MX (Mail Exchanger) record matching.

To enable antispoofing settings, you must specify your local network domains. If the header of an external e-mail contains your local domain information it is considered to be spam.

▼ To configure general anti-spam settings

1. Enter a **Name** for the Anti-Spam Settings element.

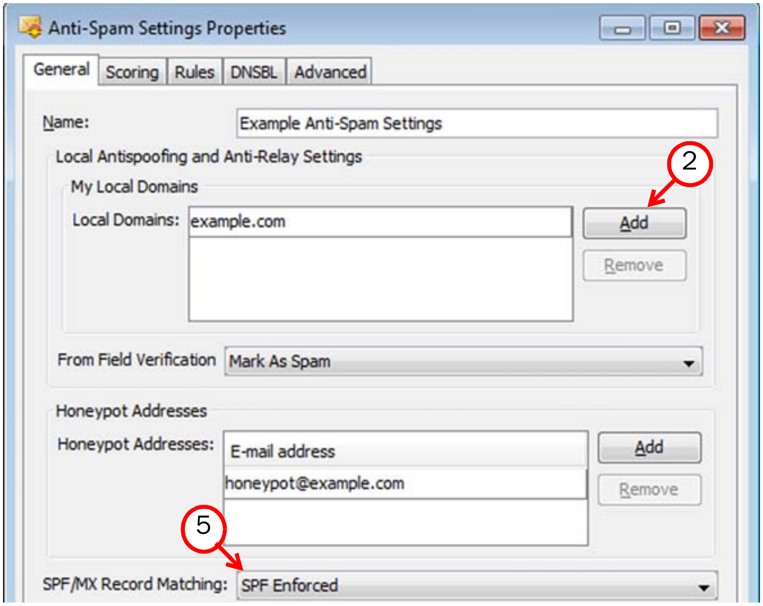

Name: Example Anti-Spam Settings

Local Domains: example.com

From Field Verification: Mark As Spam

Honeypot Addresses: honeypot@example.com

SPF/MX Record Matching: SPF Enforced

2. Click **Add** and enter your **Local Domain** (for example, example.com). The domain is added to the **Local Domains** list.
3. Select one of the following actions to apply to incoming e-mail that contains your local domain information:

Option	Description
Discard Message	The e-mail is discarded. The sender is not notified.
Mark As Spam	The message is allowed for further processing but a **Spam Label** is added to the subject header.
Use Message Only for Scoring	The message is allowed for further processing but it is assigned a spam score.

4. Click **Add** and enter a **Honeypot Address** (for example, honeypot@example.com).
   - IP addresses of all clients that send e-mails to honeypot addresses are blacklisted. If **Automatic Blacklisting** is enabled on the Advanced tab, all e-mails that come to honeypot addresses are blacklisted for the specified **Blacklist Time**.

5. Select the **SPF/MX Record Matching** settings:

Option	Description
Off	SPF/MX Record Matching is disabled.
Check	SPF/MX Records are checked and scored.
SPF Enforced	SPF Record Matching is enforced. This is the default setting.
SPF and MX Enforced	Both SPF and MX Record Matching are enforced.

- An SPF record is an entry in the DNS that specifies which hosts are allowed to send e-mails from certain domains. Mail exchangers use SPF records to check if an e-mail is sent from a legitimate host.
- An MX record specifies hosts that are responsible for handling e-mails for a certain domain.
- If SPF/MX record matching is enabled, the credibility of the messages that pass SPF/MX verification is increased and the credibility of the messages that do not pass the verification is decreased.

What's Next?

▶ If you want to define additional criteria for Anti-Spam filtering, proceed to the next relevant section below.

- Defining Scoring Settings for Anti-Spam (page 544)
- Defining Spam Filtering Rules (page 545)
- Defining DNSBL Settings (page 547)
- Modifying Advanced Anti-Spam Settings (page 548)

▶ If you are finished defining Anti-Spam settings and have not yet defined which traffic is filtered for spam, modify the IPv4 Access rules in the Firewall Policy as instructed in Editing Access Rules (page 672).

▶ Otherwise, the configuration is finished. Click **OK** and refresh the Firewall's policy to transfer the changes.

Defining Scoring Settings for Anti-Spam

Each incoming e-mail message is assigned a spam score to determine the likelihood of its being spam. By default, a spam label is added to the headers of all e-mails with a score of 2 and above, and all e-mails with a score of 8 and above are rejected. You can adjust the scoring settings to determine the score of messages that are marked as spam or that are rejected. You can fine-tune spam scoring in the individual spam rules. See Defining Spam Filtering Rules (page 545).

▼ To define scoring settings for anti-spam

1. Switch to the **Scoring** tab in the Anti-Spam Settings Properties.

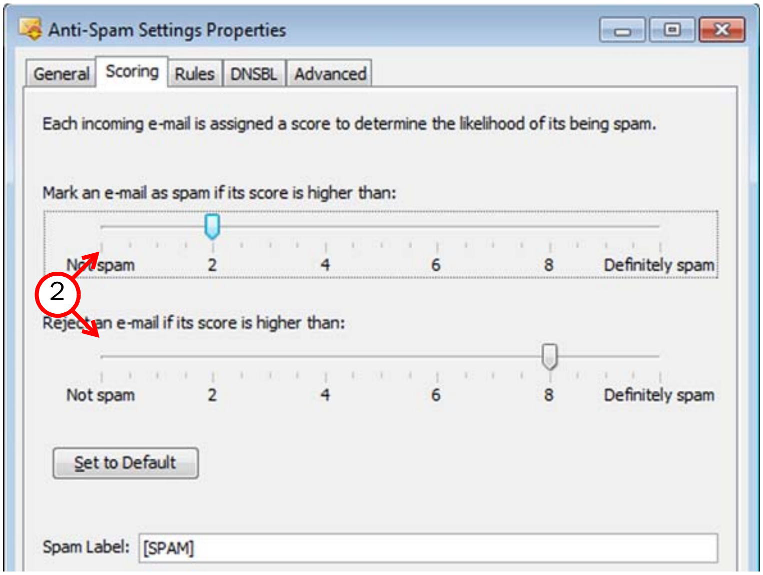

Each incoming e-mail is assigned a score to determine the likelihood of its being spam.

Mark an e-mail as spam if its score is higher than:

Reject an e-mail if its score is higher than:

Spam Label: [SPAM]

2. Drag the score markers to define when e-mails are marked as spam or rejected, or click **Set to Default** to apply the default settings.

3. *(Optional)* Modify the default **Spam Label** that is added to the subject header of an e-mail that is considered to be spam.

What's Next?

▶ If you want to define additional criteria for Anti-Spam filtering, proceed to the next relevant section below.

- Defining Spam Filtering Rules (page 545)
- Defining DNSBL Settings (page 547)
- Modifying Advanced Anti-Spam Settings (page 548)

▶ If you are finished defining Anti-Spam settings and have not yet defined which traffic is filtered for spam, modify the IPv4 Access rules in the Firewall Policy as instructed in Editing Access Rules (page 672).

▶ Otherwise, the configuration is finished. Click **OK** and refresh the Firewall's policy to transfer the changes.

Defining Spam Filtering Rules

You can define specific rules for checking the content of spam messages. Specific rules eliminate the need for further processing and prevent unnecessary use of resources.

▼ To define spam filtering rules

1. Switch to the **Rules** tab in the Anti-Spam Settings Properties.

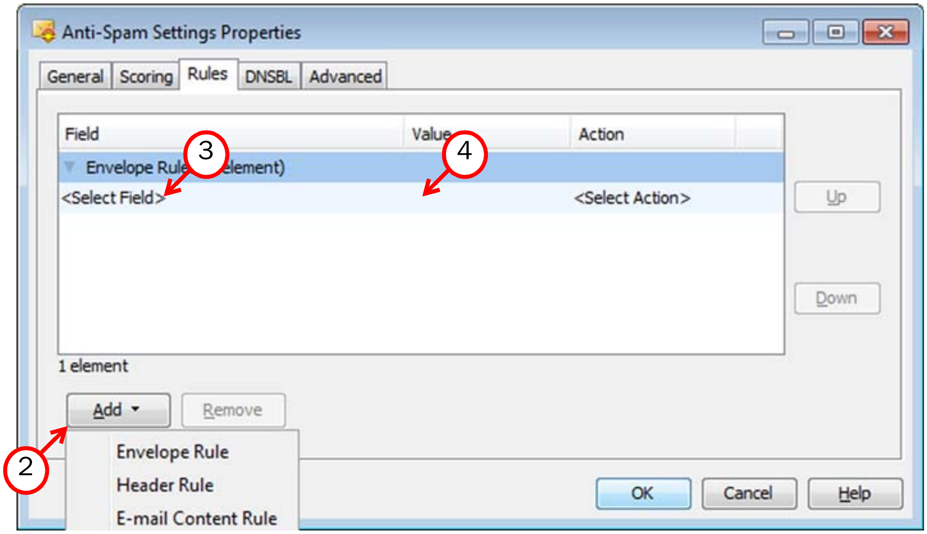

2. Click **Add** and select the type of rule:
   - **Envelope Rule:** the rule applies to the envelope of an e-mail.
   - **Header Rule:** the rule applies to the header of an e-mail.
   - **E-mail Content Rule:** the rule applies to the body of an e-mail.
3. *(Envelope and header rules only)* Click **<Select Field>** to select the envelope or header fields that the rule is applied to.
   - If the Header Field that you want to select is not listed, create a new Header Field by selecting **New** and enter a **Name** for the Header Field in the dialog that opens.

Note – An E-mail Content Rule is automatically applied to the content in the body of an e-mail.

4. Double-click the **Value** column. The Edit Value dialog opens.

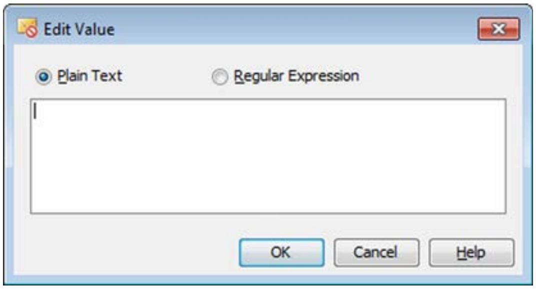

5. Select **Plain Text** or **Regular Expression** and enter the value as instructed below:

Setting	Configuration
Plain Text	Enter the text you want to match. The supported character sets are Latin, Russian, and Chinese.
Regular Expression	Enter the Regular Expression using the POSIX 1003.2 regular expression format. - The first character of an argument is interpreted as the delimiter, and all characters up to the next occurrence of the same delimiter are interpreted as the regular expression. - Any character except spaces and tabs can be used as the delimiter. However, character escaping is not supported in the regular expression. Use a character that does not occur in the regular expression as the delimiter. - Two delimiters with no arguments between them are interpreted as an empty regular expression. Empty regular expressions always match. This can be used to match only some of the arguments in rules that require multiple arguments. - Optionally, the following flags can be used after the closing delimiter: **e**: Interprets the regular expression as an extended regular expression. **i**: Ignores upper/lower case when evaluating the regular expression. **n**: Negates the regular expression. The rule matches when the regular expression does not match.

6. Click **<Select Action>** and select the **Action**.

Action	Description
Allow	The e-mail is allowed without further processing.
Discard	The e-mail is discarded. The sender is not notified.
Reject	The e-mail is rejected. The sender receives the following notification: Message undeliverable (SMTP 554 error).
Graylist	The e-mail is graylisted. The sender receives the following notification: Try again later (SMTP 440 error). If the message is retransmitted and does not fail further checks, it is allowed.
Blacklist	The message is rejected and the sender's IP address is blacklisted. Selecting this action opens a dialog where you can enter the Duration for which the sender is blacklisted.
Mark as Spam	The e-mail is allowed but a **Spam Label** is added to the subject header of the e-mail.
Change Score	Selecting this action opens a dialog where you can change the score of an e-mail. Select a negative score to increase the credibility of the e-mail and a positive score to decrease the credibility of the e-mail.

7. *(Optional)* Repeat Step 2 to Step 6 to add more rules.

What's Next?

- If you want to define additional criteria for Anti-Spam filtering, proceed to the next relevant section below.
  - Defining Scoring Settings for Anti-Spam (page 544)
  - Defining DNSBL Settings
  - Modifying Advanced Anti-Spam Settings (page 548)
- If you are finished defining Anti-Spam settings and have not yet defined which traffic is filtered for spam, modify the IPv4 Access rules in the Firewall Policy as instructed in Editing Access Rules (page 672).
- Otherwise, the configuration is finished. Click **OK** and refresh the Firewall's policy to transfer the changes.

Defining DNSBL Settings

DNS-based Blackhole Lists (DNSBLs) are used for publishing IP addresses and networks that are suspected of sending spam. You can configure the Firewall to reject messages sent from domain names and IP addresses that appear on DNSBLs.

▼ To define DNSBL settings

1. Switch to the **DNSBL** tab in the Anti-Spam Settings Properties.

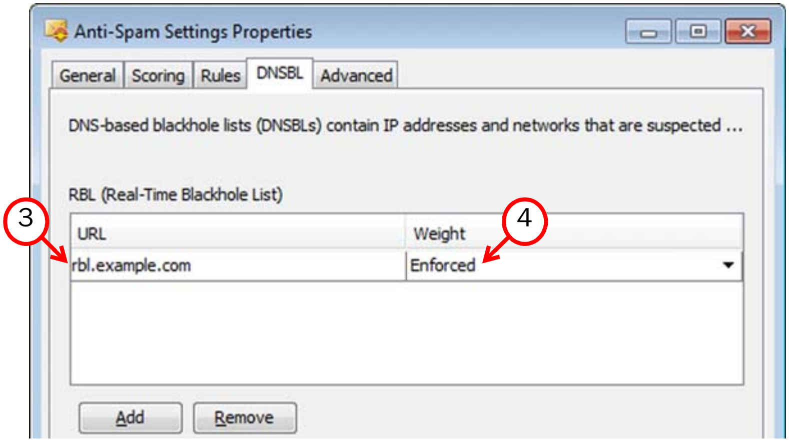

Anti-Spam Settings Properties

DNS-based blackhole lists (DNSBLs) contain IP addresses and networks that are suspected ...

RBL (Real-Time Blackhole List)

URL	Weight
rbl.example.com	Enforced

2. Click **Add** below the list that you want to modify:

Option	Description
RBL (Real-Time Blackhole List)	Includes URLs of RBLs that list IP addresses of servers responsible for sending spam or that are hijacked for spam relay.
URI DNSBL (Uniform Resource Identifier DNSBL)	Includes URLs of DNSBLs that list domain names and IP addresses of links found in the body of spam e-mails.

3. Double-click the empty field in the **URL** column and enter the URL of the RBL or DNSBL.

4. Select the **Weight** for the defined URL:
   - The weight values indicate how credible you consider the list.
   - The weight values are used to determine whether the message reaches the threshold value. The threshold value is a median of all responses received from the active lists. The messages that reach the threshold value are rejected.

Value	Description
Off	The list is not checked.
Low	The list has low credibility and may have a high false positive rate.
Medium	The list has medium credibility.
High	The list has high credibility.
Enforced	The list is fully trusted. If the sender is found on the list, the message is rejected.

What's Next?

▶ If you want to define additional criteria for Anti-Spam filtering, proceed to the next relevant section below.
- Defining Scoring Settings for Anti-Spam (page 544)
- Defining Spam Filtering Rules (page 545)
- Modifying Advanced Anti-Spam Settings

▶ If you are finished defining Anti-Spam settings and have not yet defined which traffic is filtered for spam, modify the IPv4 Access rules in the Firewall Policy as instructed in Editing Access Rules (page 672).

▶ Otherwise, the configuration is finished. Click **OK** and refresh the Firewall's policy to transfer the changes.

Modifying Advanced Anti-Spam Settings

In the advanced anti-spam settings, you can fine-tune graylisting and blacklisting options. You can also enable automatic graylisting and blacklisting, and enforcement of the SMTP protocol.

▼ To define advanced anti-spam settings

1. Switch to the **Advanced** tab in the Anti-Spam Settings Properties.

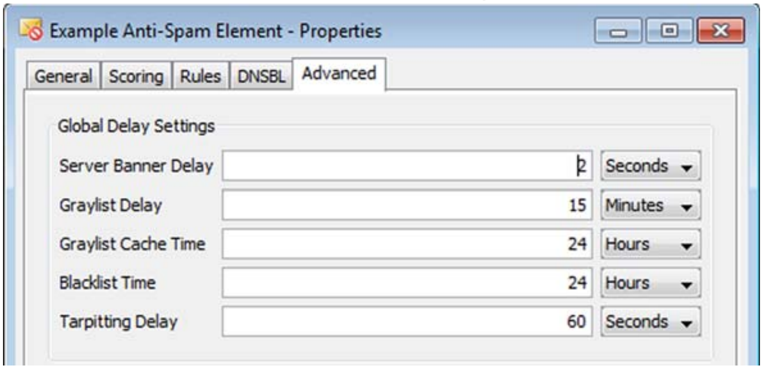

Global Delay Settings

Server Banner Delay 2 Seconds

Graylist Delay 15 Minutes

Graylist Cache Time 24 Hours

Blacklist Time 24 Hours

Tarpitting Delay 60 Seconds

2. (*Optional*) Modify the **Global Delay Settings**.

Option	Description
Server Banner Delay	The period of time for which the Firewall delays forwarding an SMTP banner to the client. The option is used to verify whether the SMTP client follows the protocol standard. Every client that starts a transmission before the banner is received fails this verification. Such a client is given a spam score. By default, set to 2 seconds.
Graylist Delay	The period of time for which an e-mail is graylisted. After this period the message is accepted if it passes further checks. By default, set to 15 minutes.
Graylist Cache Time	The period of time for which the IP address of the connecting host, the e-mail address of the message sender, and the e-mail address of the message recipient retrieved during the delivery attempt are cached. By default, set to 24 hours.
Blacklist Time	The period of time for which the IP address of the message sender is blacklisted. By default, set to 24 hours.
Tarpitting Delay	The period of time for which the reject response coming from the Firewall is delayed. The option is used to drain resources of detected spammers. The protected SMTP server is not affected. By default, set to 60 minutes.

3. Enable/disable **Enforce SMTP** to enforce SMTP protocol checks for traffic. If SMTP protocol violations are detected, the message is given a spam score.
4. Enable/disable **Automatic Graylisting.**
   - If the option is enabled, the Firewall automatically graylists all messages that pass anti-spam checks and that are not whitelisted, graylisted, blacklisted, or rejected.
   - The period of time for which the e-mails are graylisted is defined in the **Graylist Delay** field in the Global Delay Settings.
5. Enable/disable **Automatic Blacklisting.**
   - If the option is enabled, the IP addresses of senders of all discarded messages are automatically blacklisted.
   - The period of time for which the IP addresses are blacklisted is defined in the **Blacklist Time** field in the Global Delay Settings.

What's Next?

▶ If you want to define additional criteria for Anti-Spam filtering, proceed to the next relevant section below.

- Defining Scoring Settings for Anti-Spam (page 544)
- Defining Spam Filtering Rules (page 545)
- Defining DNSBL Settings (page 547)

▶ If you are finished defining Anti-Spam settings and have not yet defined which traffic is filtered for spam, modify the IPv4 Access rules in the Firewall Policy as instructed in Editing Access Rules (page 672).

▶ Otherwise, the configuration is finished. Click **OK** and refresh the Firewall's policy to transfer the changes.

Modifying Anti-Spam Settings Elements

Note – Modifying an Anti-Spam Settings element affects all Firewalls where the Anti-Spam Settings element is used.

▼ To modify an Anti-Spam Settings element

1. Select **Configuration→Configuration→Security Engine** from the menu. The Security Engine Configuration view opens.

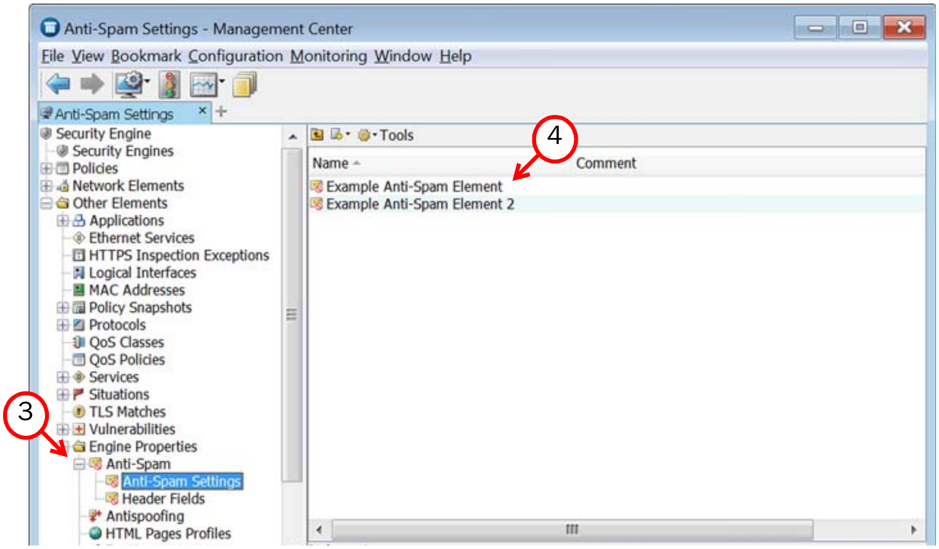

2. Browse to **Other Elements→Engine Properties→Anti-Spam.**

3. Expand the **Anti-Spam Settings** branch. The available elements are displayed in the right panel.

4. Right-click the Anti-Spam Settings element that you want to modify and select **Properties**.

5. Modify the element. See the sections listed below for more details:
   - Defining General Anti-Spam Settings (page 542)
   - Defining Scoring Settings for Anti-Spam (page 544)
   - Defining Spam Filtering Rules (page 545)
   - Defining DNSBL Settings (page 547)
   - Modifying Advanced Anti-Spam Settings (page 548)

6. Click **OK**.

What's Next?

▶ Refresh policies of all affected Firewalls to transfer the changes.

Chapter 34

Advanced Engine Settings

This section explains the settings on the Advanced tab in the properties of Firewall, IPS, and Layer 2 Firewall elements.

The following sections are included:

Getting Started with Advanced Engine Settings

Prerequisites: Creating New Engine Elements

Advanced settings cover various system parameters related to different features. These parameters generally have default values that are appropriate for most installations without any need for adjustments, or values that can be overridden in policies rule-by-rule when exceptions are needed.

❗ **Caution – Improper adjustments to some of the advanced settings may seriously degrade the performance of the system. Do not adjust the advanced settings unless you have a specific need to do so.**

▼ To open the advanced settings

1. Right-click an engine element and select **Properties**. The Properties dialog for that element opens.
2. Switch to the **Advanced** tab and adjust the settings as explained in the sections below.
   - Adjusting Firewall System Parameters
   - Adjusting Firewall Traffic Handling Parameters (page 554)
   - Adjusting IPS Engine System Parameters (page 560)
   - Adjusting IPS Engine Traffic Parameters (page 561)
   - Adjusting Layer 2 Firewall System Parameters (page 564)
   - Adjusting Layer 2 Firewall Traffic Parameters (page 565)
   - Adjusting Master Engine System Parameters (page 568)
   - Adjusting Master Engine Traffic Handling Parameters (page 569)
   - Adjusting Master Engine Clustering Options (page 571)
   - Adjusting Virtual Security Engine System Parameters (page 575)
   - Adjusting Virtual Security Engine Traffic Handling Parameters (page 575)
   - Configuring Inspection of Tunneled Traffic (page 577)
   - Configuring Default SYN Rate Limits (page 579)
   - Configuring Default Log Handling Settings (page 580)

Adjusting Firewall System Parameters

Prerequisites: See Getting Started with Advanced Engine Settings

The System Parameters section on a Firewall element's Advanced tab contains settings related to how the single or clustered firewall behaves under certain conditions. The table below explains the available settings.

Table 34.1 Firewall Advanced Settings - System Parameters Section

Setting	Description	Notes
Encrypt Configuration Data	By default, the configuration of the engine is stored in an encrypted format. Disable the encryption only if instructed to do so by Stonesoft technical support.	
Contact Node Timeout	Used when opening any kind of communication between the Management Server and the engines. A consistently slow network connection may require increasing this value. The default value is 60 seconds.	Setting the timeout value too short or too long may delay or prevent contact between the Management Server and the engines.
Automatic Reboot Timeout	The length of time after which an error situation is considered non-recoverable and the engine automatically reboots. The default value is 10 seconds.	Set to 0 to disable.
Policy Handshake	When the policy handshake feature is on, the nodes automatically roll back to using the previously installed policy if connectivity is lost after installing a new policy. Without this feature, you must switch to the previous configuration manually through the engine's boot menu.	We recommend adjusting the timeout (next setting below) rather than disabling this feature completely if there is a need to make changes.
Rollback Timeout	The time the engine waits for a management connection before it rolls back to the previously installed policy when the Policy Handshake option is active. The default value is 60 seconds.	
Automated Node Certificate Renewal	When selected, the engine's certificate for system communications is automatically renewed before it expires. Otherwise, the certificate must be renewed manually. Each certificate for system communications is valid for three years. If the certificate expires, other components refuse to communicate with the engine.	Does not renew VPN certificates. Automatic certificate renewal for internally signed VPN certificates is set separately in the Internal Security Gateway element properties. See Defining Security Gateways (page 923).
FIPS-Compatible Operating Mode	Activates a mode that is compliant with the Federal Information Processing Standard FIPS 140-2.	See the *Common Criteria User's Guide* for more information.
Log Handling	Settings related to adjusting logging when the log spool on the engines fills up or when the number of Antispoofing and Discard logs grows too high. See Configuring Default Log Handling Settings (page 580).	You can adjust the logging of Antispoofing and Discard logs also for specific interfaces. See Configuring Advanced Interface Properties for Firewalls (page 436).

Table 34.1 Firewall Advanced Settings - System Parameters Section (Continued)

Setting	Description	Notes
Clustering (*Firewall Clusters only*)	Settings related to the communications between cluster members and load-balancing between the nodes. See Adjusting Firewall Clustering Options (page 556).	

Adjusting Firewall Traffic Handling Parameters

Prerequisites: See Getting Started with Advanced Engine Settings

The Traffic Handling Parameters section on a Firewall element's Advanced tab contains settings that relate to how the single or clustered firewall treats certain traffic conditions. The table below explains the available settings.

Table 34.2 Firewall Advanced Settings - Traffic Handling Parameters Section

Setting	Description	Notes
Strict Connection Tracking	When selected, the engine does not permit TCP traffic to pass through before a complete, valid TCP handshake is performed. Leave this setting deselected unless you have a specific need to enable it.	This option is always active when the firewall is in FIPS-compliant operating mode. You can override this general setting and configure connection tracking for TCP, UDP, and ICMP traffic in Access rules. See Defining Firewall Allow Action Options (page 678).
Virtual Defragmenting	When the engine receives fragmented packets, it defragments the packets for inspection. The original fragments are queued on the engine until the inspection is finished. This setting defines how the fragmented packets are sent onwards after inspection if the packets are allowed. If this option is selected, the fragmented packets are sent onwards using the same fragmentation as when they arrived at the firewall. If the option is not selected, the packets are sent onwards as if they had arrived unfragmented.	

Table 34.2 Firewall Advanced Settings - Traffic Handling Parameters Section (Continued)

Setting	Description	Notes
Strict TCP Mode for Deep Inspection	Provides enhanced protection against TCP evasion attempts. The engine controls the progress of a TCP connection and checks that the TCP handshake proceeds according to the TCP protocol. The same node must be able to see all the packets in the connection. The engine also enforces the order of the packets and the packet direction (for example, SYN and SYN-ACK packets are not allowed from the same interface). The engine also modifies the TCP packets' header, especially the TCP window size. It removes TCP options that are not supported by the Strict TCP mode (for example, timestamps) from the packets. When Strict TCP mode is first activated, currently open TCP connections are inspected according to Normal TCP mode and Strict TCP mode is only applied to new TCP connections.	**Caution!** Strict TCP mode may significantly decrease throughput because it limits the maximum window size, and may delay or drop packets to enforce compliance with TCP standards.
Default Connection Termination in Inspection Policy	Defines how connections that match rules with the Terminate action in the Inspection Policy are handled. “Terminate and Log Connection (Default)” stops matching connections. This is the default. “Only Log Connection” does not stop matching connections, but creates a special log entry Terminate (passive), which allows you to filter the logs to see which type of connections are dropped when you change the option to actually stop traffic. This is useful for testing purposes.	You can override this general setting in the Inspection Policy.
VPN Settings	Settings related to adjusting IPsec VPNs performance-wise (see Advanced VPN Tuning (page 991)), and settings related to managing VPN client addresses (see Managing VPN Client IP Addresses (page 1000)).	
Policy Routing	Policy routing settings. See Defining Policy Routing (page 596).	
Idle Timeouts	Settings for general connection timeouts. See Setting Connection Timeouts (page 578).	
SYN Rate Limits	Settings for configuring limits for SYN packets sent to the engine. See Configuring Default SYN Rate Limits (page 579).	You can also configure SYN Rate Limits for specific interfaces. See Configuring Advanced Interface Properties for Firewalls (page 436).

Adjusting Firewall Clustering Options

Prerequisites: See Getting Started with Advanced Engine Settings

You can check and adjust some of the settings related to how the firewall cluster behaves and exchanges information from node to node as explained in Adjusting General Firewall Clustering Options.

If you are an advanced user, it is also possible to tune settings related to how the traffic load is balanced between several online nodes as explained in Tuning the Firewall Load Balancing Filter (page 558). However, we strongly discourage you from modifying the load balancing settings unless Stonesoft's technical support instructs you to make particular changes.

Adjusting General Firewall Clustering Options

▼ To adjust basic clustering settings

1. Right-click a Firewall Cluster element and select **Properties**. The Properties dialog for the firewall cluster opens.
2. Switch to the **Advanced** tab and click **Clustering**. The Clustering Properties dialog opens on the Cluster tab.
3. Configure the settings as described below:.

Table 34.3 General Firewall Clustering Options

Setting	Description
Cluster Mode	Defines whether the clustered engines are all online balancing the traffic or whether one node is online at a time and the other engines are used as backups. **Balancing**: All nodes are simultaneously online providing enhanced performance and high availability in case of node failure. Balancing is the default mode. **Standby**: Only one node can be online at a time. It is recommended to have at least one other node on standby to allow automatic takeover in case of failure. Several nodes can be on standby at a time. A randomly selected standby node is turned online when the online node fails.
Heartbeat Message Period	Defines how often clustered engines send heartbeat messages to each other (notifying that they are up and running). Enter the value in milliseconds. The default value is 1000 milliseconds (one second). **Caution!** Setting the Heartbeat Message Period too low may result in unnecessary heartbeat failures. Setting the Heartbeat Message Period too high may cause unnecessary service outages when a failure occurs.
Heartbeat Failover Time	Defines the time from the previous heartbeat message after which a node is considered as failed. Enter the value in milliseconds. The failover time must be at least twice as long as the Heartbeat Message Period. The default value is 5000 milliseconds. **Caution!** Setting the Heartbeat Failover Time too low may result in unnecessary heartbeat failures. Setting the Heartbeat Failover Time too high may cause unnecessary service outages when a failure occurs.

Table 34.3 General Firewall Clustering Options (Continued)

Setting	Description
State Sync	Defines how the nodes exchange information about the traffic that they process. **All**: (*recommended*) Both full and incremental synchronization messages are sent. This allows frequent updates without consuming resources excessively. Regular full synchronization ensures that all nodes stay synchronized even if some incremental messages are not delivered. **Full Only**: (*not recommended*) Only full synchronization messages are sent. Incremental updates are not sent in between, so nodes may not have the same information on connections unless the full sync interval is significantly reduced. **Note!** We strongly recommend using Access rule options to disable state synchronization for specific traffic rather than adjusting the State Sync settings for the cluster. See Defining Access Rule Action Options (page 676) for more information.
Full Sync Interval Incr Sync Interval	Define how frequently the full synchronizations and incremental synchronizations are done. Do not set the values much higher or lower than their defaults (5000 ms for full, 50 ms for incremental) **Caution!** Adjusting the Sync Intervals has significant impact on the cluster's performance. Inappropriate settings seriously degrade the firewall's performance.
Sync Security Level	**None**: no security features. Do not select this options unless the heartbeat traffic uses a dedicated, secure network that does not handle other traffic. **Sign**: (*default*) transmissions are authenticated to prevent outside injections of connection state information. **Encrypt and Sign**: transmissions are authenticated and encrypted. This option increases the overhead compared to the default option, but is strongly recommended if the node-to-node communications are relayed through insecure networks (for example, if the backup heartbeat is configured on an interface that handles other traffic).
Heartbeat IP	Enter an IP address between 224.0.0.0 and 239.255.255.255 if you want to change the multicast IP addresses used for node-to-node communications (default: 225.1.1.1). This multicast IP address must not be used for other purposes on any of the network interfaces.
Synchronization IP	Enter an IP address between 224.0.0.0 and 239.255.255.255 if you want to change the multicast IP addresses used for node-to-node communications (default: 225.1.1.2). This multicast IP address must not be used for other purposes on any of the network interfaces.

What's Next?

- If your cluster is in load-balancing mode and you want to change settings related to how traffic is balanced, see Tuning the Firewall Load Balancing Filter (page 558).
- Otherwise, refresh the Firewall's policy to transfer the changes.

Tuning the Firewall Load Balancing Filter

Caution – Do not manually tune the load balancing filter unless you are absolutely certain it is necessary. Normally, there is no need to tune the load balancing filter because the configuration generates all required entries automatically. Unnecessary tuning may adversely affect the operation of the filter.

You can tune the Firewall Cluster's load balancing filter by adjusting some parameters and filter entries. You can define connections where packets are passed or blocked by the filter. You can select certain connections to be passed on one node while they would be blocked on another node. You can also configure port numbers to be ignored in load balancing.

What's Next?

▶ Manually Tuning the Load Balancing Filter

Manually Tuning the Load Balancing Filter

▼ To tune the load balancing filter manually

1. Right-click a Firewall Cluster and select **Properties**. The Firewall Cluster Properties dialog opens.
2. Switch to the **Advanced** tab and click **Clustering**. The Clustering Properties dialog opens.
3. Switch to the **Manual LB Filters** tab.
4. Select the **Filter Mode**:
   - **Static**: packet ownership (the node to which the connection/packet belongs) may change only when nodes are added or removed from the cluster, or when they switch from an offline state to an online state or vice versa.
   - **Dynamic**: traffic is balanced to avoid node overloads and existing connections are moved between nodes whenever overload is detected.
5. (*Optional*) Select **Load Balancing Filter Uses Ports** to include a port value for selecting between all nodes.
   - This setting decreases the granularity of VPN load balancing, and increases the granularity of other traffic load balancing. In typical networks, traffic is balanced based on IP information only. If there is a dominating pair of communication IP addresses, apply the **Use Ports** option in the load balancing filter entry only to their traffic and not globally.

Caution – Enabling the "Load Balancing Filter Uses Ports" option is not compatible with some features, such as client-to-gateway VPNs.

What's Next?

▶ If required, continue with the definition of filter entries in Adding Load Balancing Filter Entries (page 559).

▶ Otherwise, refresh the Firewall's policy to transfer the changes.

Adding Load Balancing Filter Entries

▼ To define load balancing filter entries

1. Click **Add** to generate a new filter entry row.
2. Double-click the **IP Address** field. The Load-Balancing Filter IP Entry dialog opens.
3. Select whether you want to filter either a **Single** IP address or a **Range** of IP addresses.
4. Enter the IP address or address range and the netmask and click **OK**. The Load-Balancing Filter IP Entry dialog closes.
5. Click the **Action** cell and select one of the following actions:
   - **None**: no action is performed for the IP address specified in this entry (used with the Replacement IP, Use Ports, NAT Enforce, Use IPsec, or Ignore Other options).
   - **Replace by**: the original IP address is replaced by the IP address in the Replacement IP cell. This is the default action.
   - **Pass on All Nodes**: the filter entry allows packets to all nodes.
   - **Block on All Nodes**: the filter entry blocks packets to all nodes.
   - **Pass on Node [n]**: the filter entry forces node [n] (where n is the node ID number) to handle all the packets belonging to the connection specified in this entry.
6. If you selected **Replace by** as the action, click the **Replacement IP** field and enter the replacement IP address.
7. (*Optional*) Select additional options as described below:

Option	Description
Use Ports	Overrides the global **Load Balancing Filter Uses Ports** option. For example, if two hosts send the majority of traffic through the Firewall, you can set the Use Ports option for one of them to divide the traffic between the cluster nodes, improving granularity. Using this option for IP addresses in a VPN site may reduce the granularity of VPN load balancing and prevent VPN client connections involving those IP addresses.
NAT Enforce	Do not enable this option unless instructed to do so by Stonesoft technical support. The option enables a specific NAT-related process in the load balancing filter.
Use IPsec	Specifies addresses receiving IPsec traffic on the node itself. Do not enable this option unless instructed to do so by Stonesoft technical support. The option enables a specific load-balancing process for all IPsec traffic directed to the IP address specified in the filter entry.
Ignore Other	Forces the handling of packets to and from the specified IP address(es) by one node at a time.

8. Click **OK**.

What's Next?

▶ Refresh the Firewall's policy to transfer the changes.

Adjusting IPS Engine System Parameters

Prerequisites: See Getting Started with Advanced Engine Settings

The System Parameters section on an IPS element's Advanced tab contains settings that relate to how the Single IPS or IPS Cluster behaves under certain conditions. The table below explains the available settings.

Table 34.4 IPS Advanced Settings - System Parameters Section

Setting	Description	Notes
Encrypt Configuration Data	By default, the configuration of the engine is stored in an encrypted format. Disable the encryption only if instructed to do so by Stonesoft technical support.	
Bypass Traffic on Overload	This option maintains connectivity when an inline IPS engine is near the limits of its performance. When the option is selected, the IPS engine dynamically reduces the number of inspected connections if the load is too high. Some traffic may pass through without any access control or inspection if this option is selected. Bypassed traffic is not counted when a possible license throughput limit is enforced. The bypass does not affect traffic subject to TLS Inspection. If this option is not selected, the IPS engine inspects all connections. Some connections may not get through if the inline IPS engine gets overloaded.	
Contact Node Timeout	Used when opening any kind of communication between the Management Server and the engines. A consistently slow network connection may require increasing this value. The default value is 60 seconds.	Setting the timeout value too short or too long may delay or prevent contact between the Management Server and the engines.
Automatic Reboot Timeout	The length of time after which an error situation is considered non-recoverable and the engine automatically reboots. The default value is 10 seconds.	Set to 0 to disable.
Policy Handshake	When the policy handshake feature is on, the nodes automatically roll back to using the previously installed policy if connectivity is lost after installing a new policy. Without this feature, you must switch to the previous configuration manually through the engine's boot menu.	We recommend adjusting the timeout (next setting below) rather than disabling this feature completely if there is a need to make changes.
Rollback Timeout	The time the engine waits for a management connection before it rolls back to the previously installed policy when the Policy Handshake option is active. The default value is 60 seconds.	

Table 34.4 IPS Advanced Settings - System Parameters Section (Continued)

Setting	Description	Notes
Automated Node Certificate Renewal	With the option selected, the engine's certificate for system communications is automatically renewed before it expires. Otherwise, the certificate must be renewed manually. Each certificate for system communications is valid for three years. If the certificate expires, other components refuse to communicate with the engine.	
Log Handling	Settings related to adjusting logging when the log spool on the engines fills up or when the number of Discard logs grows too high. See Configuring Default Log Handling Settings (page 580).	You can adjust the logging of Discard logs also for specific interfaces. See Configuring Advanced Interface Properties for IPS Engines (page 469).
Clustering (*IPS Clusters only*)	Settings for adjusting how an IPS Cluster behaves. See Adjusting IPS Clustering Options (page 562).	

Adjusting IPS Engine Traffic Parameters

Prerequisites: See Getting Started with Advanced Engine Settings

The Traffic Handling Parameters section on an IPS element's Advanced tab contains settings that relate to how the Single IPS or IPS Cluster treats certain traffic conditions. The table below explains the available settings

Table 34.5 IPS Advanced Settings - Traffic Handling Parameters Section

Setting	Description	Notes
Strict TCP Mode for Deep Inspection	Provides enhanced protection against TCP evasion attempts. The engine controls the progress of a TCP connection and checks that the TCP handshake proceeds according to the TCP protocol. The same node must be able to see all the packets in the connection. The engine also enforces the order of the packets and the packet direction (for example, SYN and SYN-ACK packets are not allowed from the same interface). The engine also modifies the TCP packets' header, especially the TCP window size. It removes TCP options that are not supported in Strict TCP mode (for example, timestamps) from the packets. When Strict TCP mode is first activated, currently open TCP connections are inspected according to Normal TCP mode and Strict TCP mode is only applied to new TCP connections.	Strict TCP mode is only available if the IPS is installed in inline mode. **Caution!** Strict TCP mode may significantly decrease throughput because it limits the maximum window size, and may delay or drop packets to enforce compliance with TCP standards.

Table 34.5 IPS Advanced Settings - Traffic Handling Parameters Section (Continued)

Setting	Description	Notes
Default Connection Termination in Inspection Policy	Defines how connections that match rules with the Terminate action in the Inspection Policy are handled. “Terminate and Log Connection (Default)” stops matching connections if they are travelling through an Inline interface. This is the default. “Only Log Connection” does not stop matching connections, but creates a special log entry Terminate (passive), which allows you to filter the logs to see which type of connections are dropped when you change the option to actually stop traffic. This is useful for testing purposes.	You can override this general setting in the Inspection Policy.
Idle Timeouts	Settings for general connection timeouts. See Setting Connection Timeouts (page 578).	
SYN Rate Limits	Settings for configuring limits for SYN packets sent to the engine. See Configuring Default SYN Rate Limits (page 579).	You can also configure SYN Rate Limits for specific interfaces. See Configuring Advanced Interface Properties for IPS Engines (page 469).
Tunneled Traffic	Settings for inspecting tunneled traffic. See Configuring Inspection of Tunneled Traffic (page 577).	

Adjusting IPS Clustering Options

IPS Clusters operate by default in *load balancing* mode. This means that all the configured nodes in an IPS Cluster are online simultaneously and the traffic is distributed among the operational nodes. The load balancing aims to keep the traffic load as evenly distributed as possible.

Alternatively, the IPS Cluster can run in *standby mode*. In that case, only one IPS node at a time is online and processing traffic, while the others are in standby mode. Only if the online node fails, one of the standby nodes goes online to take over the connections being handled by the failed node.

▼ To adjust IPS clustering settings

1. Right-click an IPS Cluster element and select **Properties**. The Properties dialog for the IPS Cluster opens.
2. Switch to the **Advanced** tab.
3. Click **Clustering**. The Clustering Properties dialog opens.

4. Configure the settings as described below:

Setting	Description
Cluster Mode	Defines whether the clustered engines are all online balancing the traffic or whether one node is online at a time and the other engines are used as backups. **Balancing**: All nodes are simultaneously online providing enhanced performance and high availability in case of node failure. Balancing is the default mode. **Standby**: Only one node can be online at a time. It is recommended to have at least one other node on standby to allow automatic takeover in case of failure. Several nodes can be on standby at a time. A randomly selected standby node is turned online when the online node fails.
Filter Mode	Defines how traffic is balanced between the nodes. **Static**: packet ownership (the node to which the connection/packet belongs) may change only when nodes are added or removed from the cluster, or when they switch from an offline state to an online state or vice versa. **Dynamic**: traffic is balanced to avoid node overloads and existing connections are moved between nodes whenever overload is detected.
Heartbeat Message Period	Defines how often clustered engines send heartbeat messages to each other (notifying that they are up and running). Enter the value in milliseconds. The default value is 1000 milliseconds (one second). **Caution!** Setting the Heartbeat Message Period too low may result in unnecessary heartbeat failures. Setting the Heartbeat Message Period too high may cause unnecessary service outages when a failure occurs.
Heartbeat Failover Time	Defines the time from the previous heartbeat message after which a node is considered as failed. Enter the value in milliseconds. The failover time must be at least twice as long as the Heartbeat Message Period. The default value is 5000 milliseconds. **Caution!** Setting the Heartbeat Failover Time too low may result in unnecessary heartbeat failures. Setting the Heartbeat Failover Time too high may cause unnecessary service outages when a failure occurs.
Heartbeat IP	Enter an IP address between 224.0.0.0 and 239.255.255.255 if you want to change the multicast IP addresses used for node-to-node communications (default: 225.1.1.1). This multicast IP address must not be used for other purposes on any of the network interfaces.

5. Click **OK**.

What's Next?

▶ Refresh the IPS Cluster's policy to transfer the configuration changes.

Adjusting Layer 2 Firewall System Parameters

Prerequisites: See Getting Started with Advanced Engine Settings

The System Parameters section on a Layer 2 Firewall element's Advanced tab contains settings that relate to how the Single Layer 2 Firewall or a Layer 2 Firewall Cluster behaves under certain conditions. The table below explains the available settings.

Table 34.6 Layer 2 Firewall Advanced Settings - System Parameters Section

Setting	Description	Notes
Encrypt Configuration Data	By default, the configuration of the engine is stored in an encrypted format. Disable the encryption only if instructed to do so by Stonesoft technical support.	
Contact Node Timeout	Used when opening any kind of communication between the Management Server and the engines. A consistently slow network connection may require increasing this value. The default value is 60 seconds.	Setting the timeout value too short or too long may delay or prevent contact between the Management Server and the engines.
Automatic Reboot Timeout	The length of time after which an error situation is considered non-recoverable and the engine automatically reboots. The default value is 10 seconds.	Set to 0 to disable.
Policy Handshake	When the policy handshake feature is on, the nodes automatically roll back to using the previously installed policy if connectivity is lost after installing a new policy. Without this feature, you must switch to the previous configuration manually through the boot menu on the engine.	We recommend adjusting the timeout (next setting below) rather than disabling this feature completely if there is a need to make changes.
Rollback Timeout	The time the engine waits for a management connection before it rolls back to the previously installed policy when the Policy Handshake option is active. The default value is 60 seconds.	
Automated Node Certificate Renewal	With the option selected, the engine's certificate for system communications is automatically renewed before it expires. Otherwise, the certificate must be renewed manually. Each certificate for system communications is valid for three years. If the certificate expires, other components refuse to communicate with the engine.	
Log Handling	Settings related to adjusting logging when the log spool on the engines fills up or when the number of Discard logs grows too high. See Configuring Default Log Handling Settings (page 580).	You can adjust the logging of Discard logs also for specific interfaces. See Configuring Advanced Interface Properties for Layer 2 Firewalls (page 482).
Clustering (*Layer 2 Firewall Clusters only*)	Settings for adjusting how a Layer 2 Firewall Cluster behaves. See Adjusting Layer 2 Firewall Clustering Options (page 566).	

Adjusting Layer 2 Firewall Traffic Parameters

Prerequisites: See Getting Started with Advanced Engine Settings

The Traffic Handling Parameters section on a Layer 2 Firewall element's Advanced tab contains settings that relate to how the Single Layer 2 Firewall or a Layer 2 Firewall Cluster treats certain traffic conditions. The table below explains the available settings

Table 34.7 Layer 2 Firewall Advanced Settings - Traffic Handling Parameters Section

Setting	Description	Notes
Strict TCP Mode for Deep Inspection	Provides enhanced protection against TCP evasion attempts. The engine controls the progress of a TCP connection and checks that the TCP handshake proceeds according to the TCP protocol. The same node must be able to see all the packets in the connection. The engine also enforces the order of the packets and the packet direction (for example, SYN and SYN-ACK packets are not allowed from the same interface). The engine also modifies the TCP packets' header, especially the TCP window size. It removes TCP options that are not supported by the Strict TCP mode (for example, timestamps) from the packets. When Strict TCP mode is first activated, currently open TCP connections are inspected according to Normal TCP mode and Strict TCP mode is only applied to new TCP connections.	**Caution!** Strict TCP mode may significantly decrease throughput because it limits the maximum window size, and may delay or drop packets to enforce compliance with TCP standards.
Default Connection Termination in Inspection Policy	Defines how connections that match rules with the Terminate action in the Inspection Policy are handled. "Terminate and Log Connection (Default)" stops matching connections if they are travelling through an Inline interface. This is the default. "Only Log Connection" does not stop matching connections, but creates a special log entry Terminate (passive), which allows you to filter the logs to see which type of connections are dropped when you change the option to actually stop traffic. This is useful for testing purposes.	You can override this general setting in the Inspection Policy.
Idle Timeouts	Settings for general connection timeouts. See Setting Connection Timeouts (page 578).	
SYN Rate Limits	Settings for configuring limits for SYN packets sent to the engine. See Configuring Default SYN Rate Limits (page 579).	You can also configure SYN Rate Limits for specific interfaces. See Configuring Advanced Interface Properties for Layer 2 Firewalls (page 482).
Tunneled Traffic	Settings for inspecting tunneled traffic. See Configuring Inspection of Tunneled Traffic (page 577).	

Adjusting Layer 2 Firewall Clustering Options

Prerequisites: Creating New Engine Elements

Layer 2 Firewall Clusters operate in *active-standby* mode. Only one Layer 2 Firewall node at a time is online and processing traffic, while the others are standby. Only if the online node fails, one of the standby nodes goes online to take over the connections being handled by the failed node.

▼ To set the cluster mode of the Layer 2 Firewall Cluster

1. Right-click a Layer 2 Firewall Cluster element and select **Properties**. The Properties dialog opens.
2. Switch to the **Advanced** tab.
3. Click the **Clustering**. The Clustering Properties dialog opens.
4. Configure the settings as described below:

Table 34.8 Layer 2 Firewall Clustering Options

Setting	Description
Heartbeat Message Period	Defines how often clustered engines send heartbeat messages to each other (notifying that they are up and running). Enter the value in milliseconds. The default value is 1000 milliseconds (one second). **Caution!** Setting the Heartbeat Message Period too low may result in unnecessary heartbeat failures. Setting the Heartbeat Message Period too high may cause unnecessary service outages when a failure occurs.
Heartbeat Failover Time	Defines the time from the previous heartbeat message after which a node is considered as failed. Enter the value in milliseconds. The failover time must be at least twice as long as the Heartbeat Message Period. The default value is 5000 milliseconds. **Caution!** Setting the Heartbeat Failover Time too low may result in unnecessary heartbeat failures. Setting the Heartbeat Failover Time too high may cause unnecessary service outages when a failure occurs.
State Sync	Defines how the nodes exchange information about the traffic that they process. **All**: (*recommended*) Both full and incremental synchronization messages are sent . This allows frequent updates without consuming resources excessively. Regular full synchronization ensures that all nodes stay synchronized even if some incremental messages are not delivered. **Full Only**: (*not recommended*) Only full synchronization messages are sent. Incremental updates are not sent in between, so nodes may not have the same information on connections unless the full sync interval is significantly reduced. **Note!** We strongly recommend using Access rule options to disable state synchronization for specific traffic rather than adjusting the State Sync settings for the cluster. See Defining Access Rule Action Options (page 676) for more information.

Table 34.8 Layer 2 Firewall Clustering Options (Continued)

Setting	Description
Full Sync Interval Incr Sync Interval	Define how frequently the full synchronizations and incremental synchronizations are done. Do not set the values much higher or lower than their defaults (5000 ms for full, 50 ms for incremental) **Caution!** Adjusting the Sync Intervals has significant impact on the cluster's performance. Inappropriate settings seriously degrade the firewall's performance.
Sync Security Level	**None**: no security features. Do not select this options unless the heartbeat traffic uses a dedicated, secure network that does not handle other traffic. **Sign**: (*default*) transmissions are authenticated to prevent outside injections of connection state information. **Encrypt and Sign**: transmissions are authenticated and encrypted. This option increases the overhead compared to the default option, but is strongly recommended if the node-to-node communications are relayed through insecure networks (for example, if the backup heartbeat is configured on an interface that handles other traffic).
Heartbeat IP	Enter an IP address between 224.0.0.0 and 239.255.255.255 if you want to change the multicast IP addresses used for node-to-node communications (default: 225.1.1.1). This multicast IP address must not be used for other purposes on any of the network interfaces.
Synchronization IP	Enter an IP address between 224.0.0.0 and 239.255.255.255 if you want to change the multicast IP addresses used for node-to-node communications (default: 225.1.1.2). This multicast IP address must not be used for other purposes on any of the network interfaces.

5. Click **OK**.

What's Next?

▶ Refresh the Layer 2 Firewall Cluster's policy to transfer the configuration changes.

Adjusting Master Engine System Parameters

Prerequisites: See Getting Started with Advanced Engine Settings

The System Parameters section on a Master Engine element's Advanced tab contains settings that relate to how the Master Engine behaves under certain conditions. The table below explains the available settings.

Table 34.9 Master Engine Advanced Settings - System Parameters Section

Setting	Description	Notes
Encrypt Configuration Data	By default, the configuration of the engine is stored in an encrypted format.	**Caution!** Disable the encryption only if instructed to do so by Stonesoft technical support.
Contact Node Timeout	Used when opening any kind of communication between the Management Server and the engines. A consistently slow network connection may require increasing this value. The default value is 60 seconds.	Setting the timeout value too short or too long may delay or prevent contact between the Management Server and the engines.
Automatic Reboot Timeout	The length of time after which an error situation is considered non-recoverable and the engine automatically reboots. The default value is 10 seconds.	Set to 0 to disable.
Policy Handshake	When the policy handshake feature is on, the nodes automatically roll back to using the previously installed policy if connectivity is lost after installing a new policy. Without this feature, you must switch to the previous configuration manually through the boot menu on the engine.	We recommend adjusting the Rollback Timeout (next setting below) rather than disabling this feature completely if there is a need to make changes.
Rollback Timeout	The time the engine waits for a management connection before it rolls back to the previously installed policy when the Policy Handshake option is active. The default value is 60 seconds.	
Automated Node Certificate Renewal	With the option selected, the engine's certificate for system communications is automatically renewed before it expires. Otherwise, the certificate must be renewed manually. Each certificate for system communications is valid for three years. If the certificate expires, other components refuse to communicate with the engine.	
FIPS-Compatible Operating Mode	Activates a mode that is compliant with the Federal Information Processing Standard FIPS 140-2.	See the *Common Criteria User's Guide* for more information.
Log Handling	Settings related to adjusting logging when the log spool on the engines fills up or when the number of Antispoofing and Discard logs grows too high. See Configuring Default Log Handling Settings (page 580).	You can adjust the logging of Antispoofing and Discard logs also for specific interfaces. See Configuring Advanced Interface Properties for Master Engines (page 495).

Table 34.9 Master Engine Advanced Settings - System Parameters Section (Continued)

Setting	Description	Notes
Clustering	Settings related to the communications between cluster members and load-balancing between the nodes. See Adjusting Master Engine Clustering Options (page 571).	

Adjusting Master Engine Traffic Handling Parameters

Prerequisites: See Getting Started with Advanced Engine Settings

The Traffic Handling Parameters section on a Master Engine element's Advanced tab contains settings that relate to how the Master Engine treats certain traffic conditions. The table below explains the available settings.

Table 34.10 Master Engine Advanced Settings - Traffic Handling Parameters Section

Setting	Description	Notes
Strict Connection Tracking	When selected, the engine does not permit TCP traffic to pass through before a complete, valid TCP handshake is performed. Leave this setting deselected unless you have a specific need to enable it.	This option is always active when the Master Engine is in FIPS-compliant operating mode. You can override this general setting and configure connection tracking for TCP, UDP, and ICMP traffic in Access rules. See Defining Firewall Allow Action Options (page 678).
Virtual Defragmenting	When the engine receives fragmented packets, it defragments the packets for inspection. The original fragments are queued on the engine until the inspection is finished. This setting defines how the fragmented packets are sent onwards after inspection if the packets are allowed. If this option is selected, the fragmented packets are sent onwards using the same fragmentation as when they arrived at the engine. If the option is not selected, the packets are sent onwards as if they had arrived unfragmented.	

Table 34.10 Master Engine Advanced Settings - Traffic Handling Parameters Section (Continued)

Setting	Description	Notes
Strict TCP Mode for Deep Inspection	Provides enhanced protection against TCP evasion attempts. The engine controls the progress of a TCP connection and checks that the TCP handshake proceeds according to the TCP protocol. The same node must be able to see all the packets in the connection. The engine also enforces the order of the packets and the packet direction (for example, SYN and SYN-ACK packets are not allowed from the same interface). The engine also modifies the TCP packets' header, especially the TCP window size. It removes TCP options that are not supported by the Strict TCP mode (for example, timestamps) from the packets. When Strict TCP mode is first activated, currently open TCP connections are inspected according to Normal TCP mode and Strict TCP mode is only applied to new TCP connections.	**Caution!** Strict TCP mode may significantly decrease throughput because it limits the maximum window size, and may delay or drop packets to enforce compliance with TCP standards.
Default Connection Termination in Inspection Policy	Defines how connections that match rules with the Terminate action in the Inspection Policy are handled. "Terminate and Log Connection (Default)" stops matching connections. This is the default. "Only Log Connection" does not stop matching connections, but creates a special log entry Terminate (passive), which allows you to filter the logs to see which type of connections are dropped when you change the option to actually stop traffic. This is useful for testing purposes.	You can override this general setting in the Inspection Policy.
Policy Routing	Policy routing settings. See Defining Policy Routing (page 596).	
Idle Timeouts	Settings for general connection timeouts. See Setting Connection Timeouts (page 578).	
SYN Rate Limits	Settings for configuring limits for SYN packets sent to the engine. See Configuring Default SYN Rate Limits (page 579).	You can also configure SYN Rate Limits for specific interfaces. See Configuring Advanced Interface Properties for Firewalls (page 436).

Adjusting Master Engine Clustering Options

Prerequisites: See Getting Started with Advanced Engine Settings

You can check and adjust some of the settings related to how the Master Engine behaves and exchanges information from node to node as explained in Adjusting General Master Engine Clustering Options.

If you are an advanced user, it is also possible to tune settings related to how the traffic load is balanced between several online nodes as explained in Tuning the Master Engine Load Balancing Filter (page 573). However, we strongly discourage you from modifying the load balancing settings unless Stonesoft Support instructs you to make particular changes.

Adjusting General Master Engine Clustering Options

▼ To adjust basic clustering settings

1. Right-click a Master Engine element and select **Properties**. The Properties dialog opens.
2. Switch to the **Advanced** tab and click **Clustering**. The Clustering Properties dialog opens on the Cluster tab.
3. Configure the settings as described below:.

Table 34.11 General Master Engine Clustering Options

Setting	Description
Cluster Mode	Defines whether the clustered engines are all online balancing the traffic or whether one node is online at a time and the other engines are used as backups. **Balancing**: All nodes are simultaneously online providing enhanced performance and high availability in case of node failure. Balancing is the default mode. **Standby**: Only one node can be online at a time. It is recommended to have at least one other node on standby to allow automatic takeover in case of failure. Several nodes can be on standby at a time. A randomly selected standby node is turned online when the online node fails.
Heartbeat Message Period	Defines how often clustered engines send heartbeat messages to each other (notifying that they are up and running). Enter the value in milliseconds. The default value is 1000 milliseconds (one second). **Caution!** Setting the Heartbeat Message Period too low may result in unnecessary heartbeat failures. Setting the Heartbeat Message Period too high may cause unnecessary service outages when a failure occurs.
Heartbeat Failover Time	Defines the time from the previous heartbeat message after which a node is considered as failed. Enter the value in milliseconds. The failover time must be at least twice as long as the Heartbeat Message Period. The default value is 5000 milliseconds. **Caution!** Setting the Heartbeat Failover Time too low may result in unnecessary heartbeat failures. Setting the Heartbeat Failover Time too high may cause unnecessary service outages when a failure occurs.

Table 34.11 General Master Engine Clustering Options (Continued)

Setting	Description
State Sync	Defines how the nodes exchange information on the traffic that they process. **All**: (*recommended*) Both full and incremental synchronization messages are sent. This allows frequent updates without consuming resources excessively. Regular full synchronization ensures that all nodes stay synchronized even if some incremental messages are not delivered. **Full Only**: (*not recommended*) Only full synchronization messages are sent. Incremental updates are not sent in between, so nodes may not have the same information on connections unless the full sync interval is significantly reduced. **Note!** We strongly recommend using Access rule options to disable state synchronization for specific traffic rather than adjusting the State Sync settings for the cluster. See Defining Access Rule Action Options (page 676) for more information.
Full Sync Interval Incr Sync Interval	Define how frequently the full synchronizations and incremental synchronizations are done. Do not set the values much higher or lower than their defaults (5000 ms for full, 50 ms for incremental) **Caution!** Adjusting the Sync Intervals has significant impact on the cluster's performance. Inappropriate settings seriously degrade the engine's performance.
Sync Security Level	**None**: no security features. Do not select this options unless the heartbeat traffic uses a dedicated, secure network that does not handle other traffic. **Sign**: (*default*) transmissions are authenticated to prevent outside injections of connection state information. **Encrypt and Sign**: transmissions are authenticated and encrypted. This option increases the overhead compared to the default option, but is strongly recommended if the node-to-node communications are relayed through insecure networks (for example, if the backup heartbeat is configured on an interface that handles other traffic).
Heartbeat IP	Enter an IP address between 224.0.0.0 and 239.255.255.255 if you want to change the multicast IP addresses used for node-to-node communications (default: 225.1.1.1). This multicast IP address must not be used for other purposes on any of the network interfaces.
Synchronization IP	Enter an IP address between 224.0.0.0 and 239.255.255.255 if you want to change the multicast IP addresses used for node-to-node communications (default: 225.1.1.2). This multicast IP address must not be used for other purposes on any of the network interfaces.

What's Next?

- If your cluster is in load-balancing mode and you want to change settings related to how traffic is balanced, see Tuning the Master Engine Load Balancing Filter.
- Otherwise, refresh the Master Engine's policy to transfer the changes.

Tuning the Master Engine Load Balancing Filter

❗ **Caution – Do not manually tune the load balancing filter unless you are absolutely certain it is necessary. Normally, there is no need to tune the load balancing filter because the configuration generates all required entries automatically. Unnecessary tuning may adversely affect the operation of the filter.**

You can tune the Master Engine's load balancing filter by adjusting some parameters and filter entries. You can define connections where packets are passed or blocked by the filter. You can select certain connections to be passed on one node while they would be blocked on another node. You can also configure port numbers to be ignored in load balancing.

What's Next?

▶ Manually Tuning the Load Balancing Filter (page 558)

Manually Tuning the Load Balancing Filter

▼ **To tune the load balancing filter manually**

1. Right-click a Master Engine and select **Properties**. The Properties dialog opens.
2. Switch to the **Advanced** tab and click **Clustering**. The Clustering Properties dialog opens.
3. Switch to the **Manual LB Filters** tab.
4. Select the **Filter Mode**:
   - **Static**: packet ownership (the node to which the connection/packet belongs) may change only when nodes are added or removed from the cluster, or when they switch from an offline state to an online state or vice versa.
   - **Dynamic**: traffic is balanced to avoid node overloads and existing connections are moved between nodes whenever overload is detected.
5. (*Optional*) Select **Load Balancing Filter Uses Ports** to include a port value for selecting between all nodes.
   - This setting decreases the granularity of VPN load balancing, and increases the granularity of other traffic load balancing. In typical networks, traffic is balanced based on IP information only. If there is a dominating pair of communication IP addresses, apply the **Use Ports** option in the load balancing filter entry only to their traffic and not globally.

❗ **Caution – Enabling the "Load Balancing Filter Uses Ports" option is not compatible with some features, such as client-to-gateway VPNs.**

What's Next?

▶ If required, continue with the definition of filter entries in Adding Load Balancing Filter Entries.

▶ Otherwise, refresh the Master Engine's policy to transfer the changes.

Adding Load Balancing Filter Entries

▼ To define load balancing filter entries

1. Click **Add** to generate a new filter entry row.
2. Double-click the **IP Address** field. The Load-Balancing Filter IP Entry dialog opens.
3. Select whether you want to filter either a **Single** IP address or a **Range** of IP addresses.
4. Enter the IP address or address range and the netmask and click **OK**. The Load-Balancing Filter IP Entry dialog closes.
5. Click the **Action** cell and select one of the following actions:
   - **None**: no action is performed for the IP address specified in this entry (used with the Replacement IP, Use Ports, NAT Enforce, Use IPsec, or Ignore Other options).
   - **Replace by**: the original IP address is replaced by the IP address in the Replacement IP cell. This is the default action.
   - **Pass on All Nodes**: the filter entry allows packets to all nodes.
   - **Block on All Nodes**: the filter entry blocks packets to all nodes.
   - **Pass on Node [n]**: the filter entry forces node [n] (where n is the node ID number) to handle all the packets belonging to the connection specified in this entry.
6. If you selected **Replace by** as the action, click the **Replacement IP** field and enter the replacement IP address.
7. (*Optional*) Select additional options as described below:

Option	Description
Use Ports	Overrides the global **Load Balancing Filter Uses Ports** option. For example, if two hosts send the majority of traffic through the engine, you can set the Use Ports option for one of them to divide the traffic between the cluster nodes, improving granularity. Using this option for IP addresses in a VPN site may reduce the granularity of VPN load balancing and prevent VPN client connections involving those IP addresses.
NAT Enforce	Do not enable this option unless instructed to do so by Stonesoft technical support. The option enables a specific NAT-related process in the load balancing filter.
Use IPsec	Specifies addresses receiving IPsec traffic on the node itself. Do not enable this option unless instructed to do so by Stonesoft technical support. The option enables a specific load-balancing process for all IPsec traffic directed to the IP address specified in the filter entry.
Ignore Other	Forces the handling of packets to and from the specified IP address(es) by one node at a time.

8. Click **OK**.

What's Next?

▶ Refresh the Master Engine's policy to transfer the changes.

Adjusting Virtual Security Engine System Parameters

Prerequisites: See Getting Started with Advanced Engine Settings

The System Parameters section on a Virtual Security Engine element's Advanced tab contains settings that relate to how the Virtual Security Engine behaves under certain conditions. The table below explains the available settings.

Table 34.12 Virtual Security Engine Advanced Settings - System Parameters Section

Setting	Description	Notes
Encrypt Configuration Data	By default, the configuration of the engine is stored in an encrypted format.	**Caution!** Disable the encryption only if instructed to do so by Stonesoft technical support.
Log Handling	Settings related to adjusting logging when the log spool on the engines fills up or when the number of Antispoofing and Discard logs grows too high. See Configuring Default Log Handling Settings (page 580).	You can adjust the logging of Antispoofing and Discard logs also for specific interfaces. See Configuring Advanced Interface Properties for Virtual Security Engines (page 506).

Adjusting Virtual Security Engine Traffic Handling Parameters

Prerequisites: See Getting Started with Advanced Engine Settings

The Traffic Handling Parameters section on a Virtual Security Engine element's Advanced tab contains settings that relate to how the Virtual Security Engine treats certain traffic conditions. The table below explains the available settings.

Table 34.13 Master Engine Advanced Settings - Traffic Handling Parameters Section

Setting	Description	Notes
Strict Connection Tracking	When selected, the engine does not permit TCP traffic to pass through before a complete, valid TCP handshake is performed. Leave this setting deselected unless you have a specific need to enable it.	You can override this general setting and configure connection tracking for TCP, UDP, and ICMP traffic in Access rules. See Defining Firewall Allow Action Options (page 678).
Virtual Defragmenting	When the engine receives fragmented packets, it defragments the packets for inspection. The original fragments are queued on the engine until the inspection is finished. This setting defines how the fragmented packets are sent onwards after inspection if the packets are allowed. If this option is selected, the fragmented packets are sent onwards using the same fragmentation as when they arrived at the engine. If the option is not selected, the packets are sent onwards as if they had arrived unfragmented.	

Table 34.13 Master Engine Advanced Settings - Traffic Handling Parameters Section (Continued)

Setting	Description	Notes
Strict TCP Mode for Deep Inspection	Provides enhanced protection against TCP evasion attempts. The engine controls the progress of a TCP connection and checks that the TCP handshake proceeds according to the TCP protocol. The same node must be able to see all the packets in the connection. The engine also enforces the order of the packets and the packet direction (for example, SYN and SYN-ACK packets are not allowed from the same interface). The engine also modifies the TCP packets' header, especially the TCP window size. It removes TCP options that are not supported by the Strict TCP mode (for example, timestamps) from the packets. When Strict TCP mode is first activated, currently open TCP connections are inspected according to Normal TCP mode and Strict TCP mode is only applied to new TCP connections.	**Caution!** Strict TCP mode may significantly decrease throughput because it limits the maximum window size, and may delay or drop packets to enforce compliance with TCP standards.
Default Connection Termination in Inspection Policy	Defines how connections that match rules with the Terminate action in the Inspection Policy are handled. “Terminate and Log Connection (Default)” stops matching connections. This is the default. “Only Log Connection” does not stop matching connections, but creates a special log entry Terminate (passive), which allows you to filter the logs to see which type of connections are dropped when you change the option to actually stop traffic. This is useful for testing purposes.	You can override this general setting in the Inspection Policy.
VPN Settings	Settings related to adjusting IPsec VPNs performance-wise (see Advanced VPN Tuning (page 991)), and settings related to managing VPN client addresses (see Managing VPN Client IP Addresses (page 1000)).	
Policy Routing	Policy routing settings. See Defining Policy Routing (page 596).	
Idle Timeouts	Settings for general connection timeouts. See Setting Connection Timeouts (page 578).	
SYN Rate Limits	Settings for configuring limits for SYN packets sent to the engine. See Configuring Default SYN Rate Limits (page 579).	You can also configure SYN Rate Limits for specific interfaces. See Configuring Advanced Interface Properties for Virtual Security Engines (page 506).

Configuring Inspection of Tunneled Traffic

Prerequisites: See Getting Started with Advanced Engine Settings

If an IPS engine or a Layer 2 Firewall inspects traffic that is tunneled using IP-in-IP or Generic Routing Encapsulation (GRE), the contents of the tunneling packet can be checked against the IPv6 Access rules and/or IPv4 Access rules several times according to the number and type of layers in the tunnel. You can define in the Advanced engine settings how the IPS engine or Layer 2 Firewall inspects tunneled traffic.

▼ To configure inspection of tunneled traffic on an IPS engine or Layer 2 Firewall

1. Right-click an IPS or Layer 2 Firewall element and select **Properties**. The Properties dialog opens.
2. Switch to the **Advanced** tab.
3. Click **Tunneled Traffic**. The Inspection of Tunneled Traffic dialog opens.
4. Define the settings as described in the table below:

Settings	Description
Limit for Rematching Tunneled Traffic	Define how many times the contents of tunneled packets can be rematched against the IPv6 Access rules and/or IPv4 Access rules when several layers of tunneling are encountered. The default is 1. When the limit is reached, the action defined in the next setting below is taken.
Action if Limit is Exceeded	Define whether remaining encapsulated packets inside the tunneling packet are allowed without further inspection or discarded. The default is to discard the remaining packets. When this action is triggered, you are notified according to the Log Level setting below.
Log Level	Define whether you are notified through a normal (stored) log entry or an Alert when the limit for rematching tunneled traffic is reached.
Alert (*Only if Alert is selected as Log Level*)	If you selected Alert as the Log Level, select the Alert element that is used when an event triggers an alert. The Alert elements can be used for matching in Alert Policies.

5. Click **OK**.

What's Next?

▶ Refresh the engine's policy to transfer the configuration changes.

Setting Connection Timeouts

Prerequisites: Creating New Engine Elements

You can define general timeouts for removing idle connections from the state table, including non-TCP communications that are handled like connections. The timeout prevents wasting engine resources on storing information about abandoned connections. Timeouts are a normal way to clear traffic information with protocols that have no closing mechanism. The communicating client and server also have timeouts for closing inactive connections.

You can set timeouts by protocol and by TCP connection state. Idle timeouts set in Access rules override these global settings.

Timeouts do not affect active connections. The connections are kept in the state table as long as the interval of packets within a connection is shorter than the timeouts set.

Caution – Setting excessive timeouts for a large number of connections consumes resources excessively and may disturb the operation of the engine.

▼ To set connection timeouts

1. Right-click a Firewall, IPS, or Layer 2 Firewall element and select **Properties**. The Properties dialog opens.
2. Switch to the **Advanced** tab.
3. Click **Idle Timeouts** in the Traffic Handling section. The Idle Timeouts dialog opens.
4. Click the **Timeout(s)** column and enter the timeout value for the protocol in seconds. The default values for the predefined protocols are:
   - **UDP**: 5
   - **Other**: 180
   - **TCP**: 1800
   - **ICMP**: 5
5. (*Optional*) Click **Add** to add a new protocol to the list and enter the timeout for the protocol.
6. Click **OK**.

What's Next?

▶ Refresh the engine's policy to transfer the configuration changes.

Configuring Default SYN Rate Limits

Prerequisites: Creating New Engine Elements

You can configure SYN Rate Limits to reduce the risk of SYN flood attacks against the Firewall, IPS engine, 2 Firewall, Master Engine, or Virtual Security Engine.

The global SYN Rate Limits that you define in the security engine properties as explained below are applied as default settings on all interfaces. You can also define SYN Rate Limits that override the global settings in each interface's properties. For more information, see Configuring Advanced Interface Properties for Firewalls (page 436), Configuring Advanced Interface Properties for IPS Engines (page 469), Configuring Advanced Interface Properties for Layer 2 Firewalls (page 482), Configuring Advanced Interface Properties for Master Engines (page 495), and Configuring Advanced Interface Properties for Virtual Security Engines (page 506).

▼ To configure default SYN Rate Limits

1. Right-click an engine element and select **Properties**. The Properties dialog opens.
2. Switch to the **Advanced** tab.
3. Click **SYN Rate Limits** in the Traffic Handling section. The Default SYN Rate Limits Settings dialog opens.
4. Select the **SYN Rate Limits Mode**:
   - **Off**: SYN Rate Limits are disabled. This is the default setting.
   - **Automatic**: This is the recommended mode. The engine automatically calculates the number of **Allowed SYNs per Second** (the number of allowed SYN packets per second) and the **Burst Size** (the number of allowed SYNs before the engine starts limiting the SYN rate) for the interface based on the engine's capacity and memory size.
   - **Custom**: Enter the desired values for **Allowed SYNs per Second** and **Burst Size**. We recommend that the Burst Size be at least one tenth of the **Allowed SYNs per Second** value. If the Burst Size is too small, SYN Rate Limits do not work. For example, if the value for Allowed SYNs per Second is 10000, the Burst Size must be at least 1000.

Caution – The recommended values for the SYN Rate Limits depend on your network environment. If the Custom settings are not carefully configured, the capacity of the security engine may suffer or SYN Rate Limits may not work correctly.

5. Click **OK**.

What's Next?

▶ Refresh the engine's policy to transfer the changes.

Related Tasks

▶ Configuring Default Log Handling Settings (page 580)

Configuring Default Log Handling Settings

Prerequisites: Creating New Engine Elements

Log Handling settings allow you to adjust logging when the log spool on the Firewall, IPS, and Layer 2 Firewall engines fills up. Logs are spooled locally when the Log Server is not available.

You can also configure Log Compression to save resources on the engine. By default, each generated Antispoofing and Discard log entry is logged separately and displayed as a separate entry in the Logs view. Log Compression allows you to define the maximum number of separately logged entries. When the defined limit is reached, a single Antispoofing log entry or Discard log entry is logged. The single entry contains information on the total number of the generated Antispoofing log entries or Discard log entries. After this, logging returns to normal and all the generated entries are once more logged and displayed separately.

The general Log Compression settings you define in the engine properties are applied as default settings on all interfaces. You can also define Log Compression and override the global settings in each interface's properties. For more information, see Configuring Advanced Interface Properties for Firewalls (page 436), Configuring Advanced Interface Properties for IPS Engines (page 469), and Configuring Advanced Interface Properties for Layer 2 Firewalls (page 482).

▼ To configure Log Handling settings

1. Right-click a Firewall, IPS, or Layer 2 Firewall element and select **Properties**. The Properties dialog opens.
2. Switch to the **Advanced** tab.
3. Click **Log Handling** in the System Parameters section. The Log Handling Settings dialog opens.
4. Select the **Log Spooling Policy** to define what happens when the engine's log spool becomes full:
   - **Stop Traffic**: The node goes offline. If the node is part of a cluster, the other nodes of the cluster take over the connections.
   - **Discard Log**: Log entries are discarded in four stages, according to available space. Monitoring data is discarded first, followed by log entries marked as Transient and Stored, and finally log entries marked as Essential. The engine keeps processing traffic. This is the default option.
5. (*Optional*) In the Log Compression table, enter values for Antispoofing entries (*Firewall only*) and for Discard entries.
   - Do not enable Log Compression if you want all the Antispoofing and Discard entries to be logged as separate log entries (for example, for reporting purposes or for statistics).
   - **Log Rate (Entries/s):** the maximum number of entries per second. The default value for Antispoofing entries is 100 entries/s. By default, Discard log entries are not compressed.
   - **Burst Size (Entries):** the maximum number of matching entries in a single burst. The default value for Antispoofing entries is 1000. By default, Discard log entries are not compressed.
6. Click **OK**.

What's Next?

▶ Refresh the engine's policy to transfer the changes.

Chapter 35

Setting up SNMP for Engines

SNMP is a standard protocol that different equipment can use to send network management related information to each other. Security Engines can be set up to send out SNMP traps to external equipment.

The following sections are included:

Getting Started with SNMP Configuration

Prerequisites: None

Security engines can send SNMP traps on system events, such as when a test fails, for example, to a central network monitoring system. An SNMP Agent element defines the settings according to which the engines send SNMP trap messages to compatible external software. A single SNMP Agent can be used on multiple engines.

What Do I Need To Know Before I Begin?

The SNMP Agent supports SNMPv1 (RFC1157), SNMPv2c (RFCs 1901 and 3416), and SNMPv3 (RFC 3414). Check the documentation of the receiving software for information on which version to use.

The MIBs are included on the Management Center DVD. For more information on the traps the SNMP Agent sends, see SNMP Traps and MIBs (page 1193)

What's Next?

- ▶ Configuring SNMP Version 1 or 2c
- ▶ Configuring SNMP Version 3 (page 583)

Configuring SNMP Version 1 or 2c

Prerequisites: None

▼ To configure the monitoring settings for SNMPv1 and v2c

1. Select **Configuration→Configuration→Security Engine** from the menu. The Security Engine Configuration view opens.
2. Browse to **Other Elements→Engine Properties→SNMP Agents**.
3. Right-click **SNMP Agents** and select **New SNMP Agent**. The SNMP Agent Properties dialog opens.
4. Enter a unique **Name** for the SNMP Agent and select **v1** or **v2c** as the **Version**.
5. (*Optional*) Click **Add** in the Monitoring section and enter the community string in the dialog that opens. The community string is used for authentication in monitoring.
6. (*Optional*) Enter the UDP **Listening Port** number that the SNMP agent listens to. The default port is 161.
7. Enter the **Contact** information for the person responsible for the engines. This string is returned to queries from the SNMPv2-MIB object.

What's Next?

- ▶ Configuring What Triggers SNMP Traps (page 583)

Configuring SNMP Version 3

Prerequisites: None

▼ To configure the monitoring settings for SNMPv3

1. Select **Configuration→Configuration→Security Engine** from the menu. The Security Engine Configuration view opens.
2. Browse to **Other Elements→Engine Properties**.
3. Right-click **SNMP Agents** and select **New SNMP Agent**. The SNMP Agent Properties dialog opens.
4. Enter a unique **Name** for the SNMP Agent and select **v3** as the **Version**.
5. Click **Add** in the **User Names** section. The SNMP User Properties dialog opens.
6. Enter the **User Name**, select the Authentication and Privacy options, and click **OK**.
7. Repeat from Step 5 to add more users as necessary.
8. Click **Add** in the Monitoring section and select the user for monitoring in the dialog that opens. Repeat as necessary to select more users for monitoring.
9. (*Optional*) Enter the UDP **Listening Port** number that the SNMP agent listens to. The default port is 161.
10. Enter the **Contact** information for the person responsible for the engines. This string is returned to queries from the SNMPv2-MIB sysContact object.

What's Next?

▶ Configuring What Triggers SNMP Traps

Configuring What Triggers SNMP Traps

Prerequisites: Configuring SNMP Version 1 or 2c / Configuring SNMP Version 3

The trap parameters control where and how SNMP traps are sent. The Destinations field is important. If it is left empty, no traps are sent, and the other trap parameters are ignored. If the Destinations field has a value, the rest of the trap parameters must also have a value.

▼ To activate a specific trap

1. Specify a community or user name in the Traps section:
   - In SNMPv1 or SNMPv2c, enter the **Community** string.
   - In SNMPv3, select a **User Name**.
2. Click **Add** in the Traps section and enter the IP address and UDP port where the traps are sent. The default port is 162.

3. Select the events for which you want to set a trap in the Active Traps section. The possible events are:
   - **Boot** (for example, startup of the agent process by minit)
   - **Shutdown**
   - **Going Online**
   - **Going Offline**
   - **Policy Applied**
   - **User Login** (via console or with SSH)
   - **Hardware Alerts** (for appliances that support hardware monitoring).
4. Click **OK**. The SNMP Agent is created.

In addition to these general events, the Tester on each engine can send SNMP traps when a test fails.

What's Next?

▶ Activating the SNMP Agent on Engines

Activating the SNMP Agent on Engines

Prerequisites: See Getting Started with SNMP Configuration

The SNMP Agent is responsible for SNMP-related tasks on the engines.

▼ To activate the SNMP Agent

1. Open the properties of the Firewall, IPS engine, or Layer 2 Firewall element.
2. Select the **SNMP Agent** you want to activate.
3. Define the **SNMP Location** string that is returned on queries to the SNMPv2-MIB or SNMPv2-MIB-sysLocation object.
4. Click **OK**.

What's Next?

▶ Refresh the engine's policy to transfer the changes.

Related Tasks

▶ If you want to create a Test Entry that sends a notice using the SNMP Agent, see Adding Engine Tests (page 522)

▶ For more information on the traps the SNMP Agent sends, see SNMP Traps and MIBs (page 1193)

ROUTING

In this section:

Chapter 36

Configuring Routing

Routing and the related antispoofing configuration is done entirely in the Management Client. For the most part, this information is automatically filled in according to the interfaces defined for each engine.

The following sections are included:

Getting Started with Routing

Prerequisites: None

How Routing Works

Routing is configured mainly in the Management Client in the tree displayed in the *Routing view*. Basic routing information for directly connected networks is added automatically to both routing and antispoofing based on IP addresses you have defined for the engines' interfaces. You must manually define routing for firewall Tunnel Interfaces used in the Route-Based VPN. IPS engines and Layer 2 Firewalls do not route traffic. You may need to define a default route for IPS engines and Layer 2 Firewalls if other system components are not located on a directly connected network.

Antispoofing is used only on the Firewall. Antispoofing configuration is added and updated automatically based on the routing configuration. Firewalls always enforce antispoofing, which by default allows a connection to originate only from the interface that has a routing entry for the connection's source IP address.

Limitations

- Selection of which traffic is routed through policy-based VPNs is not determined by the basic routing configuration; routing is checked after policy-based VPN traffic is encapsulated inside encrypted packets with different source and destination IP address information.
- IPv6 addresses are not supported in DHCP relay, multicast routing, and policy routing.
- Multicast routing is not supported on Wireless Interfaces.
- IGMP Proxy mode for multicast routing is not supported on Tunnel Interfaces.
- Dynamic routing cannot be configured using the Management Client. You must configure dynamic routing on the Firewall engine command line. See Configuring Dynamic Routing (page 236).

Configuration Overview

1. Add the default route. See the relevant section below:
   - Adding Routes for Firewalls (page 589)
   - Adding Routes for Master Engines (page 598)
   - Adding Routes for Virtual Firewalls (page 599)
   - Defining a Multi-Link Route (page 600)
   - Adding Routes for IPS Engines and Layer 2 Firewalls (page 605)
2. Add routes to networks that are not directly connected, but require a next-hop gateway. See the relevant section below:
   - Adding Routes for Firewalls (page 589)
   - Adding Routes for Master Engines (page 598)
   - Adding Routes for Virtual Firewalls (page 599)
3. Add routes to networks that are reachable through the Tunnel Interfaces used in the Route-Based VPN. See Defining Routing for the Route-Based VPN (page 604).

Once you have configured routing you can monitor it as explained in Monitoring Connections, Blacklists, VPN SAs, Users, and Routing (page 111).

Related Tasks

- ▶ To configure the firewall to relay DHCP messages, proceed to Routing DHCP Messages (page 590).
- ▶ To define a route for multicast traffic, proceed to Routing Multicast Traffic (page 593).
- ▶ To define a route that sends connections to different destinations based on both source and destination, proceed to Defining Policy Routing (page 596).

Adding Routes for Firewalls

Prerequisites: None

Firewalls route network traffic. You must add a default route and any routes through next-hop gateways to networks that are not directly connected to the firewalls. On firewall clusters, the routing information is configured simultaneously for all cluster nodes and all nodes always have identical routing tables. If you are configuring the Route-Based VPN, you must manually define the routing for the Tunnel Interfaces.

What’s Next?

- ▶ To define only one route to the same destination or to define a single route for interfaces that belong to an aggregated link, proceed to Defining a Single-Link Route for a Firewall.
- ▶ To define two or more alternative routes to the same destination, proceed to Defining a Multi-Link Route (page 600).
- ▶ To define routing for Tunnel Interfaces used in the Route-Based VPN, proceed to Defining Routing for the Route-Based VPN (page 604).

Defining a Single-Link Route for a Firewall

Follow these instructions to add a default route or static routes to a particular network. If there are two or more alternative routes to the same destination, create a Multi-Link route instead. See Defining a Multi-Link Route (page 600).

Routing decisions are made for each packet by matching from the most specific route definition to the most general. However, any policy routing entries you define override the other routing information. For packets subject to address translation, routing is always done after NAT using the translated addresses.

▼ To define a single-link route

1. Select **Configuration→Routing** from the menu. The Routing view opens.
2. Select the firewall for which you want to configure routing from the list above the Routing tree.
3. Expand the tree of the correct interface so that the Network element is displayed below the correct interface.
4. Add a Router to the Network element.
   - If the Router element you want to add already exists, drag and drop it from the Resources panel onto the Network in the Routing tree.
   - If the Router element does not exist, create a new Router element by right-clicking the Network element under the correct interface in the Routing tree and selecting

New→Router. The Router element for dynamic interfaces is marked as “DHCP Assigned”.

5. Add the element that contains the IP addresses that are routed through this interface:
   - To add a default route, right-click the Router element you just added in the Routing tree and select **New→Any Network**. The (existing) default Any Network element is added below the Router in the Routing tree.
   - To add a route using some element that is already configured, drag and drop the element from the Resources panel.
   - To create a new Host or Network element, use the right-click menu for the Router element you just added in the routing tree.

Note – Placing the Any network element behind two different basic router elements does not create true router redundancy and load balancing. Use NetLink elements instead. See Defining a Multi-Link Route.

What’s Next?

▶ Install the policy to transfer the changed configuration to the firewall.

Routing DHCP Messages

The firewall can relay DHCP messages. If DHCP messages are routed through the firewall (from a network segment to some other, isolated segment), you must enable DHCP relay on the firewall interface properties for the interface where the DHCP requests are originating from (client’s network).

Only IPv4 addresses are supported in DHCP relay. You can enable DHCP relay only on firewall interfaces that have at least one IPv4 address.

This DHCP relay configuration does not relay DHCP messages from the IPsec VPN Client. See Managing VPN Client IP Addresses (page 1000).

What’s Next?

▶ If you have already configured the DHCP server that you want to use for relaying DHCP messages, start by Enabling DHCP Relay (page 592).

▶ Otherwise, start by Defining a DHCP Server (page 591).

Defining a DHCP Server

A DHCP Server dynamically assigns IP addresses.

▼ To define a DHCP Server element

1. Select **Configuration→Configuration→Security Engine** from the menu. The Security Engine Configuration view opens.
2. Expand the **Network Elements** tree.
3. Right-click **Servers** and select **New→DHCP Server**. The DHCP Server Properties dialog opens.

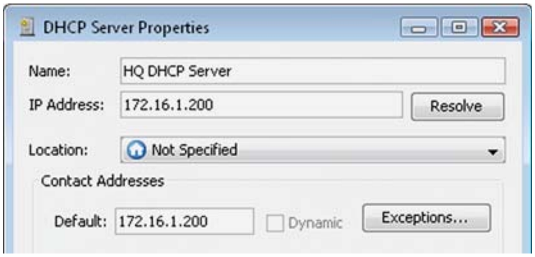

DHCP Server Properties

Name: HQ DHCP Server

IP Address: 172.16.1.200

Location: Not Specified

Contact Addresses

Default: 172.16.1.200 □ Dynamic

4. Enter a unique **Name** and **IP Address**.

Note – Only IPv4 addresses are supported for DHCP relay in the Management Center.

5. A **Location** and **Contact Address** is needed if there is a NAT device between a firewall and the server, so that the firewall cannot connect directly to the IP address defined for the interface. See Defining Locations (page 66) and Defining Contact IP Addresses (page 66) for more information.
6. Click **OK** to apply the changes to the DHCP Server properties.

What’s Next?

- Configure the firewall to relay DHCP messages as explained in Enabling DHCP Relay (page 592).

Related Tasks

- To use the DHCP server to assign virtual IP addresses to VPN clients, see Configuring the Gateway for Virtual IP Address Clients (page 1002).

Enabling DHCP Relay

You must select which firewall interfaces perform DHCP relay. Activate the relay on the interface towards the DHCP clients.

▼ To enable DHCP relay

1. Right-click the Firewall element and select **Properties**. The Firewall Properties dialog opens.
2. Switch to the **Interfaces** tab.
3. Right-click a physical or a VLAN interface and select **Edit Physical Interface** or **Edit VLAN Interface**. The Interface Properties dialog opens.

Note – When configuring VLAN interfaces, you must set the DHCP relay separately for each VLAN.

4. Switch to the **DHCP** tab, and select **DHCP Relay** as the **DHCP Mode**.
5. Select the correct DHCP server from the list of servers on the left and click **Add**. The DHCP server is added to the list on the right.
   - You can also create a new DHCP Server element by clicking the New icon above the left panel. See Defining a DHCP Server (page 591).
6. Adjust the maximum allowed packet size if you have a specific need to do so.
7. (*Optional*) Select the CVI or IP address you want to use for **DHCP Relay**.
8. Click **OK** in both open dialogs.

Access rules required to relay DHCP messages are not enabled by default, but there is a ready-made DHCP Relay Sub-Policy that you can start using.

What's Next?

▶ To allow DHCP relay in the Access rules, proceed to Activating the DHCP Relay Sub-policy (page 592).

Activating the DHCP Relay Sub-policy

DHCP relay is not allowed in the Firewall Template. However, there is a ready-made DHCP relay sub-policy that you can refer to from your own Firewall Policy.

Note – The sub-policy contains rules for both local DHCP relay between internal networks and the DHCP relay for VPN clients. If just one of these DHCP relay features is active, the rules for the other feature are invalid and ignored when the policy is installed.

▼ To activate the DHCP relay sub-policy

1. Open the Firewall Policy or Template Policy for editing.
2. Add a new IPv4 Access rule with the following properties:

Source	Destination	Service	Action
ANY	ANY	BOOTPC (UDP) BOOTPS (UDP)	Jump DHCP Relay

Caution – If you use a Zone element in the Source cell of a rule that jumps to the DHCP Relay sub-policy, you must also add the Node-internal Zone element to the Source cell, or allow connections from the Node-internal Zone elsewhere in the policy. Otherwise, DHCP relay does not work.

What's Next?

- Refresh the policies of all affected engines to transfer the changed configuration.

Related Tasks

- To monitor the DHCP relay more closely during testing, you can activate DHCP relay diagnostic logging on the firewall. See Enabling/Disabling Firewall/VPN Diagnostics (page 222).

Routing Multicast Traffic

The Firewall supports static multicast and IGMP-based multicast forwarding (IGMP proxying). IGMP-based multicast forwarding is only supported in tree topology networks. Only IPv4 addresses are supported in multicast routing.

Static multicast allows you to configure static routes for multicast traffic between a source IP address and firewall interface pair, and a destination (multicast) IP address and firewall interface pair. Static multicast is often used for enduring configurations, such as mutually agreed multicast traffic between organizations (for example, multicast news feeds and video conferences).

In IGMP-based multicast forwarding, the firewall maintains a list of hosts that are subscribed to the multicast host group. It forwards the traffic sent to the multicast IP address only to the currently subscribed hosts.

Note – In addition to configuring the Firewall, routers and other network devices must be configured to allow IP multicasting along the path to the client machines.

What's Next?

- To define static IP multicast entries, continue by Defining Static Multicast (page 594).
- To configure IGMP Proxy multicast routing, continue by Defining IGMP-Based Multicast Forwarding (page 595).

Defining Static Multicast

Note – Make sure your IPv4 Access rules allow the static multicast traffic to pass through the firewall (see Editing Access Rules (page 672)).

▼ To define static multicast

1. Open the Properties for the Firewall element for which you want to define static multicast routing.
2. Switch to the **Interfaces** tab.
3. Click the **Multicast Routing** button below the interface table. The Multicast Routing Properties dialog opens.
4. Select **Static** as the **Multicast Routing Mode**. The options for the mode are enabled.

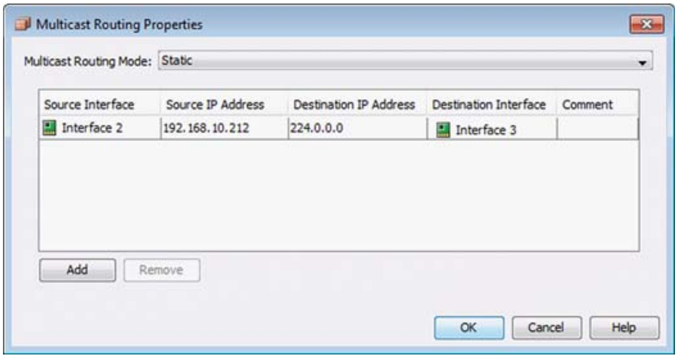

Multicast Routing Mode: Static

Source Interface	Source IP Address	Destination IP Address	Destination Interface	Comment
Interface 2	192.168.10.212	224.0.0.0	Interface 3	

5. Click **Add**. A new entry appears in the table.
6. Configure the values for the entry:

Cell	How to Configure
Source Interface	Select the Firewall interface to use for multicast routing.
Source IP Address	Enter the unicast IP address of the multicast source.
Destination IP Address	Enter the multicast destination IP address. The destination address must be within the multicast range of 224.0.0.0 to 239.255.255.255.
Destination Interface	Right-click **Destination** Interface and select **Edit Destination** Interface to select the interface(s) where you want this multicast traffic forwarded.

Note – Wireless Interfaces cannot be used in multicast routing.

7. Repeat from Step 5 to define any additional static IP multicast entries.
8. Click **OK**.

What's Next?

▶ Refresh the policies of all affected engines to transfer the changed configuration.

Defining IGMP-Based Multicast Forwarding

IGMP-based multicast forwarding (IGMP proxying) is implemented on the Firewall based on RFC 4605. IGMP-based multicast forwarding is only supported in tree topology networks. RFC 4605 includes support for source-specific multicast (SSM) with IGMP version 3. SSM is not supported on the Firewall. However, you can configure Access rules that filter multicast traffic based on the source.

The firewall maintains a membership database of the subscriptions from the downstream networks and sends unsolicited reports or leaves on the upstream interface when the subscription database changes. It also sends IGMP membership reports when queried on the upstream interface.

▼ To define IGMP-based multicast forwarding

1. Right-click the Firewall element for which you want to define IGMP-based multicast forwarding and select **Properties**. The Properties dialog opens.
2. Switch to the **Interfaces** tab.
3. Click the **Multicast Routing** button below the interface table. The Multicast Routing Properties dialog opens.

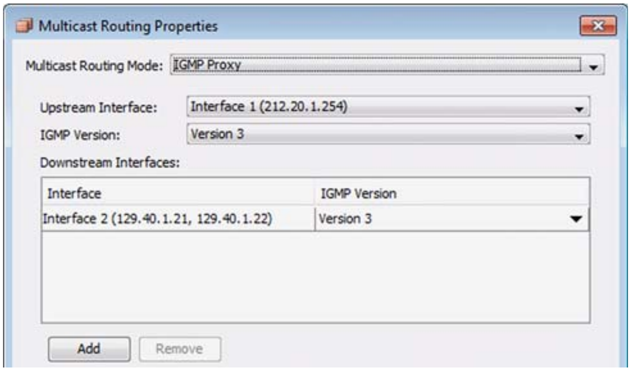

Multicast Routing Properties

Multicast Routing Mode: IGMP Proxy

Upstream Interface: Interface 1 (212.20.1.254)

IGMP Version: Version 3

Downstream Interfaces:

Interface	IGMP Version
Interface 2 (129.40.1.21, 129.40.1.22)	Version 3

4. Select **IGMP Proxy** as the **Multicast Routing Mode**. The options for the mode are enabled.
5. (*Optional*) Select the **Upstream Interface** and the **IGMP Version** for it.
   - If the multicast servers and the hosts are in the local networks, or if you want to limit the multicast to the local networks, it is not necessary to define the Upstream Interface. In that case, leave **Not Set** selected for **Upstream Interface**.
   - (*Firewall Clusters only*) You can only select as the Upstream Interface an interface that has a Cluster Virtual IP Address (CVI). You cannot select a Heartbeat Interface as the Upstream Interface.
   - The default IGMP version is version 3. You may need to select another IGMP version, for example, to troubleshoot multicast accessibility on hosts, or if some hosts use a lower IGMP version.
6. Click **Add** to define **Downstream Interfaces**. A new entry appears in the table.
   - The firewall periodically queries the downstream networks for hosts that want to join or leave the multicast host group.
7. Click the **Interface** cell and select the Downstream Interface from the list.
   - You can use each interface only once in the IGMP proxy configuration.
   - (*Firewall Clusters only*) The interface that you select as a Downstream Interface must have Node Dedicated IP Addresses (NDIs). It cannot be a Heartbeat Interface. It is recommended that the Node Dedicated IP Addresses increase in the same order on each

node: for example, 192.168.1.10 and 192.168.2.10 for node A, and 192.168.1.11 and 192.168.2.11 for node B.

Note – The downstream interfaces must have the lowest IP addresses among all the IGMP queriers in the local networks.

8. Click the **IGMP Version** cell and select the IGMP version for the downstream interface. The default version is IGMP version 3. You may need to select another IGMP version, for example, to troubleshoot multicast accessibility on hosts, or if some hosts use a lower IGMP version.
9. Repeat from Step 6 to define more Downstream Interfaces.
10. Click **OK**.

Note – Make sure your IPv4 Access rules allow this traffic to pass through the firewall (see Editing Access Rules (page 672)).

What's Next?

▶ Refresh the policies of all affected engines to transfer the changed configuration.

Defining Policy Routing

Policy routing is used to resolve duplicate IP addressing issues. Policy routing allows you to route traffic based on both source and destination IP address. Policy routing is useful if the same IP address is in use on more than one physical host in different networks and it is not possible to resolve the situation by changing IP address or by using NAT. In most systems, there is no need to add any policy routing entries.

Note – Policy routing is only used if duplicate IP addresses are used in different networks. To configure alternative routes to the same destination, use Multi-Link routing (see Defining a Multi-Link Route (page 600)).

Policy routing entries are applied before the regular routes defined in the Routing tree (overriding those configurations if matches are found), and they are processed exactly in the order specified in the Policy Routing window; the first matching policy routing entry is applied to a connection and any further entries are ignored.

Only IPv4 addresses are supported in policy routing.

▼ To configure policy routing

1. Right-click the Firewall element for which you want to define policy routing and select **Properties**. The Properties dialog opens.

2. Switch to the **Advanced** tab.

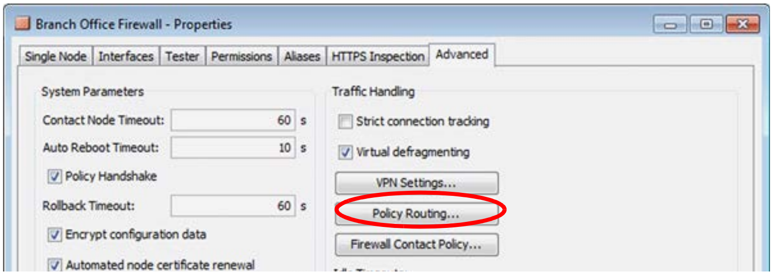

Contact Node Timeout: 60 s

Auto Reboot Timeout: 10 s

☑ Policy Handshake

Rollback Timeout: 60 s

☑ Encrypt configuration data

☑ Automated node certificate renewal

□ Strict connection tracking

☑ Virtual defragmenting

3. Click the **Policy Routing** button. The Policy Routing dialog opens.

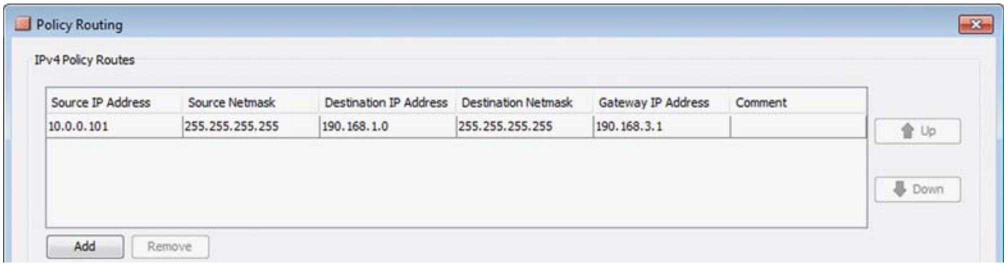

Source IP Address	Source Netmask	Destination IP Address	Destination Netmask	Gateway IP Address	Comment
10.0.0.101	255.255.255.255	190.168.1.0	255.255.255.255	190.168.3.1	

4. Click **Add**. A new entry appears in the table.

5. Double-click the cells and enter the required routing information (note that Routing is done after NAT):

Setting	Description
Source IP Address	This is always something other than the default 0.0.0.0 that matches any IP address (such configurations can be handled more easily with the normal routing tools in the Routing view).
Source Netmask	Enter the netmask for the source IP address.
Destination IP Address	Enter the destination IP address.
Destination Netmask	Enter the netmask for the destination IP address.
Gateway IP Address	The IP address of the device to which packets that match the source/ destination pair are forwarded.

6. Repeat from Step 4 to create any additional policy routing entries, then click **OK**. The Policy Routing dialog closes.

7. Click **OK**. The Properties dialog closes.

What's Next?

- Refresh the policies of all affected engines to transfer the changed configuration.

Related Tasks

- Policy routing entries are not automatically added to the Antispoofing view, so you may need to update the antispoofing information manually. See Modifying Antispoofing (page 606).

Adding Routes for Master Engines

Prerequisites: Creating New Master Engine Elements

Basic routing information for directly connected networks is added automatically to both routing and antispoofing based on IP addresses you have defined for the engines' interfaces. You must add a default route and any routes through next-hop gateways to networks that are not directly connected to the Master Engine.

What's Next?

- ▶ To define only one route to the same destination or to define a single route for interfaces that belong to an aggregated link, proceed to Defining a Single-Link Route for a Master Engine.
- ▶ To define two or more alternative routes to the same destination, proceed to Defining a Multi-Link Route (page 600).
- ▶ To define routing for Tunnel Interfaces used in the Route-Based VPN, proceed to Defining Routing for the Route-Based VPN (page 604).

Defining a Single-Link Route for a Master Engine

Follow these instructions to add a default route or static routes to a particular network. If there are two or more alternative routes to the same destination, create a Multi-Link route instead. See Defining a Multi-Link Route (page 600).

Routing decisions are made for each packet by matching from the most specific route definition to the most general. However, any policy routing entries you define override the other routing information. For packets subject to address translation, routing is always done after NAT using the translated addresses.

▼ To define a single-link route for a Master Engine

1. Select **Configuration→Routing** from the menu. The Routing view opens.
2. Select the Master Engine for which you want to configure routing from the list above the Routing tree.
3. Expand the tree of the correct interface so that the Network element is displayed below the correct interface.
4. Add a Router to the Network element.
   - If the Router element you want to add already exists, drag and drop it from the Resources panel onto the Network in the Routing tree.
   - If the Router element does not exist, create a new Router element by right-clicking the Network element under the correct interface in the Routing tree and selecting **New→Router**. The Router element for dynamic interfaces is marked as "DHCP Assigned".
5. Add the element that contains the IP addresses that are routed through this interface:
   - To add a default route, right-click the Router element you just added in the Routing tree and select **New→Any Network**. The (existing) default Any Network element is added below the Router in the Routing tree.
   - To add a route using some element that is already configured, drag and drop the element from the Resources panel.
   - To create a new Host or Network element, use the right-click menu for the Router element you just added in the routing tree.

Note – Placing the Any network element behind two different basic Router elements does not create true router redundancy and load balancing. Use NetLink elements instead. See Defining a Multi-Link Route (page 600).

What's Next?

- ▶ Install the policy to transfer the changed configuration to the Virtual Firewall.

Adding Routes for Virtual Firewalls

Prerequisites: Creating New Virtual Security Engines

Basic routing information for directly connected networks is added automatically to both routing and antispoofing based on IP addresses you have defined for the engines' interfaces. You must add a default route and any routes through next-hop gateways to networks that are not directly connected to the Virtual Firewall. If you are configuring the Route-Based VPN, you must manually define the routing for the Tunnel Interfaces.

What's Next?

- ▶ To define only one route to the same destination or to define a single route for interfaces that belong to an aggregated link, proceed to Defining a Single-Link Route for a Virtual Firewall.
- ▶ To define two or more alternative routes to the same destination, proceed to Defining a Multi-Link Route (page 600).
- ▶ To define routing for Tunnel Interfaces used in the Route-Based VPN, proceed to Defining Routing for the Route-Based VPN (page 604).

Defining a Single-Link Route for a Virtual Firewall

Follow these instructions to add a default route or static routes to a particular network. If there are two or more alternative routes to the same destination, create a Multi-Link route instead. See Defining a Multi-Link Route (page 600).

Routing decisions are made for each packet by matching from the most specific route definition to the most general. However, any policy routing entries you define override the other routing information. For packets subject to address translation, routing is always done after NAT using the translated addresses.

▼ To define a single-link route for a Virtual Firewall

1. Select **Configuration→Routing** from the menu. The Routing view opens.
2. Select the Virtual Firewall for which you want to configure routing from the list above the Routing tree.
3. Expand the tree of the correct interface so that the Network element is displayed below the correct interface.

4. Add a Router to the Network element.
   - If the Router element you want to add already exists, drag and drop it from the Resources panel onto the Network in the Routing tree.
   - If the Router element does not exist, create a new Router element by right-clicking the Network element under the correct interface in the Routing tree and selecting **New→Router**. The Router element for dynamic interfaces is marked as “DHCP Assigned”.
5. Add the element that contains the IP addresses that are routed through this interface:
   - To add a default route, right-click the Router element you just added in the Routing tree and select **New→Any Network**. The (existing) default Any Network element is added below the Router in the Routing tree.
   - To add a route using some element that is already configured, drag and drop the element from the Resources panel.
   - To create a new Host or Network element, use the right-click menu for the Router element you just added in the routing tree.

Note – Placing the Any network element behind two different basic Router elements does not create true router redundancy and load balancing. Use NetLink elements instead. See Defining a Multi-Link Route (page 600).

What’s Next?

▶ Install the policy to transfer the changed configuration to the Virtual Firewall.

Defining a Multi-Link Route

When you use Multi-Link routing, traffic is routed through different network connections (NetLinks) according to traffic conditions, the availability of the different links, and explicitly configured preferences you set. Most often, Multi-Link is used for Internet connections, but Multi-Link can be used to provide alternative routes to any other network as well. You can define a Multi-Link route for firewalls and Virtual Firewalls.

Configuration Overview

1. Create a NetLink for each alternative route as instructed in Creating NetLinks (page 601).
2. Add Networks below the NetLink(s) in the Routing tree to define a route as explained in Adding a Multi-Link Route (page 603).

What’s Next?

▶ Creating NetLinks (page 601)

Creating NetLinks

NetLink elements are used to represent alternative routes that lead to the same destination IP addresses. NetLinks usually represent Internet connections, but can be used for other communications links as well.

Although the Management Client may not require all settings to be filled in at this point, we recommend that you fill in all settings as instructed here to achieve a fully functional configuration and to avoid having to readjust the settings later. To monitor the status of the network links, you must define the *probe IP addresses* in the NetLink properties.

> **Note – Only NetLinks that are using in an Outbound Multi-Link element are probed. Status monitoring is not available for NetLinks that are only used in Routing. See Getting Started with Outbound Traffic Management (page 610) for more information about Outbound Multi-Link configuration.**

▼ To define a NetLink element

1. Select **Configuration→Routing** from the menu. The Routing view opens.
2. Select the firewall or Virtual Firewall for which you want to configure routing from the list above the Routing tree.
3. Expand the tree under the interface under which you want to create the NetLink.
4. Right-click the network to which the NetLink belongs and select **New→Static NetLink** or **New→Dynamic NetLink**. The NetLink Properties dialog opens.

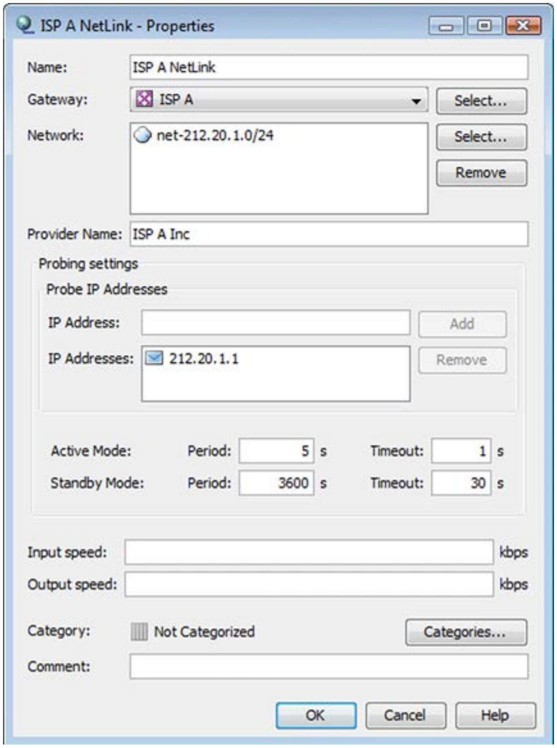

ISP A NetLink - Properties

Name: ISP A NetLink

Gateway: ISP A

Network: net-212.20.1.0/24

Provider Name: ISP A Inc

Probing settings

Probe IP Addresses

IP Address:

IP Addresses: 212.20.1.1

Active Mode: Period: 5 s Timeout: 1 s

Standby Mode: Period: 3600 s Timeout: 30 s

Input speed: kbps

Output speed: kbps

Category: Not Categorized

Comment:

5. Give the NetLink a unique **Name**.

6. (*Static NetLink only*) Select the next-hop **Gateway** the NetLink uses.
   - The **Select** button opens a dialog that allows you to create a new element if needed.
   - The Gateway is usually added as a Router element, but you can also choose another Firewall according to your network configuration.
7. (*Static NetLink only*) Click **Select** next to the **Network** field and define the address space that the ISP has assigned to your organization.

Note – Do not select the Any Network element for the NetLink. Select or create Network elements that represent the specific public IP addresses assigned by your ISP and routed to the Internet provider's router.

8. (*Optional*) Enter the **Provider Name** of your ISP.
9. (*Optional*) Enter or adjust the Probing Settings:

Setting	Description
IP Address	Enter an IP address and click **Add**. The IP addresses that are probed with ICMP echo requests (ping) to determine if the link is up. Repeat this for each IP address you want to add. We recommend entering more than one address to avoid excluding the link in case the host that is probed goes down.
Period	Define how often the NetLink is probed when it is in Active/Standby mode. Leave the setting for Standby Mode as 0 if you prefer not to test this link when it is on standby (this is meant to minimize costs for links that are charged based on use rather than at a fixed rate).
Timeout	Define how long the firewall waits before it considers the probe failed. Change the setting for Standby Mode to 0 if you prefer not to test this link when it is on standby.

Note – Select IP addresses that give reliable results of end-to-end connectivity to destinations that need to be reached through the link (such as a remote server across the internet instead of an internal router or server of your ISP in the case of an Internet connection).

10. (*Optional*) If you plan to use the ratio-based load balancing method, fill in the **Input Speed** and **Output Speed** based on the real-life bandwidth this network connection provides. These values are used to calculate how much traffic each link receives in relation to the other links.
11. Click **OK**.

What's Next?

- ▶ Repeat these steps to add any other NetLinks to the Routing tree.
- ▶ If you are finished adding NetLinks, define a route through the NetLinks as explained in Adding a Multi-Link Route (page 603).

Adding a Multi-Link Route

Note – The network interfaces for the NetLinks must have a Node Dedicated IP Address defined for all nodes in clusters. Otherwise, the IP address of the interface marked as the default IP address for outgoing connections is used, which may lead to incorrect load balancing. See Configuring Firewall Cluster IP Addresses (page 444).

▼ To add a Multi-Link route

1. Select **Configuration→Routing** from the menu. The Routing view opens.
2. Select the firewall or Virtual Firewall for which you want to configure routing from the list above the Routing tree.
3. Browse to the NetLink through which you want to route traffic in the Routing Tree.
4. Add the element that contains the IP addresses that are routed through this interface:
   - To add a default route, right-click the NetLink element you just added in the Routing tree and select **New→Any Network**. The (existing) default Any Network element is added below the NetLink in the Routing tree.
   - To add a route using some element that is already configured, drag and drop the element from the Resources panel.
   - To create a new Host or Network element, use the right-click menu for the NetLink element you just added in the Routing tree.

Caution – If you use Multi-Link with IGMP proxy multicast routing, make sure that you do not create routing loops. If you add a NetLink to the upstream interface, do not add a NetLink to any downstream interface.

What's Next?

- ▶ Repeat these steps for any additional NetLinks.
- ▶ Otherwise, the Multi-link Routing configuration is complete.

Related Tasks

- ▶ To use Multi-Link to balance outbound traffic, see Getting Started with Outbound Traffic Management (page 610)

Defining Routing for the Route-Based VPN

Prerequisites: Defining Tunnel Interfaces for Firewalls

In a Route-Based VPN, the routing defines which traffic is sent through the VPN tunnel. The routing configuration also determines the physical network interfaces on the engine to which the Tunnel Interfaces are automatically mapped.

Routing for the Route-Based VPNs does not use Router or NetLink elements. Instead, you add the remote networks that are reachable through the VPN tunnel directly to the Tunnel Interfaces as if they were directly connected networks. Routes to local networks that are directly connected to the Tunnel Interface are automatically added if the Tunnel Interface has an IP address with a netmask other than the /32 host netmask.

Example **A Tunnel Interface has the following IP address and netmask: 10.1.1.1/24. An automatic route to the 10.1.1.0/24 network is added for the Tunnel Interface.**

▼ To define a route for the Route-Based VPN

1. Select **Configuration→Routing** from the menu. The Routing view opens.
2. Select the firewall or Virtual Firewall for which you want to configure routing from the list above the Routing tree.
3. Add the Network element(s) that represent the networks that are reachable through the Tunnel Interface(s):
   - To add a route using a Network element that is already configured, drag and drop the element from the Resources panel directly to the Tunnel Interface in the Routing tree.
   - To create a new Network element, use the right-click menu for the Tunnel Interface in the routing tree.

What’s Next?

- ▶ If you are defining routing as part of a new Route-Based VPN configuration, continue as explained in the Configuration Overview (page 919).
- ▶ Otherwise, install the policy to transfer the changed configuration to the firewall.

Adding Routes for IPS Engines and Layer 2 Firewalls

Prerequisites: None

IPS engines and Layer 2 Firewalls do not route traffic. Interfaces on IPS engines and Layer 2 Firewalls that pick up traffic for inspection do not need any manual routing configuration. You may need to define a default route if the Management Center (Management Servers and Log Servers) and other system components are not located on a directly connected network. Other routes may be needed in addition to the default route if one or more system components are not directly connected and cannot be reached through the default gateway.

▼ To add a route for an IPS engine or Layer 2 Firewall

1. Select **Configuration→Routing** from the menu. The Routing view opens.
2. Select the correct IPS engine or Layer 2 Firewall element from the list above the Routing tree.
3. Expand the tree and browse to the interface from/to which traffic is to be routed. A Network element corresponding to the interface's IP address is shown.
4. Add a Router to the Network element.
   - If the Router element you want to add already exists, drag and drop it onto the Network in the Routing tree.
   - If the Router element does not exist, create a new Router element by right-clicking the Network element under the correct interface in the Routing tree and selecting **New→Router**.
5. Add the element that contains the IP addresses that are routed through this interface:
   - To add a default route, right-click the Router element you just added in the Routing tree and select **New→Any Network**. The (existing) default Any Network element is added below the Router in the Routing tree.
   - To add a route using some element that is already configured, drag and drop the element from the Resources panel.
   - To create a new Host or Network element, use the right-click menu for the Router element you just added in the Routing tree.

What's Next?

▶ Refresh the policy on the IPS engine or Layer 2 Firewall to transfer the changed configuration.

Related Tasks

▶ You may want to test your routing configuration with a few IP addresses. See Checking Routes (page 608).

Removing Routes

Prerequisites: Adding Routes for Firewalls/Adding Routes for IPS Engines and Layer 2 Firewalls

You can remove any element you have added to the routing tree yourself at any time. Elements that are added automatically based on the IP addressing of the interfaces cannot be removed while they are still in use. Use the Access rules to control access to/from such addresses (see Editing Access Rules (page 672)).

Automatically added elements corresponding to a previous configuration are not automatically removed when the IP address of an interface is changed to prevent any possible manual routing definitions from being deleted accidentally. Instead, the elements that belong to the old configuration are shown as invalid and you must remove the obsolete elements manually.

Removing elements from routing does not delete the elements from the system.

▼ To remove an element from the Routing tree

↳ Right-click the element in the Routing view and select **Remove**. The element is removed from routing.

Modifying Antispoofing

Prerequisites: Adding Routes for Firewalls

Antispoofing is an important security measure. Antispoofing is intended to prevent malicious attempts to use a legitimate internal IP address to gain access from lower-security networks to higher-security networks. Antispoofing determines which addresses are valid source addresses for the network(s) connected to each interface. If an interface receives a packet with a source address that is not a valid source address for the network(s) that are connected to that interface, the packet is considered to come from a spoofed IP address.

Firewalls, Master Engines, and Virtual Firewalls automatically determine antispoofing rules based on the routing configuration. In most cases, there is no need to change the antispoofing configuration in any way.

What's Next?

- ▶ If you want antispoofing to allow connections that would otherwise be dropped, see Deactivating Antispoofing for an IP Address/Interface Pair (page 607).
- ▶ If you want antispoofing to drop connections that it currently allows, see Activating Antispoofing for Routable IP Addresses (page 608).

Deactivating Antispoofing for an IP Address/Interface Pair

In rare cases you may need to modify the default antispoofing definitions to make exceptions to the antispoofing rules (for example, if you have defined policy routing manually). Manually modified entries are marked with a plus sign (+) (active entries) or a minus sign (–) (disabled entries) in the Antispoofing view to distinguish them from the automatically generated antispoofing entries.

Note – Errors in the routing configuration (in the Management Client or in the surrounding network) can cause legitimate packets to be incorrectly identified as coming from a spoofed IP address. Always check that the routing is configured correctly before modifying antispoofing. For example, routing loops generate log messages about spoofed packets, which cannot be removed by any antispoofing changes.

By default, the engine interprets the antispoofing tree by picking the most specific entry defined in the view (for example, a definition of a single IP address is picked over a definition of a whole network). In the default mode, if some IP address must be allowed access through two or more different interfaces, the definition for each interface must be at the same level of detail for the IP address in question.

Example If Interface A contains a Host element for 192.168.10.101 and Interface B contains a Network element for 192.168.10.0/24, connections from 192.168.10.101 are considered by default spoofed if they enter through Interface B, even though the address is included in the Network element. The antispoofing configuration must be modified to allow the address also from Interface B.

▼ To configure antispoofing

1. Select **Configuration→Antispoofing** from the menu. The Antispoofing view opens.
2. Select the correct Firewall element from the list above the Antispoofing tree.
3. Add the element that represents the IP addresses for which you want to create an exception to the Antispoofing tree in one of the following ways:
   - Drag and drop an existing element on top of the interface through which you want the element's IP address(es) to be available.
   - Create a new element by right-clicking the Network element under the correct interface and selecting **New**→(*type of element*).
4. (*Optional*) If you want to allow all connections from a network through a specific interface, right-click the network under the correct interface and select **Absolute**. All the IP addresses that belong to that network are now allowed for the interface, although more specific antispoofing definitions for some addresses in the network may be defined for other interfaces.

Caution – Never mark the Any Network element as Absolute. Disabling antispoofing in this way is a security risk. Resolve large-scale antispoofing conflicts with specific antispoofing definitions or by changing routing.

What's Next?

▶ Refresh the policies of all affected engines to transfer the changed configuration.

Activating Antispoofing for Routable IP Addresses

The Antispoofing view shows the allowed addresses for each interface. Only the first interface selected for an aggregated link is shown in the Antispoofing view. Manually modified entries are marked with a plus sign (+) (active entries) or a minus sign (–) (disabled entries) in the Antispoofing view to distinguish them from the automatically generated antispoofing entries.

There is rarely any need to change the automatically added entries. The preferred way of preventing routing for IP addresses is to change the Routing view instead.

Note – Disabling or removing elements in the Antispoofing view *prevents access*. The IP addresses a disabled or removed element represents are considered spoofed addresses.

▼ To remove an element from the Antispoofing tree

↪ Right-click the element and select **Remove**.

▼ To disable an element in the Antispoofing tree

↪ Right-click the element and select **Disable**. You can again re-enable the element in the same way when needed.

What's Next?

▶ Refresh the policies of all affected engines to transfer the changed configuration.

Checking Routes

Prerequisites: None

The Route Query tool allows you to check where packets with a certain IP address are routed. The tool helps you check that the routing is correct and to quickly find a particular branch in the Routing tree. Policy routing is not considered in this query. The route query uses the configuration that is currently stored on the Management Server (shown in the Management Client). You must refresh the policy of the affected engines after completing the configuration to transfer the changed routing information.

▼ To check routes with a Route Query

1. Select **Configuration→Routing** from the menu. The Routing view opens.
2. Select the correct engine element from the list above the Routing tree.
3. Right-click the engine element and select **Query Route**. The Query Route dialog opens.

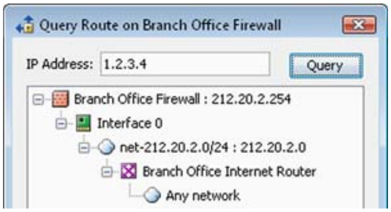

Query Route on Branch Office Firewall

IP Address: 1.2.3.4

Branch Office Firewall : 212.20.2.254
Interface 0
net-212.20.2.0/24 : 212.20.2.0
Branch Office Internet Router
Any network

4. Enter an **IP address** and click **Query**. The route is shown in the space below.

Related Tasks

▶ You can monitor routing as explained in Monitoring Connections, Blacklists, VPN SAs, Users, and Routing (page 111).

Chapter 37

Outbound Traffic Management

The Multi-Link feature of Single Firewalls, Firewall Clusters, and Master Engines allows you to distribute outbound traffic between two or more Internet connections to provide high availability and load balancing for outbound traffic.

The following sections are included:

Getting Started with Outbound Traffic Management

Prerequisites: See the Configuration Overview below

What Outbound Traffic Management Does

Single Firewalls and Firewall Clusters can balance outbound traffic between two or more network links (NetLinks) using the Multi-Link feature. NetLinks are combined into Multi-Link elements. The NetLinks may be of different types and they may have different speeds. You can also use Multi-Link with aggregated links.

Multi-Link allows you to:

- Balance outbound traffic between two or more alternative network links to increase the available bandwidth.
- Ensure that outbound network connectivity remains available even if network links fail. When a network link fails, the firewall detects this and stops forwarding traffic through the failed link.

Limitations

- Multi-Link is only supported on Single Firewalls and Firewall Clusters.
- VPN traffic is balanced independently from the settings covered in this configuration (when the firewall is the VPN end-point).
- If you configure IGMP proxy multicast routing and use a NetLink on a firewall’s upstream interface, you cannot use NetLinks on any of the firewall’s downstream interfaces.
- Multi-Link is only supported for IPv4 traffic.

What Do I Need To Know Before I Begin?

For a truly redundant configuration, we recommend connecting each link through different physical network equipment (such as routers).

Configuration Overview

Illustration 37.1 Elements in the Configuration

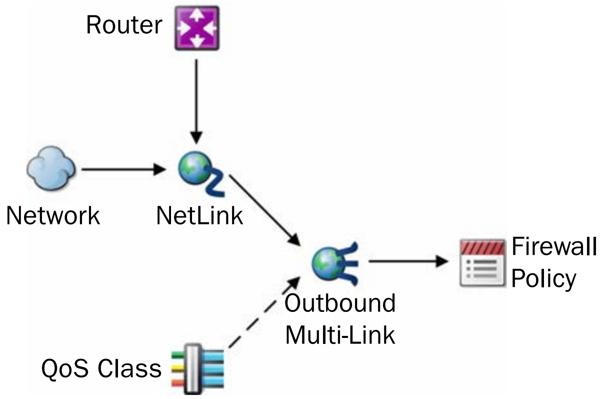

1. Configure routing with at least two NetLinks. See Defining a Multi-Link Route (page 600).
2. If you want the firewall to select the NetLink based on the type of traffic, create QoS Classes and assign them to traffic. See Getting Started with QoS (page 782).
3. Create an Outbound Multi-Link element to group your NetLinks and define traffic management settings. See Configuring Outbound Multi-Link Settings.
4. In the Firewall Policy, create a NAT rule in for outbound load balancing. See To define NAT rules for outbound load balancing (page 614).

Configuring Outbound Multi-Link Settings

Prerequisites: Defining a Multi-Link Route, (*optional*) Creating QoS Classes

Outbound Multi-Link elements are used for combining NetLinks and for setting options for the high availability and load balancing features.

There are two methods for link selection. If link selection is based on *Round Trip Time,* engines send a connection to all the links, and then select the link that delivers the response the fastest. The link selection can also be based on a *Ratio* of the relative bandwidth of each NetLink. You must configure the basic settings for these methods in the properties of each NetLink (see Creating NetLinks (page 601)).

What's Next?

▶ Combine your NetLinks into Outbound Multi-Link element(s). Proceed to Creating an Outbound Multi-Link Element (page 612).

Creating an Outbound Multi-Link Element

The Outbound Multi-Link element collects together the NetLinks and sets options for the load balancing. The Outbound Multi-Link elements you create do not work on their own; you must use them in the Firewall Policy's NAT rules to select traffic for outbound load balancing.

▼ To define an Outbound Multi-Link element

1. Select **Configuration→Configuration→Security Engine**. The Security Engine Configuration view opens.
2. Browse to **Network Elements→Traffic Handlers**.
3. Right-click **Traffic Handlers** and select **New→Outbound Multi-Link**. The Outbound Multi-Link Properties dialog opens.

Illustration 37.2 Outbound Multi-Link Properties

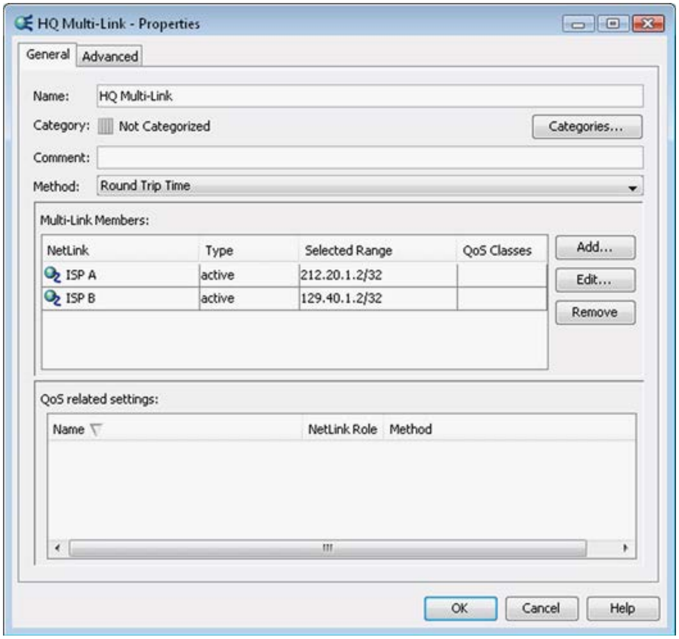

Name: HQ Multi-Link

Category: Not Categorized

Comment:

Method: Round Trip Time

Multi-Link Members:

NetLink	Type	Selected Range	QoS Classes
ISP A	active	212.20.1.2/32	
ISP B	active	129.40.1.2/32	

QoS related settings:

Name	NetLink Role	Method

4. Give the Outbound Multi-Link element a unique **Name**.
5. (*Optional*) Add a free-form **Comment** for your own reference.
6. Specify the **Method** for link selection:
   - **Round Trip Time**: The firewall periodically probes the NetLinks to test them for speed and selects the fastest available active NetLink for each new outbound connection.
   - **Ratio**: Traffic is distributed between all of the available active NetLinks according to the relative bandwidth of each NetLink. The NetLink with the highest bandwidth is assigned the largest portion of the traffic.

What's Next?

▶ Continue by Selecting NetLinks for an Outbound Multi-Link (page 613).

Selecting NetLinks for an Outbound Multi-Link

▼ To add a NetLink to an Outbound Multi-Link element

1. Click the **Add** button. The Multi-Link Member dialog opens.

Illustration 37.3 Multi-Link Member Dialog

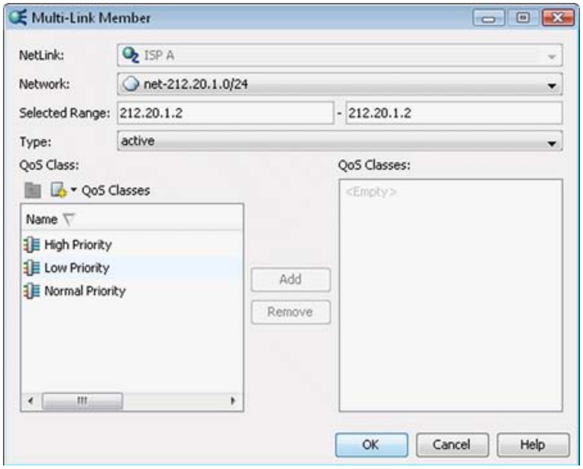

Multi-Link Member

NetLink: ISP A

Network: net-212.20.1.0/24

Selected Range: 212.20.1.2 - 212.20.1.2

Type: active

QoS Class:

Name
High Priority
Low Priority
Normal Priority

QoS Classes: <Empty>

2. Select a **NetLink**.

3. Select the **Network** element that represents the IP address space in the directly connected external network of this network link.

4. In the **Selected Range** fields, specify the IP address range for dynamic source address translation (NAT) for the internal source IP addresses on this NetLink. To define a single IP address, enter the same address in both fields.
   - Since Multi-Link works without external arrangements, the address translation is necessary for correct routing of packets.

5. Select the **Type**:
   - **Active**: traffic is routed through the NetLink according to the Method you specify in the Outbound Multi-Link element properties.
   - **Standby**: traffic is only routed through the NetLink if all primary (active) NetLinks are unavailable.

6. (*Optional*) Select the QoS Classes for traffic handled by this NetLink and click **Add**. You can use the QoS classes to assign the NetLink with or without activating the actual QoS features. For more information, see Getting Started with QoS (page 782).
   - You can select the same QoS class for several NetLinks to balance the traffic between the NetLinks. If none of the NetLinks with the appropriate QoS class are available, or if the traffic has not been assigned a QoS class, the traffic is distributed between the NetLinks according to the Method you specify in the Outbound Multi-Link element properties.
   - QoS classes are assigned based on ToS codes in network traffic or in the Access rules. Traffic that has been assigned the selected QoS class uses this NetLink if the NetLink is available.

7. Click **OK**. The Multi-Link Member dialog closes, and the NetLink is listed in the Multi-Link Members list.

8. Repeat these steps for each NetLink you want to add to the Outbound Multi-Link element. When you are finished, click **OK** in the Multi-Link Properties dialog.

What's Next?

- ▶ If you want to define how information on NetLink performance is cached, continue by Defining Destination Cache Settings.
- ▶ Otherwise, proceed to Creating Outbound Load Balancing NAT Rules (page 614).

Related Tasks

- ▶ Changing NetLink State Manually (page 224)

Defining Destination Cache Settings

Information about the performance of each NetLink is cached. No new measurement is made if a new connection is opened to the same destination within a short time period. You can define the duration of the cached information.

▼ To define destination cache settings

1. Switch to the **Advanced** tab in the Outbound Multi-Link Properties.
2. Deselect **Default**.
3. Enter the **Timeout** (in seconds) after which a new measurement of NetLink performance is made. The default is 3600 seconds.
4. Enter the number of **Maximum Retries** for checking each NetLink. The default is 2.
5. Click **OK**.

What's Next?

- ▶ Continue by Creating Outbound Load Balancing NAT Rules (page 614).

Creating Outbound Load Balancing NAT Rules

Prerequisites: Defining a Multi-Link Route, Configuring Outbound Multi-Link Settings

To balance outbound connections between the NetLinks, you must define NAT rules using the source NAT addressing defined in the Outbound Multi-Link element.

You can create several Outbound Multi-Link elements to balance different types of traffic differently. If there is some traffic that you do not want to balance at all, you can direct traffic through a specific NetLink with a normal NAT rule that translates the source using a particular NetLink's address space and matches the desired connections before the balancing rule.

▼ To define NAT rules for outbound load balancing

1. Right-click the policy and select **Edit**. The policy opens for editing.
2. Switch to the **IPv4 NAT** tab.
3. Add a new rule and specify the **Source**, **Destination**, and **Service** according to the traffic that you want to balance using a particular Outbound Multi-Link element.
4. Double-click the **NAT** cell. The Network Address Translation dialog opens on the Source Translation tab.

5. Select **Dynamic** as the **Translation Type**.
6. Click the **Select** button next to the IP Address Pool field. The Select Element dialog opens.
7. Browse to **Network Elements→Traffic Handlers** and select the Outbound Multi-Link element.
8. (*Optional*) Specify a port range for source port translation by entering port numbers in the **First Port to Use** and **Last Port to Use** fields. The default is to use all “free” high ports from 1024 to 65535.

> **Note – Make sure that the IP address pool and port range are large enough to translate the number of simultaneous connections that match this NAT rule.**

9. Select/deselect **Automatic Proxy ARP**:
   - Make sure that this option is selected if you want proxy ARP entries to be generated.
   - For proxy ARP to work, the untranslated IP address of all hosts whose IP address is translated must be included in the routing configuration in the Routing view (that is, there is a CVI with an address from the same network defined in the Firewall element’s properties or the Network or Host is manually added to the Routing view).
   - Deselect the option to disable proxy ARP.
10. Click **OK**.

Illustration 37.4 Outbound Load Balancing NAT Rule - Example

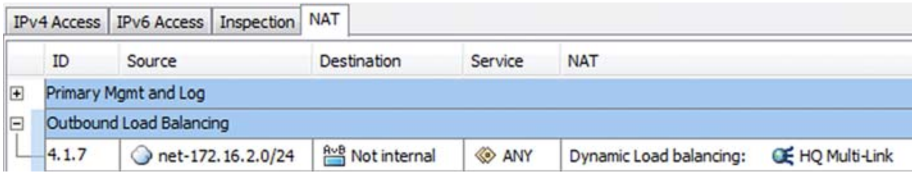

ID	Source	Destination	Service	NAT
Primary Mgmt and Log				
Outbound Load Balancing				
4.1.7	net-172.16.2.0/24	Not internal	ANY	Dynamic Load balancing: HQ Multi-Link

This completes the configuration outlined in the Configuration Overview (page 611). Save and install the policy to transfer the changes to the engines.

What’s Next?

- To follow up on how your Multi-Link configuration works, see Monitoring And Testing Outbound Traffic Management (page 616).
- To configure inbound traffic management, see Getting Started with Inbound Traffic Management (page 618).
- VPNs have independent link selection. To configure active and standby links (tunnels) for VPNs, see Defining VPN Tunnel Settings for Policy-Based VPNs (page 953).

Monitoring And Testing Outbound Traffic Management

Prerequisites: See Getting Started with Outbound Traffic Management

The status of the different NetLinks can be monitored in the System Status view. See Getting Started with System Monitoring (page 94). Status is displayed based on the results of status probes to the Probing addresses configured in the NetLink elements.

You can test that the outbound traffic management configuration works as intended by generating traffic and then disabling network connections, for example, by unplugging cables (in a cluster, disable the same link on all active nodes to simulate ISP failure).

Note – By default, if you unplug a cable from a node in a cluster, the automatic engine tester recognizes this as a malfunction and turns the node offline (unless it is the last node that remains online). To override this automatic test, command all nodes to “Lock Online” for the duration of the test.

Most connections that are currently open on the link that goes down are cut and must be reopened by the communicating applications, since the external IP address used in the communications changes to an address of the new link. Any type of traffic can switch to using a different link completely transparently if it is transported through a Multi-Link VPN between two Firewall/VPN engines.

If the Ratio-based load-balancing method is used, a link is used until a failure is detected based on the link status probes to the Probing addresses configured in the NetLink elements. With the Round Trip Time method, new connections are directed to working links (since the link that responds the fastest is chosen) even if probing is not configured or working.

Related Tasks

- Getting Started with the Engine Tester (page 520).

CHAPTER 38

INBOUND TRAFFIC MANAGEMENT

Firewalls allow you to set up a server pool that distributes the load of incoming traffic between a group of servers and/or allows controlling the NetLink use of incoming traffic in a Multi-Link setup. This ensures that services remain available even when one or more servers or NetLinks fail, and balances the load of incoming traffic more efficiently between a group of servers. Inbound traffic management is not supported on Layer 2 Firewalls.

The following sections are included:

Getting Started with Inbound Traffic Management

Prerequisites: None

Inbound traffic on a Firewall can be managed with a server pool.

What Inbound Traffic Management Does

The first and primary use for a server pool is to offer load-balancing and high availability for two or more servers that offer the same service to users. A second use is to control which NetLink incoming traffic uses. Inbound traffic management can:

- load-balance incoming traffic between several servers to even out their workload
- monitor the server's status so that the traffic is not directed at unavailable or overloaded servers (if the optional Monitoring Agents are installed on each server)
- send dynamic DNS (DDNS) updates to a DNS server to prevent incoming traffic from attempting to use a non-functioning NetLink in a Multi-Link setup (with or without using the other features of Server Pools).

Limitations

- DDNS updates have no access control, so the communications must be secured in other ways.
- Standby servers cannot be defined for a server pool. Only load-balancing between the servers in the server pool is supported.
- Only TCP and UDP protocols are supported for server pools.
- Source address translation is not supported for server pools.
- Server Pools are only supported for IPv4 traffic.

What do I need to know before I begin?

- Each server pool should only offer one type of service. If the same servers provide more than one service (for example, HTTP and HTTPS), create separate Server Pools for each service.

Configuration Overview

Illustration 38.1 Elements in the Configuration

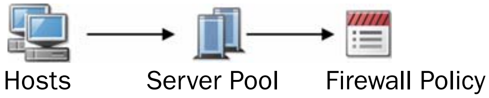

1. Create Host elements for all servers you want to include in the server pool.
2. Define a Server Pool element and define settings for inbound traffic management (see Defining a Server Pool).
3. Check that the servers' public, external IP addresses in the Server Pool properties correspond to their in the DNS server (see Entering Server Pool IP Addresses on Your DNS Server (page 637)).
4. Modify your policy to include inbound traffic management (see Creating Access Rules for Inbound Load Balancing (page 638)).
5. (*Optional*) If you want to use Server Pool Monitoring Agents, install, configure, and enable them (see Installing Monitoring Agents (page 621)).
6. (*Optional*) If you want to send dynamic DNS (DDNS) updates to a DNS server, configure the updates (see Configuring Dynamic DNS Updates (page 639)).

Defining a Server Pool

Prerequisites: None

The Server Pool element collects servers that provide a particular service into a single element and defines the settings for handling the inbound traffic.

Creating a New Server Pool Element

Add servers in Server Pools based on the services that the servers provide. The Server Pool considers all included servers to be equal. If the same servers provide more than one service, create separate Server Pools for the different services even if the same services run all servers (each Server Pool reserves an additional IP address).

▼ To create a new Server Pool element

1. Select **Configuration→Configuration→Security Engine**. The Security Engine Configuration view opens.
2. Expand the **Network Elements** branch.
3. Right-click **Traffic Handlers** and select **New→Server Pool**. The Server Pool Properties dialog opens.
4. Enter a unique **Name** for the Server Pool.

Defining External Address(es) of Server Pool

Clients make their incoming connections to the address of the Server Pool. The Firewall then decides which server is going to handle the connection and translates (in a NAT operation) the public address to the private IP address of that server. The external address or addresses of the Server Pool are defined as properties of the Server Pool element.

Note – Make sure your NAT rules do not overlap with the NAT of the Server Pool.

▼ To add External Addresses

1. Click **Add** in the External Addresses section of the Server Pool Properties dialog. The External Address dialog opens.

Illustration 38.2 External Address

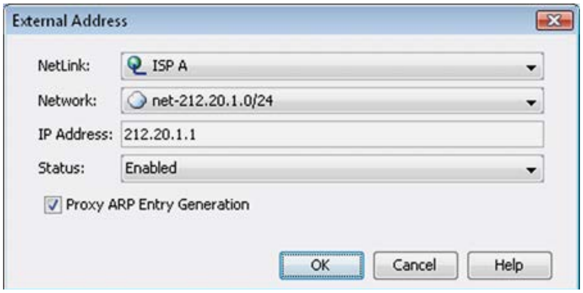

External Address

NetLink: ISP A

Network: net-212.20.1.0/24

IP Address: 212.20.1.1

Status: Enabled

☑ Proxy ARP Entry Generation

2. Select the **NetLink** you want to use, or if you want to configure load sharing for the servers but no traffic balancing between NetLinks, select **Not Specified**.
3. For the **Network**, select the Network element that is used for the Server Pool's external NATed address.

4. Define the external NATed destination **IP Address** for the Server Pool. This is the address client machines contact when accessing the service that the server(s) in the pool offer.

Note – The IP address you enter here must be reserved for NAT and it must not be used by any equipment in your network. Remember to update your DNS server with any changes in IP addressing.

5. In **Status**, select **Enabled** to use the NetLink for the Server Pool.
6. (*Recommended*) Select **Proxy ARP Entry Generation** to automatically generate a proxy ARP for the NATed address in the selected Network. Otherwise, you must define the ARP entry manually in the Firewall element properties.
7. Click **OK** to add the external address. Repeat these steps for any other NetLinks that handle traffic to this Server Pool.

Adding Server Pool Members

You must add a Host element for the internal IP address of each member of the pool. The Firewall uses these addresses to select which server handles which traffic that arrives to the Server Pool's external address.

The Server Pool can have any number of members. You can also create a one-server Server Pool to enable DDNS updates for a single server when ISP links go down if the server does not need the other Server Pool features.

▼ To add Server Pool members

1. Click **Select** in the Server Pool Members section of the Server Pool Properties dialog. The Select Element dialog opens.
2. Select the elements for the members of the pool.
   - You can use either Server elements (for servers that have some special role in the SMC configuration) or Host elements (for other servers).
   - To create new elements to represent the Server Pool members, browse to Network Elements and click the New icon above the element list.
3. In **Allocate Traffic to Servers by**, select the granularity for the server selection (defines how 'sticky' the selection is). Usually it is best to choose the least granular option that still produces an acceptable distribution of traffic. The options are (from least granular to most granular):
   - **Source Network** directs traffic coming from the same C-class network to the same server. This is a good choice when connections come from many different networks. The **Not Defined** setting has the same effect.
   - **Host** directs traffic coming from the same IP address to the same server. This is a good choice when a large portion of connections come from different hosts in the same C-class network.
   - **Connection** makes a new traffic management decision for each new connection. This choice may be necessary if a large portion of connections uses just one IP address.

Note – Consider the type of traffic when selecting the allocation method. Using the Host setting (if the host's IP address apparent to the Management Center can change) or the Connection setting (in all cases) means that connections from the same source may be directed to different servers. Depending on the services offered, this may deteriorate service quality.

4. Define the **Monitoring** mode for the availability of the servers in the Server Pool.
   - Select **Ping** if you want the Firewall to monitor the availability of the servers by using periodical ICMP echo request (ping) messages.
   - Select **Agent** if you want to use the Server Pool Monitoring Agent feature. The default port for the servers to listen on is UDP 7777, but you can select a different port in the **Port** field.
5. Click **OK**.

What's Next?

- If you are using Monitoring Agents, continue to Installing Monitoring Agents.
- Otherwise, continue to Entering Server Pool IP Addresses on Your DNS Server (page 637).

Installing Monitoring Agents

Prerequisites: Defining a Server Pool

Server Pool Monitoring Agents are installed on the servers that form the pool to report about the server's health and load. The Firewall uses the information to send traffic to the server that can best handle the additional load and not send new connections to a server that is inoperative. You must install the Server Pool Monitoring Agents onto all members of the Server Pool.

To install in Windows

1. Copy the `ServerPool_Monitoring_Agent_Installer\Windows\` directory from the Stonesoft Management Center installation DVD to the Server Pool member on which you wish to have the Monitoring Agent installed.
2. Run the self extracting installation program `sgagent.exe`.
3. Repeat these steps on all Server Pool members.

What's Next?

- Proceed to Configuring Monitoring Agents (page 623).

To install in Linux

1. Copy the `ServerPool_Monitoring_Agent_Installer/Linux/` directory from the Stonesoft Management Center installation DVD to the Server Pool member on which you wish to have the Monitoring Agent installed.
2. Go to the location where the directory was copied, and extract the `SGFWagent` package. Run the `rpm` command. For example, you can use the following commands. The name of the `.rpm` file may vary depending on the version you are using.

```
cp SGFWagent.tar /var/spool/pkg
cd /var/spool/pkg
tar -xf SGFWagent.tar
rpm -i sgagent-4.0.0.build403-1.i386.rpm
```

3. Repeat these steps on all Server Pool members.

What's Next?

▶ Proceed to Configuring Monitoring Agents (page 623).

▼ To install in Solaris

1. Copy the `ServerPool_Monitoring_Agent_Installer/Solaris/` directory from the Stonesoft Management Center installation DVD to a Server Pool member on which you wish to have the Monitoring Agent installed.
2. Go to the location where the directory was copied. Untar the `SGFWagent` package and run the package add command. For example, you may use the following commands:

```
cp SGFWagent.tar /var/spool/pkg
cd /var/spool/pkg
tar -xf SGFWagent.tar
pkgadd
```

3. Select `SGFWagent` for installation.
4. Repeat these steps on all Server Pool members.

What's Next?

▶ Proceed to Configuring Monitoring Agents (page 623).

Uninstalling Monitoring Agents

Prerequisites: If you plan to remove all Monitoring Agents, reconfigure the Server Pool not to use them

▼ To remove the Monitoring Agent in Windows

1. Open the Windows Control Panel.
2. Click **Add/Remove Programs**.
3. Select **Stonesoft Monitoring Agent** and follow the instructions.

▼ To remove the Monitoring Agent in Linux

➥ Run the following command:

- **`rpm -e sgagent`**

▼ To remove the Monitoring Agent in Solaris

➥ Run the following command:

- **`pkgrm SGFWagent`**

Configuring Monitoring Agents

Prerequisites: Installing Monitoring Agents

The Monitoring Agents are configured in two files: `sgagent.local.conf` and `sgagent.conf`. Both files are found in the following directory on each server:

- Solaris: `/opt/sgagent/etc/`
- Linux: `/etc/stonegate/`
- Windows: `Program Files\Stonesoft\StoneGate Monitoring Agent\`

The `sgagent.local.conf` file contains server-specific information, and is different for each server. If you do not want to change the default values, there is no need to edit this file.

The `sgagent.conf` file contains information that refers to the entire Server Pool.

What's Next?

▶ To define settings specific to individual servers, proceed to Editing sgagent.local.conf.

▶ If you only want to configure global settings for all servers, proceed to Editing sgagent.conf (page 624).

Editing sgagent.local.conf

The server-specific file `sgagent.local.conf` includes information to identify the Server Pool member if default values are not used. If the file is not configured, the Monitoring Agent uses the system host name as the default value to identify the member. If edited, this file must be different on each Server Pool member.

Table 38.1 Statements in sgagent.local.conf

Statement	Description
host *hostname*	Specifies *hostname* of the server. If no value is defined, the system host name is used.
set *alias value*	Defines the *value* of an *alias* that can be used as a variable in the global sgagent.conf configuration.

For example, to configure the member of the pool where `sgagent.local.conf` resides to use *server1* as its host name, and to set the alias *hostip* to 192.168.1.1:

```
host server1
set hostip 192.168.1.1
```

What's Next?

▶ Proceed to Editing sgagent.conf (page 624).

Editing sgagent.conf

The file common to all Server Pool members is called sgagent.conf. In this file, you define attributes for the entire Server Pool. In most cases, the Server Pool members are similar in configuration and function, so the tests performed on each member are also similar. Usually, the same sgagent.conf file is used on all servers. The sgagent.conf file contains two sections: the statement section and the test section. See Illustration 38.3.

Illustration 38.3 An Example sgagent.conf File

```
# --------------------------START OF STATEMENT SECTION
begin host server1
  config boot-delay 5:00  # set the config delay on the slow member
end
begin host not server1
  config boot-delay  20   # 20s config delay on other members
end
config alert-interval 1:00:00
config startup-script /etc/scripts/startup.sh
# Monitor one minute load average and set alert threshold at 5
load-index  500  load-average-1
# --------------------------START OF TEST SECTION
# Tests
test "webster listening"
  interval 2:00
  action  exclude
  recovery  10:00
command portlistening 80 %hostip%  # Check that this member's
                                   # port 80 is listening.
test "server1 running baz"
  interval 10:00
  action exclude
  host server1
  script /etc/init.d/baz start
command servicerunning baz
test "external1"
  interval 10:00
  action alert
command external 3 700 /usr/lib/webster/test.sh 212.20.2.21
```

What's Next?

- ▶ Proceed to Editing the sgagent.conf Statement Section (page 625).

Editing the sgagent.conf Statement Section

In the statement section, you can define general commands for the tests the Monitoring Agents perform. A statement can apply to all or only some members of the Server Pool.

Table 38.2 Statements in the sgagent.conf Statement Section

Statement	Description

begin host *hostname*	If you wish a statement to apply to certain Server Pool members only, the statement must start with the `begin host` command followed by the name of the Server Pool member.
	The `begin host` command may be followed by actual server names or an exclusion. In an exclusion, `begin host` is followed by `not` and the server name, instead of just a server name. This indicates that the setting is to be applied to all other servers except the one specified.
	For example, the following statement would apply to all members in the pool except server1:
	`begin host not server1`
	`config settings`
	`end`

config [**port** *port_number*] [**boot-delay** *time*] [**alert-interval** *time*] [**alert-script** *path*] [**startup-script** *path*] [**overload-script** *path*] [**listen** *IP_address* [*:port*]] [**load-index** *threshold* *index_name* [***index-parameter***]] [load-index-action exclude *time*]	Defines the option in question. See Options in the sgagent.conf Statement Section (page 626) for more information about the options. If a statement applies to all members in the pool, you can enter the setting on a single line without using begin or end.
end	Ends the statement when `begin` is used.

What's Next?

- Proceed to Options in the sgagent.conf Statement Section (page 626).

Options in the sgagent.conf Statement Section

When you want to configure a particular setting, the syntax starts with the command `config`. The available options are listed below.

Table 38.3 Options in the sgagent.conf Statement Section

Option	Description
`port` *`port_number`*	Specifies the port that the Server Pool uses for communication with the engines. All Server Pool members must use the same port number; the port setting cannot start with the `begin host` command. The port number can be from 0-65535. By default, the port is set to `7777`. If you use another port instead of the default port, you must modify the IPv4 and IPv6 Access rules to allow connections to the new port.
`boot-delay` *`time`*	Defines how long the tester waits before starting the tests after the system has been booted up or restarted to ensure that all necessary processes have completed their startup. Time is entered in the format [(hh:)mm:]ss. The default is 30 seconds.
`alert-interval` *`time`*	Defines how long the system waits before sending a new alert in response to a test failure to prevent excessive alerts from being sent. Time is entered in the format [(hh:)mm:]ss. The default is 1 hour.
`alert-script` *`path`*	Specifies the path to a custom alert script. You must use quotation marks if the path contains spaces. This is only needed when you want to use a custom alert script. There is no default value.
`startup-script` *`path`*	Specifies the path to the custom script to be run when the Monitoring Agent starts. You must use quotation marks if the path contains spaces.
`overload-script` *`path`*	Specifies the path to the custom script to be run if the load index exceeds the threshold value defined with the `load-index` statement. You must use quotation marks if the path contains spaces.
`listen` *`IP_address`* `[`*`:port`*`]`	Specifies the IP address the Monitoring Agent listens on. If no IP address is specified, the first IP address of the interface is used by default. If no port is defined, the default port (7777) is used.
`load-index` *`threshold index_name`* `[`*`index-parameter`*`]`	Defines the method that is used to measure the load. We recommend using the same method of measurement for all members of the Server Pool. If the load measurement exceeds the defined *`threshold`*, a log message and an alert are generated. The *`index_name`* can be a complex expression combining multiple methods with basic arithmetic operations (+ - * /), such as “index-name + index-name”. The index_name options are explained in Table 38.4.
`load-index-action exclude[`*`time`*`]`	Excludes the server from handling traffic for the specified *`time`* (in seconds) when the threshold value of the load-index is reached. When the server has switched to the Excluded state, the server does not return to the normal state until the specified length of time has elapsed, even if the load of the server has dropped to a normal level.

Table 38.4 Index_name Options

Option	Description
processor-load *time*	Measures the average processor load level over the specified time period. The time is entered in the format [(hh:)mm:]ss. The returned value is a percentage value from 0-100.
processor-kernel-load *time*	Measures the average processor time spent in the kernel mode over the specified time period. The time is entered in the format [(hh:)mm:]ss. The returned value is a percentage value from 0-100.
load-average-*minutes*	Measures the average number of processes waiting in the execution queue due to reserved resources over the specified number of minutes. The options for the value of *minutes* are 1, 5, or 15. The returned value is the load average for the specified time period multiplied by 100.

What's Next?

- To see some examples that illustrate the statement configuration explained so far, proceed to Monitoring Agent Statement Configuration Examples.
- To skip the examples and see how to configure the Test section, proceed to Editing the sgagent.conf Test Section (page 629).

Monitoring Agent Statement Configuration Examples

Port Option Example

To use port 5555 for the communication between the engine(s) and the Server Pool members instead of the default port:

```
config port 5555
```

Boot-delay Option Example

To set a boot delay of 20 seconds on server1:

```
begin host server1
config boot-delay 20
end
```

Alert-interval Option Example

To set an alert interval of 1 hour on servers 1 and 2:

```
begin host server1 server2
config alert-interval 1:00:00
end
```

Alert-script Option Example

To run `alertscript.sh` every time an alert is generated:

```
config alert-script /etc/test/alertscript.sh
```

Startup-script Option Example

To run `startup.sh` when the Monitoring Agent starts on server1:

```
begin host server1
config startup-script /etc/test/startup.sh
end
```

Overload-script Option Example

To run `overload.sh` script when the threshold value of the load index is reached on server2:

```
begin host server2
config overload-script /etc/test/overload.sh
end
```

Listen Option Example

To make the Monitoring Agent of server2 listen to the IP address 192.168.10.31 using the default port 7777:

```
begin host server2
config listen 192.168.10.31
end
```

Load-index Option Example

This example demonstrates how the measurement of `load-average-5` is used on all servers except server1 to compare against the *`threshold`* value 4. When the load-average is higher than 4, a log message is generated.

```
begin host not server1
config load-index 400 load-average-5
end
```

Because server1 is ignored in the previous statement, it would typically have its own statement with different options, such as the one below:

```
begin host server1
config load-index 400 load-average-1
end
```

Load-Index-Action Option Example

To set the server status to Excluded when the threshold value of the load index is reached on server2:

```
begin host server2
config load-index-action exclude 10
end
```

What's Next?

- ▶ Proceed to Editing the sgagent.conf Test Section.

Editing the sgagent.conf Test Section

The second part of the `sgagent.conf` file specifies which tests the Monitoring Agents carry out on the servers they monitor. A test begins with the `test` definition, followed by optional parameter lines, and ends with the `command` definition. The `test` and `command` parameters are always required. The `interval`, `action`, `host`, `script`, and `recovery` parameters are optional and have their default values if not specified for the test.

Table 38.5 Parameters in the sgagent.conf Test Section

Parameter	Description
test *name*	Begins the test definition. Specifies the name of the test. The *name* appears in the logs and alerts in the event of a failure, so it should be descriptive to give the administrators enough information about what has happened. The test name may be any string. You must use quotation marks if the name contains spaces. This parameter is required and has no default value.
interval *time*	Defines how often the test is executed. Time is entered in the format [(hh:)mm:]ss. The minimum interval is 3 seconds and the maximum is 1 day. If no interval is defined, the default is 60 seconds.
action `alert|exclude`	Defines what to do if the test fails. `alert` creates an alert entry and runs the alert script if one is defined for the test. `exclude` sets the server status to Excluded and creates an alert. No new connections are directed to the excluded server and the current connections are moved to the other members of the pool. If no action is set for the test, `alert` is the default action.
host *server*	Defines the *server* to be tested. If the name of the server is specified in the host field of sgagent.local.conf, that name must be used here. Otherwise, the default host name of the server must be used. See Editing sgagent.local.conf (page 623) for more information. You can also list multiple server names, separated by spaces. If no specific server is defined, the test is performed on all Server Pool members by default.

Table 38.5 Parameters in the sgagent.conf Test Section (Continued)

Parameter	Description
script [always|*repeat*] *path*	Specifies how many times the script is executed after the test fails, and the *path* to the script file. You must use quotation marks if the path contains spaces. There are two options for how many times to execute the script after the test fails: *repeat* defines the number of failures after which the script is no longer run. The counter is reset each time the server status changes, or the test succeeds. always runs the script without restrictions on the number of failures. If no argument is given, the script is run only after the first failure by default.
recovery always|*time*	Determines how long the tester reruns a failed test to see if it succeeds on the next attempt. There are two options for specifying the time period: always: there are no time limits on how long the tester waits for the test to succeed again. The server state is changed from Excluded to OK when the test succeeds. *time*: If the test succeeds in the defined time period after a failure, the server state is changed from Excluded to OK. Time is entered in the format [(hh:)mm:]ss. If no recovery options are specified, the test is never rerun after a failure, and the server state remains Excluded.
command external|*test_expression* [*retry*] [*timeout*] [*path*] [*parameters*]	Specifies which test script to use, and the options for running the test. This parameter is required and has no default value. There are two options for specifying the test script: external: indicates that an external test is used. External tests can be created by combining predefined tests with AND, OR, and NOT operators, or you can write your own script. Note that external tests must return an exit code of 0 (zero) to indicate success. Any non-zero return value is interpreted as a failure. *test_expression*: specifies the internal test to run. For more information about the internal tests, see Editing Internal Tests for Monitoring Agents (page 632). The other options are used as follows: *retry*: specifies how many times the script is run before the test is considered failed. *timeout*: specifies how long (in milliseconds) the tester waits for the test result. *path*: specifies the path to the external test script. You must use quotation marks if the path contains spaces. The path is required for external tests, and is not used for internal tests. *parameters*: specifies any possible *parameters* that a specific test may need.

What’s Next?

- To see some examples that illustrate the test configuration explained so far, proceed to Monitoring Agent Test Configuration Examples.
- To skip the examples and see which internal tests are available, proceed to Editing Internal Tests for Monitoring Agents (page 632).

Related Tasks

- Enabling Monitoring Agents (page 637)

Monitoring Agent Test Configuration Examples

Internal Test Example

In this example, the test checks that server1 is listening to traffic on port 80. In this example, the port information and the IP address of the server are required in the command definition. The test is performed every two minutes, and in case of a failure, the server is excluded from the Server Pool. If the test fails, the `/etc/init.d/port start` script is run. If the subsequent tests succeed within 10 minutes of the failure, recovery occurs and the server returns to the OK state.

```
test “port listening test”
interval 2:00
action exclude
host server1
script /etc/init.d/port start
recovery 10:00
command portlistening 80  192.168.1.1
```

External Test Example

In this example, the external test script `test.sh` is run. There are three attempts to run the test before it is considered failed, and the tester will wait for one second before the test is timed out. The test is run every 10 minutes, and in case of a failure an alert is generated.

```
test “multiple_services”
interval 10:00
action alert
command external 3 1000 /usr/lib/server/test.sh 192.168.1.1
```

What’s Next?

- Proceed to Editing Internal Tests for Monitoring Agents (page 632).

Related Tasks

- Enabling Monitoring Agents (page 637)

Editing Internal Tests for Monitoring Agents

All internal tests available for the Server Pool monitoring are listed below. The syntax for using the tests in the command definition is described in Editing the sgagent.conf Test Section (page 629).

Note – Some of the internal tests may not work on virtualization platforms.

Table 38.6 Internal Tests

Test	Description
swapfree *limit*	Checks the available swap space. If the current amount of free swap space drops below the *limit* (given in kilobytes), the test fails.
processcount *count*	Checks how many processes are currently running. If the number of processes exceeds *count*, the test fails.
multi-ping *retry* *timeout* *ip_address* [*ip_address* [...]]	Sends ICMP Echo Request messages to the defined IP addresses and waits for the replies. The test fails if none of the IP addresses answers the ping query. *retry* specifies how many ICMP Echo Request packets are sent to each IP address. *timeout* defines how long (in milliseconds) the tester waits for the reply.
filesystem *partition* *free_space*	Checks whether there is enough space left on a particular *partition*. The test fails if the amount of *free_space* (in kilobytes) is lower than the free space on the partition. The partition may also be identified by its path. File system tests are quite taxing on the operating system. We recommend that you set an interval of at least 15 minutes when using this test.
servicerunning *service_name*	Checks whether the specified service is running on the server. The test fails if the specified service/process is not running. On Windows platforms, the test checks whether the service called *service_name* is running. In UNIX environments, the test checks whether the process named *service_name* is running.
ip-listening *ip_address*	Checks whether the specified *ip_address* is listening on the server. The test fails if the host does not have the specified IP address.
portlistening *port* [*ip_address* [*protocol*]]	Checks whether the specified *port* is listening on the server(s). The test fails if the port is not in the listen state. You can optionally specify the *ip_address* at which to check for the port. The *protocol* can be TCP or UDP.

Table 38.6 Internal Tests (Continued)

Test	Description
portanswer *retry* *timeout* *query* *response* *port* [*ip_address*]	Checks whether the TCP port answers a TCP query with the expected response. The test fails if the port does not answer, or the answer differs from the expected response. Note that the Telnet protocol executes handshake procedures before beginning the connection. If you use `portanswer` to test a Telnet service, test the handshake procedure, not the Telnet prompt. *retry* specifies how many attempts are made before the test is considered failed. *timeout* specifies how long the tester waits after sending a query to the port. The timeout is entered in milliseconds. If no response is received within the time period, the test fails. *query* specifies the TCP query to send. You must use quotation marks if the query contains spaces. If no query parameter is specified, nothing is sent to the port. *response* specifies the expected response to the TCP query. You must use quotation marks if the response contains spaces. *port* specifies the port to which the query is sent. *ip_address* is the IP address to which the query is sent.
httpanswer *retry* *timeout* *URL* *file* *port* [*ip_address*]	Checks whether the server answers an HTTP request for the specified *URL* with the expected *file*. The test fails if the server does not answer the HTTP request, or if the reply does not contain the contents of the specified file. *port* specifies the port to which the request is sent. *retry* specifies how many attempts are made before the test is considered failed. *timeout* specifies how long the tester waits after sending a query to the port. The timeout is entered in milliseconds. If no response is received within the time period, the test fails. *ip_address* is the IP address to which the request is sent. If no *ip_address* is specified, the tester sends the query to the localhost address by default.
networkinterface-up *interface*	Checks that the specified *interface* is up. The test fails if the specified network interface does not exist or is down. In Windows, this test behaves similarly to the `networkinterface-linkstatus` test described next.
networkinterface-linkstatus *interface*	Checks the link-layer status of the specified *interface*. The test fails if the specified network interface is not linked.
file-exists *path*	Checks whether a file exists on the server(s). The test fails if the file is not found at the specified *path*.

What’s Next?

- To see some examples that illustrate the internal test configuration, proceed to Monitoring Agent Internal Test Examples.
- To skip the examples and see how to enable the Monitoring Agents, proceed to Enabling Monitoring Agents (page 637).

Monitoring Agent Internal Test Examples

Each test has an example configuration, which may differ from what you configure for your own environment.

Swapfree Test Example

This example checks that there is at least 1500 kilobytes of free space on server 2. The test runs at one-hour intervals and sends an alert if the test fails.

```
test “swapfree test”
interval 1:00:00
host server2
action alert
command swapfree 1500
```

Processcount Test Example

This example checks that there are a maximum of 2500 processes running on server2. The test runs at 60-second intervals and sends an alert if the test fails.

```
test process_count
interval 60
host server2
action alert
command prosesscount 2500
```

Multi-ping Test Example

This example tests the connectivity of server2. The tester sends three ICMP Echo requests to the IP addresses 192.168.1.2 and 192.168.1.254, and waits 1000 milliseconds for a reply. The test runs at 30-second intervals. If the test fails, server2 is excluded from the Server Pool.

```
test “multi-ping test”
interval 30
host server2
action exclude
command multi-ping 3 1000 192.168.1.2 192.168.1.254
```

Filesystem Test Example

This example checks that there is at least 1000 kilobytes of free space in the /var/tmp partition on server2. The test runs at 15-second intervals and sends an alert if the test fails.

```
test “free space in filesystem”
interval 15
host server2
action alert
command filesystem /var/tmp 1000
```

Servicerunning Test Example

This example checks that the `ntpd` service is running. The test runs at 15-second intervals. If the test fails, the `ntpd start` script is run. If the test succeeds again within one minute of failing, recovery occurs and the server is returned to the OK state. Otherwise, the server is excluded from the Server Pool.

```
test “is ntp running -test”
interval 15
script /etc/init.d/ntpd start
action exclude
recovery 1:00
command servicerunning ntpd
```

Ip-listening Test Example

This example checks that server2 is listening at the IP address 192.168.1.1. The test runs at 30-second intervals, and the server is excluded from the Server Pool if the test fails.

```
test ip-listening
interval 30
host server2
action exclude
command iplistening 192.168.1.1
```

Port-listening Test Example

This example checks that server2 and server3 are listening on port 80. The test runs at 30-second intervals, and the server is excluded from the Server Pool if the test fails.

```
test port-listening
interval 30
host server2 server3
action exclude
command portlistening 80
```

Httpanswer Test Example

This example checks that the server sends the contents of the file `/web/file.html` as a part of its response to a URL request for `http://www.example.com/index.html` on port 80. The request is attempted four times, and the tester waits 3500 milliseconds for a response before retrying. The test runs at 30-second intervals, and the server is excluded from the Server Pool if the test fails.

```
test http_answering
interval 30
action exclude
command httpanswer 4 3500 /www.example.com/index.html /web/file.html 80
```

Networkinterface-up Test Example

This example checks that the `eth0` interface is up. The test runs at 45-second intervals, and an alert is sent if the test fails.

```
test etc0_up
interval 45
action alert
command networkinterface-up eth0
```

Networkinterface-linkstatus Test Example

This example checks that the `eth0` interface is linked. The test runs at 45-second intervals, and an alert is sent if the test fails.

```
test “link status on eth0”
interval 45
action alert
command networkinterface-linkstatus eth0
```

File-exists Test Example

This example checks that the `/important/index.html` file exists. The test runs at 2-minute intervals, and an alert is sent if the test fails.

```
test “index -file exists”
interval 2:00
action alert
command file-exists /important/index.html
```

What’s Next?

▶ Proceed to Enabling Monitoring Agents (page 637).

Enabling Monitoring Agents

Prerequisites: Defining a Server Pool, Installing Monitoring Agents, Configuring Monitoring Agents

Caution – Do not enable the Monitoring Agents before you configure them (see Configuring Monitoring Agents (page 623)).

▼ To enable the Monitoring Agents

1. Right-click the Server Pool element for which you want to enable the Monitoring Agents and select **Properties**. The Properties dialog opens.
2. Select **Agent** in the Monitoring section.
3. (*Optional*) Enter the **Port** if you want to use a port other than the default port (7777).

Note – Remember to modify the corresponding IPv4 and IPv6 Access rules to allow connections to the new port.

4. Click **OK**.

What's Next?

- If you are using static DNS entries, continue by Entering Server Pool IP Addresses on Your DNS Server.
- Otherwise, continue in Creating Access Rules for Inbound Load Balancing (page 638).

Entering Server Pool IP Addresses on Your DNS Server

Prerequisites: Defining a Server Pool

When using static DNS entries (recommended), make sure that the IP addresses for your servers, as specified in the Server Pool properties, are properly entered into your DNS server's records.

If the servers have multiple IP addresses, make sure these are all defined on the DNS server to ensure operation as intended.

Consult the documentation of your DNS server software for instructions on how to enter the IP addresses.

What's Next?

- Continue the configuration by Creating Access Rules for Inbound Load Balancing (page 638).

Creating Access Rules for Inbound Load Balancing

Prerequisites: Defining a Server Pool

The IPv4 Access rules specify which traffic is directed to the server pool.

Note the following:

- The Server Pool does automatic NAT from the external addresses you configured in the Server Pool element to the addresses of the included servers, so make sure there are no overlapping NAT rules in the policy. You can add a NAT rule that disables further NAT for matching connections (empty NAT cell), if necessary.
- If you want to balance traffic that arrives through a VPN using a Server Pool, NAT must be enabled in the properties of the VPN element (NAT is disabled by default for traffic that uses a VPN).
- You must create a separate rule for each Server Pool.
- If the same Server Pool provides more than one service, you must create a separate rule for each Service.
- You must enable Connection Tracking for the rule that directs traffic to the Server Pool. The Server Pool uses NAT, which does not work without Connection Tracking.

▼ To create an access rule for inbound load balancing

1. Open the Firewall Policy for editing and add an IPv4 Access rule.
2. Configure the rule to match the **Source**, **Destination**, and **Service** of the traffic you want to direct to the server pool.

Note – Each rule must contain only one Service.

3. Set the **Action** to **Allow** and enable **Connection Tracking** in the **Action Options**.

Example **The following example rules direct traffic from outside networks to the HTTP Server Pool and to the HTTPS Server Pool.**

Source	Destination	Service	Action
NOT internal Expression	HTTP Server Pool	HTTP	Allow Connection tracking: normal
NOT internal Expression	HTTPS Server Pool	HTTPS	Allow Connection tracking: normal

What's Next?

- ▶ If you are using Dynamic DNS entries, continue the configuration by Defining a Dynamic DNS Rule (page 642).
- ▶ If you are using static DNS entries, the configuration is complete. Save and Install the Firewall Policy to transfer the changes.

Configuring Dynamic DNS Updates

Prerequisites: See Getting Started with Inbound Traffic Management

Firewalls support the Dynamic DNS protocol and can send DDNS updates to a specified DNS server. If a network connection specified by a NetLink element fails, the dynamic DNS updates notify the DNS, which then removes the corresponding IP address from its records.

To configure DDNS updates, you must have already defined the necessary NetLinks and the Server Pool element. To use DDNS updates, you must set up a DDNS-capable DNS server in your network. The DNS server must be configured as the primary DNS server for the domain.

Using dynamic DNS updates is a security risk. Before you start using dynamic DNS updates, read the section Improving DDNS Security.

Configuration Overview

1. Implementing DDNS may be a security risk. Read the information in Improving DDNS Security before proceeding with the configuration.
2. To enable monitoring of the status of your NetLinks, you must add probe IP addresses in the NetLinks' properties as instructed in Creating NetLinks (page 601).
3. Define the DNS Server element (Defining an External DNS Server (page 640)).
4. Modify the Server Pool element to include information on how the DNS server is updated (see Defining the Dynamic DNS Update Information (page 641)).
5. Modify the active policy to allow DDNS updates (see Defining a Dynamic DNS Rule (page 642)).

Improving DDNS Security

Caution – Although Firewalls support dynamic DNS updates, the protocol itself poses a security risk because there is no access control. If you must use dynamic DNS updates, do so only after very careful research, planning, and testing.

This section presents some actions to take to improve the security of dynamic DNS updates:

- Always place the DNS server(s) behind the Firewall for protection from IP spoofing.
- Use BIND or an equivalent DNS server that allows you to define which hosts are allowed to send dynamic updates.
- Consider using static DNS entries instead, as DDNS is not necessarily needed with inbound load balancing. In that case the DNS entries are not removed automatically from the DNS server if an ISP fails, but these problems can sometimes be solved by other means. For example, some web browsers can automatically try other IP addresses if one address does not respond.

What's Next?

▶ To start using DDNS with inbound traffic management, continue to Defining an External DNS Server.

Defining an External DNS Server

You must define an External DNS Server element in the following cases:

- (*Firewalls only*) For DDNS updates with a Multi-Link configuration.
- (*Firewalls only*) If you want to use a DNS server for resolving virus signature mirrors.
- If you want to use a DNS server for resolving domain names and web filtering categorization services on Firewalls, IPS engines, and Layer 2 Firewalls.

▼ To define an External DNS Server element

1. Select **Configuration→Configuration→Security Engine**. The Security Engine Configuration view opens.
2. Expand the **Network Elements** branch.
3. Right-click **Servers** and select **New→External DNS Server**. The External DNS Server Properties dialog opens:

Illustration 38.4 External DNS Server Properties

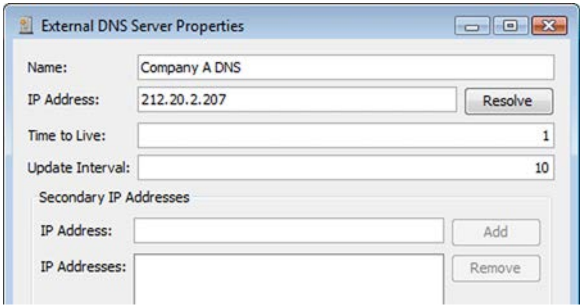

External DNS Server Properties

Name: Company A DNS

IP Address: 212.20.2.207

Time to Live: 1

Update Interval: 10

Secondary IP Addresses

IP Address:

IP Addresses:

4. Enter a unique **Name** and **IP address** for the server.
5. Enter a **Time to Live** (TTL) interval in seconds. It defines how long a DNS entry can be cached before querying the DNS server again. The default is 1 second.
6. Enter an **Update Interval** in seconds. It defines how often the DNS entries can be updated to the DNS server if the link status changes constantly. The default is 10 seconds.
7. (*Optional*) If the device has additional IP addresses, you can enter them as **Secondary IP Addresses** instead of creating additional elements. However, secondary IP addresses are only used in the Source and Destination cells in rules. They are ignored otherwise.
8. Click **OK**.

What's Next?

▶ Continue the DDNS configuration by Defining the Dynamic DNS Update Information (page 641).

Related Tasks

▶ Adding Custom Commands to Element Menus (page 55)

Defining the Dynamic DNS Update Information

Next, you must define the settings for the Dynamic DNS updates sent from your Firewall to the DNS server. This data is configured in the Server Pool element.

▼ To define DDNS update information

1. Select **Configuration→Configuration→Security Engine**. The Security Engine Configuration view opens.
2. Browse to **Network Elements→Traffic Handlers**.
3. Right-click the Server Pool and select **Properties**. The Properties dialog opens.

Illustration 38.5 Server Pool Properties

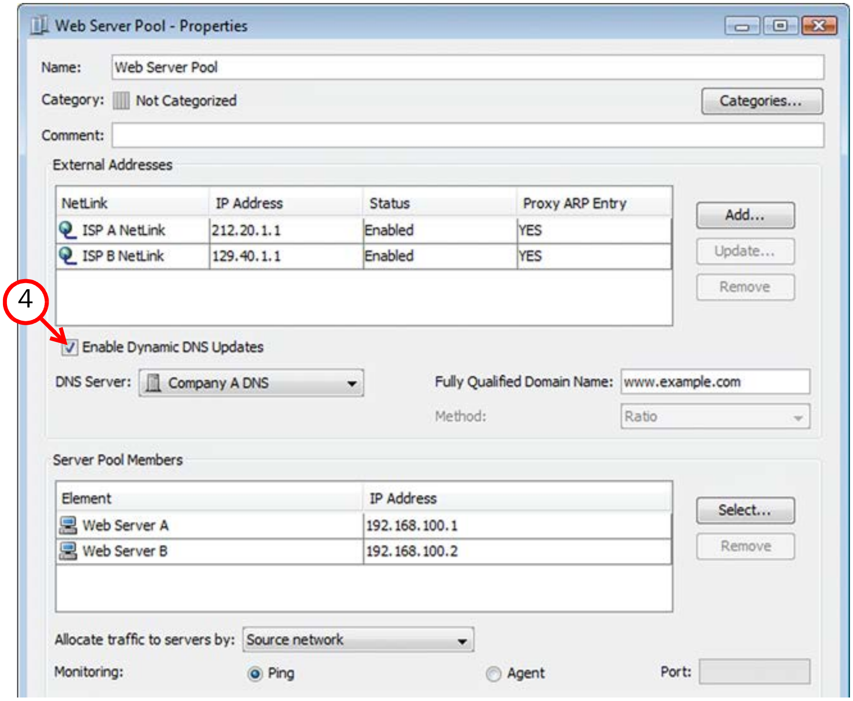

Web Server Pool - Properties

Name: Web Server Pool

Category: Not Categorized

Comment:

External Addresses

NetLink	IP Address	Status	Proxy ARP Entry
ISP A NetLink	212.20.1.1	Enabled	YES
ISP B NetLink	129.40.1.1	Enabled	YES

☑ Enable Dynamic DNS Updates

DNS Server: Company A DNS

Fully Qualified Domain Name: www.example.com

Method: Ratio

Server Pool Members

Element	IP Address
Web Server A	192.168.100.1
Web Server B	192.168.100.2

Allocate traffic to servers by: Source network

Monitoring: ● Ping ○ Agent Port:

4. Select **Enable Dynamic DNS Updates**.
5. Select the **DNS Server** element to which the DDNS updates are sent.
6. Enter the **Fully Qualified Domain Name** for the Server Pool service (for example, “`www.example.com`”).
7. Click **OK**.

What’s Next?

- ▶ Continue by Defining a Dynamic DNS Rule (page 642).

Defining a Dynamic DNS Rule

In addition to creating a rule for inbound load balancing (as explained in Creating Access Rules for Inbound Load Balancing (page 638)), you may also need to add a rule allowing the Dynamic DNS updates from your Firewall to the specified DNS server. The Firewall Template allows this traffic by default, so if your policy is based on an unmodified Firewall Template, there is no need to add the rule outlined below.

▼ To add a rule that allows Dynamic DNS updates

1. Right-click the Firewall Policy and select **Edit**. The policy opens for editing.
2. Add the following Access rule:
   - **Source**: the address (towards the DNS server) of the Firewall or NDIs of all nodes in the Firewall Cluster.
   - **Destination**: the DNS server
   - **Service**: the **DNS** Service Group.
   - **Action**: **Allow**
3. Save and install the policy.

Monitoring and Testing Monitoring Agents

Prerequisites: Installing Monitoring Agents, Configuring Monitoring Agents

You can test and monitor the Server Pool Monitoring Agents on the command line. The monitoring tool allows you to query the agents locally or remotely to check their status. The Monitoring Agent command line tools are described in Server Pool Monitoring Agent Commands (page 1164).

▼ To test a Monitoring Agent

1. On the command line, browse to the folder where the Monitoring Agent is installed.
2. Run one of the following commands:
   - **Linux**: `sgagentd -d test [filename]`
   - **Windows**: `agent test [filename]`

   Where the *filename* is a configuration you want to test. Giving the filename is optional, but if you omit the filename, the test is run with the active configuration.

▼ To monitor a Monitoring Agent

1. On the command line, browse to the folder where the Monitoring Agent monitoring tool is installed:
   - In Windows, this is the same folder as where the Monitoring Agent itself is installed.
   - On Linux and Solaris, the program is installed in `usr/bin`.
2. (*Optional*) Copy the `sgmon` program and copy it to a remote host that you want to use for testing the Monitoring Agents.
3. Run `sgmon -h` to see usage instructions and run your monitoring commands according to the syntax displayed.
   - The *`host`* argument is mandatory. To get the status locally on the server where the Monitoring Agent is installed, use **`localhost`** as the *`host`* argument.

TRAFFIC INSPECTION POLICIES

In this section:

CHAPTER 39

CREATING AND MANAGING POLICY ELEMENTS

Policy elements are containers for the rules that determine how traffic is filtered and inspected.

The following sections are included:

Getting Started with Policies

Prerequisites: None

What Policy Elements Do

The policies organize traffic processing rules hierarchically to make the administration easier and to optimize traffic inspection performance:

- Firewall, Layer 2 Firewall, and IPS Policies contain the rules according to which the Security Engines allow or block traffic. The same policy can be shared by several Firewall engines, several IPS engines, several Layer 2 Firewall engines, several Master Engines, or several Virtual Firewall engines.
- Inspection Policies contain the rules according to which the Security Engines inspect traffic. The same Inspection Policy can be shared by several Firewall Policies, IPS Policies, and Layer 2 Firewall Policies.
- Each policy must always be based on a *Template Policy*. Template Policies contain rules that are inherited into any template or policy below it in the policy hierarchy.
- You can also insert *Sub-Policies* in your policies. A Sub-Policy is a set of IPv4 or IPv6 Access rules that can be matched conditionally to a restricted part of the traffic and may therefore improve processing performance. Sub-Policies can also enforce administrative boundaries.
- Policies can share Policy Templates and Sub-Policies. In shared rules, Alias elements can represent IP addresses that depend on the environment, so that the actual values are defined separately for each component.

What Do I Need to Know Before I Begin?

- There is no separate type of policy for Master Engines and Virtual Firewalls. Master Engines and Virtual Firewalls use Firewall Policies.

Default Policy Elements

Ready-made Policies, Template Policies, and Inspection Policies provide an easy starting point for determining what kinds of rules your system needs. The system has the following default policy elements:

Table 39.1 Default Policy Elements for Security Engines

Element Type	Default Element Name	Description
Firewall Template Policy	Firewall Template	A Template Policy that contains the predefined Access rules necessary for the firewall to communicate with the Management Center and some external components. The Firewall Template uses the Inspection rules defined in the No Inspection Policy. The Firewall Template provides access control without deep inspection.
	Firewall Inspection Template	A Template Policy that is based on the Firewall Template. It uses Inspection rules defined in the High-Security Inspection Policy. The Firewall Inspection Template enables deep inspection for all traffic.

Table 39.1 Default Policy Elements for Security Engines (Continued)

Element Type	Default Element Name	Description
Firewall Sub-Policy	DHCP Relay	A Sub-Policy that contains rules that allow the firewall to relay DHCP requests from a host in one internal network to a DHCP server in a different network, as well as DHCP requests from VPN clients to an internal DHCP server.
IPS Template Policy	IPS Template	A Template Policy that contains the predefined Access rules necessary for the IPS engine to communicate with the Management Center and some external components. The IPS Template Policy uses Inspection rules from the High-Security Inspection Policy. The IPS Template Policy provides an easy starting point for determining what kinds of rules your system needs.
IPS Policy	Customized High-Security Inspection IPS Policy	An IPS Policy that is based on the IPS Template. The Customized High-Security Inspection IPS Policy contains a set of customized rules that were used when Stonesoft IPS was tested at ICSA Labs and NSS Labs.
	Default IPS Policy	An IPS Policy that is based on the IPS Template. The Default IPS Policy does not add any rules to those defined in the IPS Template. It allows you to install the predefined rules in the IPS Template on the IPS engine right after installation (Template Policies cannot be installed on the engines).
Layer 2 Firewall Template Policy	Layer 2 Firewall Template	A Template Policy that contains the predefined Access rules necessary for the Layer 2 Firewall to communicate with the Management Center and some external components. The Layer 2 Firewall Template uses Inspection rules from the No Inspection Policy. The rules in the No Inspection Policy do not enforce inspection.
	Layer 2 Firewall Inspection Template	A Template Policy that is based on the Layer 2 Firewall Template. It uses Inspection rules from the High-Security Inspection Policy. The Layer 2 Firewall Inspection Template enables deep inspection for all traffic.
Inspection Policy	No Inspection Policy	An Inspection Policy with a set of Inspection rules that do not enforce inspection.
	Medium-Security Inspection Policy	An Inspection Policy with a set of Inspection rules for detecting common threats. The Medium-Security Inspection Policy logs Situations categorized as Suspected Attacks but allows the traffic to pass. The Medium-Security Inspection Policy is suitable for Firewall, Layer 2 Firewall, Master Engine, and Virtual Security Engine deployments. It is also suitable for inline IPS deployments in asymmetrically-routed networks and IPS deployments in IDS mode. The risk of false positives is low in production use.

Table 39.1 Default Policy Elements for Security Engines (Continued)

Element Type	Default Element Name	Description
Inspection Policy *(cont.)*	High-Security Inspection Policy	An Inspection Policy with a set of Inspection rules for detecting common threats. The High-Security Inspection Policy terminates Suspected Attacks with an alert. The High-Security Inspection Policy is suitable for Firewall, Layer 2 Firewall, inline IPS, Master Engine, and Virtual Security Engine deployments where extended inspection coverage and strong evasion protection is required. The risk of false positives is moderate in production use. The High-Security Inspection Policy terminates a connection if the security engine cannot see the whole connection. We recommended that you use the High-Security Inspection Policy as a starting point for your Inspection Policies.
	Customized High-Security Inspection Policy	An Inspection Policy that is based on the High-Security Inspection Policy and contains a set of customized Inspection rules. The High-Security Inspection Policy is an example of a highly customized Inspection Policy for network environments in which unconditional inspection coverage and evasion protection are required. The risk of false positives is high in production use. The High-Security Inspection Policy was used when Stonesoft IPS was tested at ICSA Labs and NSS Labs. It provides an example of a customized Inspection Policy.

The default policy elements are introduced into the system when you import and activate a recent dynamic update package (for example, during the installation). The elements may change when you install newer update packages. None of the default policy elements can be modified. However, you can make copies of the default policies if you need to create a modified version.

See the *Firewall/VPN Reference Guide* and the *IPS and Layer 2 Firewall Reference Guide* for more background information and examples on using policies, template policies, sub-policies, and the rules in the policies.

Configuration Overview

1. (*Optional*) Create a custom template as explained in Creating a New Template Policy or a Policy (page 649) and add the rules and insert points.
2. Create a custom policy as explained in Creating a New Template Policy or a Policy (page 649) and add the rules.
3. (*Optional*) Create sub-policies according to need as explained in Creating a New Sub-Policy (page 650) and add the IPv4 Access rules.

Related Tasks

- Installing Policies (page 652)
- Tracking Policy Changes (page 654)
- Moving the Policy Under a Different Template (page 656)
- Deleting Policies, Templates, and Sub-Policies (page 657)
- Getting Started with Editing the Rules in Policies (page 660)

Creating a New Template Policy or a Policy

Prerequisites: None

All new Policies are based on a Template Policy. The template rules are inherited to lower levels in the policy hierarchy. Template Policies can enforce administrative boundaries and prevent modifications from unintentionally changing the main design of your rules. You must add `insert points` to templates for all types of rules to mark the places where rules can be inserted in Policies (or Template Policies) that use the template.

▼ To create a new Policy or Template Policy

1. Select **Configuration→Configuration→Security Engine**. The Security Engine Configuration view opens.
2. Right-click **Policies**.
3. Select **New** and the type of policy you want to create. The Properties dialog opens.

Illustration 39.1 Firewall Template Policy Properties

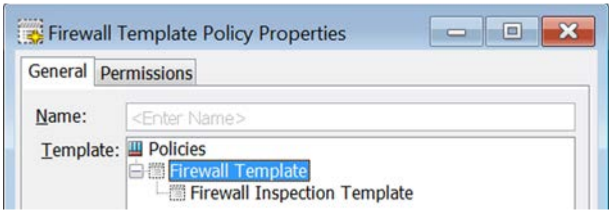

4. Give the element a **Name**.
5. Select the **Template** you want to base this template or policy on.
6. (*Optional*) Switch to the **Permissions** tab and grant rights for the template or policy.
   - To add a permission, click **Add Permission**. A new row appears on the administrator list.
   - Click the **Administrator** cell and select the Administrator.
   - Double-click the **Administrator Role** cell to select the correct role.
7. Click **OK**. The new Template Policy or Policy opens in the Policy Editing view.

If you modified administrator permissions for the policy, the changes are applied immediately. The permissions are also automatically updated in the properties of the administrator’s account.

What's Next?

- If you created a new Template Policy, proceed to Adding Insert Points in Policy Templates (page 668).
- Otherwise, add the rules. See Getting Started with Editing the Rules in Policies (page 660).

Creating a New Sub-Policy

Prerequisites: None

A Sub-Policy is a set of IPv4 or IPv6 Access rules that you can reuse in different places and in several Firewall, IPS Policies, or Layer 2 Firewall Policies.

The main goal of Sub-Policies is to match only part of the traffic against the rules in the Sub-Policy and allow other traffic to bypass the Sub-Policy. Rules in a sub-policy should have at least some identical matching criteria (source, destination, service, source VPN, and/or user authentication details) that can be used to select only a portion of the traffic for the sub-policy checks.

Sub-policies can also be used to organize the policies and to delegate administrator privileges: you can restrict specific administrators to edit/add/remove rules within a limited section of IPv4 or IPv6 Access rules (which is more restrictive than giving access to a template or policy).

There are two ways to create Sub-Policies: you can either create a Sub-Policy and add Access Rules to it manually or select Access rules in a policy and convert them into a Sub-Policy.

What's Next?

- To start from scratch, proceed to Creating a New Empty Sub-Policy.
- To use existing rules, proceed to Converting Existing Rules into a Sub-Policy (page 651).

Creating a New Empty Sub-Policy

▼ To create a new empty Sub-Policy

1. Select **Configuration→Configuration→Security Engine**. The Security Engine Configuration view opens.
2. Right-click the **Policies** branch and select **New→Firewall Sub-Policy**, **New→IPS Sub-Policy** or **New→Layer 2 Firewall Sub-Policy**. The Sub-Policy Properties dialog opens.

Note – To create a Sub-Policy for use on Master Engines or Virtual Firewalls, select Firewall Sub-Policy.

3. Enter a unique **Name**.
4. (*Optional*) Switch to the **Permissions** tab and adjust the Access Control Lists at the top part of the dialog to include the Sub-Policy on one or more custom Access Control Lists.

5. Click **OK**. The new sub-policy opens in the Policy Editing view.

What's Next?

- Add the rules. See Editing Access Rules (page 672).

Converting Existing Rules into a Sub-Policy

You can convert IPv4 and IPv6 Access rules in an existing policy into a Sub-Policy. The rules do not have to be consecutive, but if you add several references to a sub-policy in the same policy, all sub-policy rules are checked at each reference point—even if those rules were already checked at a previous reference point. This can be avoided, for example, by adding a rule at the end of the Sub-Policy that stops all connections that did not match the other rules.

▼ To create a sub-policy from existing Access rules

1. Right-click the policy or template and select **Edit Firewall Policy**, **Edit IPS Policy**, or **Edit Layer 2 Firewall Policy**.
2. Switch to the **IPv4 Access** or **IPv6 Access** tab.
3. Select Access rules that you want to add to the sub-policy.
4. Right-click one of the selected rules and select **Create Sub-Policy**. The Sub-Policy dialog opens.
5. Enter a **Name** for the Sub-Policy and click **OK**. The Sub-Policy element is created, a new Jump rule that references the Sub-Policy is automatically added to the policy, and the selected rules are moved to the Sub-Policy.
6. Edit the Jump rule cells to be as specific as possible so that traffic is not unnecessarily matched to the sub-policy. If necessary, you can add more references to the Sub-Policy, for example, by copy-pasting the Jump rule.
7. (*Optional*) Add the Sub-Policy to a custom Access Control List:

   7a. Right-click the **Action** cell in the Jump rule and select **Properties**. The Properties dialog for the Sub-Policy opens.

   7b. Switch to the **Permissions** tab and adjust the Access Control Lists at the top part of the dialog.

   7c. Click **OK**.

Related Tasks

- Getting Started with Editing the Rules in Policies (page 660)

Installing Policies

Prerequisites: The policy you want to install must be allowed in the engine's configuration

Changes to an engine's rules and most other configurations are activated in a policy installation. Only Policy elements can be installed; Template Policy and Sub-Policy rules are installed as part of the main Policy. A policy snapshot is automatically created each time you install or refresh a policy. You can install a policy through the Policy element or through the engine element. The procedure below explains the first method.

▼ To install a policy

1. Select **Configuration→Configuration→Security Engine**. The Security Engine Configuration view opens.
2. Expand the **Policies** branch and select the appropriate type of policies.
3. Right-click the policy you want to install and select **Install Policy** from the menu. The Task Properties dialog opens.

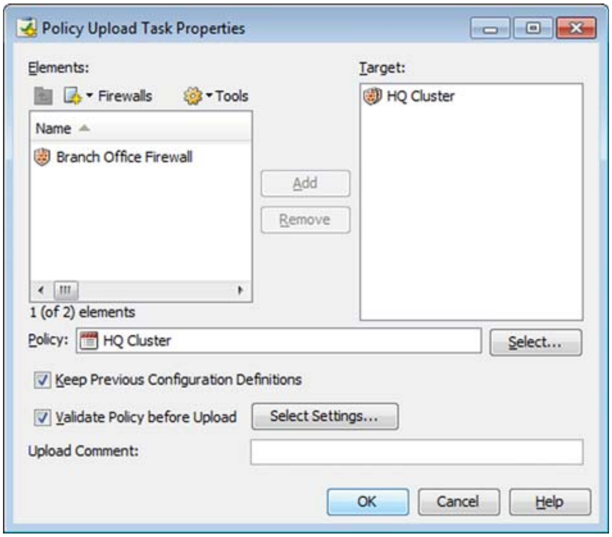

Branch Office Firewall

1 (of 2) elements

Target: HQ Cluster

Policy: HQ Cluster

☑ Keep Previous Configuration Definitions

☑ Validate Policy before Upload

4. Select the engine(s) on which you want to install the same policy and click **Add**. You can install the same policy on several engines in one operation.
   - The system tracks the policies installed on engines and pre-fills the Target when possible.
5. (*Optional*) Leave **Keep Previous Configuration Definitions** selected to allow established connections to continue using existing configurations (such as NAT rules) until they finish.
   - If the previous configurations are erased, connections that use them are cut.
   - All previously established connections that are not allowed by the newly installed policy are always dropped regardless of this setting.
6. (*Optional*) Leave **Validate Policy before Upload** selected to validate the rules in the policy. See Validating Rules Automatically (page 714) for more information on the related settings.

Note – You cannot validate the policy if you are installing the policy on several engines.

7. (*Optional*) Add an **Upload Comment**. The comment provides history information to administrators and Web Portal users when viewing Policy Snapshots.

8. Click **OK** to validate and install the policy. The progress is displayed on a new tab.

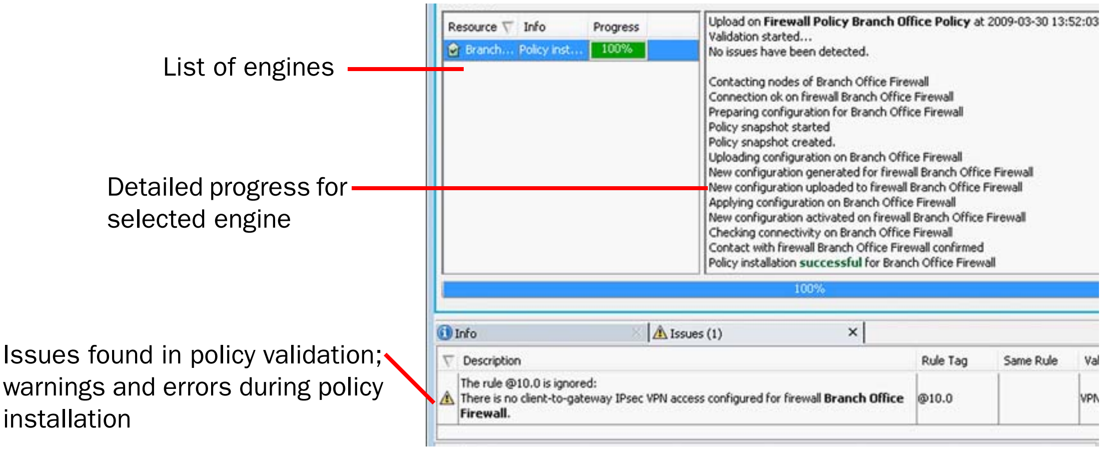
List of engines
Detailed progress for selected engine
Issues found in policy validation; warnings and errors during policy installation
100%
Info
Issues (1)

Resource	Info	Progress
Branch...	Policy inst...	100%

```
Upload on Firewall Policy Branch Office Policy at 2009-03-30 13:52:03
Validation started...
No issues have been detected.

Contacting nodes of Branch Office Firewall
Connection ok on firewall Branch Office Firewall
Preparing configuration for Branch Office Firewall
Policy snapshot started
Policy snapshot created.
Uploading configuration on Branch Office Firewall
New configuration generated for firewall Branch Office Firewall
New configuration uploaded to firewall Branch Office Firewall
Applying configuration on Branch Office Firewall
New configuration activated on firewall Branch Office Firewall
Checking connectivity on Branch Office Firewall
Contact with firewall Branch Office Firewall confirmed
Policy installation successful for Branch Office Firewall
```

Description	Rule Tag	Same Rule	Val
The rule @10.0 is ignored: There is no client-to-gateway IPsec VPN access configured for firewall **Branch Office Firewall**.	@10.0		VPN

9. If validation issues are found, read the Issues tab in the Info panel and take action:
   - Double-click an issue to view the corresponding configuration or click **Continue**.
   - For more information, including steps to disable warnings, see Viewing Policy Validation Issues (page 717).

10. Monitor the installation and make sure it is successful. With multiple engines, the progress is indicated through colored icons on the left (click the icon to view the details):
    - Yellow: ongoing installation.
    - Blue: waiting for the installation(s) on other component(s) to finish.
    - Red: failure
    - Green: success

Related Tasks

- ▶ Tracking Policy Changes (page 654)
- ▶ Selecting Permitted Policies for Engines (page 531)

Tracking Policy Changes

Prerequisites: None

Checking the Currently Installed Policy

▼ To check the currently installed policy

1. Right-click the engine icon (for clusters, click the cluster icon) and select **Current Policy→Info**.
2. A message is displayed on the screen with the following information:
   - Name of the installed policy
   - Name of the administrator who installed the policy
   - Date (year-month-day) and time of the policy installation.

Illustration 39.2 Current Security Policy Message Box

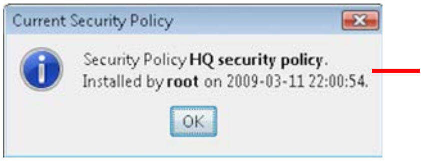

Current Security Policy

Security Policy **HQ security policy**.
Installed by **root** on 2009-03-11 22:00:54.

Date format is year-month-day

3. Click **OK** to close the message box.

Previewing the Currently Installed Policy

▼ To preview the currently installed policy

➥ Right-click the engine icon (for clusters, click the cluster icon) and select **Current Policy→Preview**. A preview of the policy in its current format on the Management Server opens.

Note – The policy on the Management Server may be different from the actual currently installed policy if the policy has been modified after it was installed.

Checking and Comparing Policy Versions

Each time a policy is successfully installed, a record of that configuration is stored in the upload history. This allows you to the check which policies were uploaded and when they were uploaded, and allows you to run an automatic check for policy refresh needs. For example, you can compare two policy snapshots, or check which policy was in place at a particular time.

Note – The policy snapshot is only stored when the policy is successfully installed. No snapshot is stored if the policy installation fails.

What's Next?

▶ Viewing Policy Snapshots (page 655)

▶ Comparing Two Policy Snapshots (page 655)

Viewing Policy Snapshots

▼ To view a policy snapshot

1. Select **Configuration→Configuration→Security Engine**. The Security Engine Configuration view opens.
2. Browse to **Other Elements→Policy Snapshots**. A list of policy snapshots with the upload times and dates appears.
3. Double-click the policy snapshot you want to view. The details of the selected policy snapshot opens.

Tip – You can also view the snapshot in HTML format: right-click a snapshot and select **Tools→Save as HTML File**.

4. Select elements in the other panel to see their details. These are organized in the following groups:
   - Policy: the policy whose installation created the snapshot.
   - Target: properties of the engine on which the policy has been saved.
   - Elements: elements included in the policy.

Tip – Select **View→Show XML** to view the policy snapshot in the XML format.

Comparing Two Policy Snapshots

You can compare any two policy snapshots from the list to check for changes between policy installations.

▼ To compare policy snapshots

1. Open the comparison in one of the following ways:
   - Select two policy snapshots you want to compare, right-click, and select **Compare Snapshots** to compare any two snapshots to each other.
   - Right-click the policy snapshot that you want to compare and select **Compare→Compare to Engine's Current Policy** to compare it to the most recent policy snapshot.
   - Right-click the policy snapshot that you want to compare and select **Compare→Compare to Latest Saved Version** to compare it to the current working version stored on the Management Server.
2. Check the changes that are highlighted in the policy in colors and as a summary below.
3. Select objects in the summary to focus on them. For example, if you select a network element, its properties are shown above. The objects are organized in the following groups:
   - Targets: The engine on which the policy has been saved.
   - New Elements: Elements added to the policy.
   - Removed Elements: Elements that have been removed from the policy.
   - Modified Elements: Elements that exist in both policies but have been modified.

Checking for Untransferred Configuration Changes

Most changes done in any part of the configuration are transferred to the engines when the Security Policy for that engine is refreshed. You can view a list of engines for which the Security Policy has not been refreshed after the latest changes that affect the engine's operation. You can also view a list of the Log Servers for which the Alert Policy needs to be refreshed.

▼ Checking for policy refresh needs

1. In the System Status view, select the element(s) that you want to include in the check.
2. Right-click the selected category, and select **Configuration→Policy Refresh Check**. The Policy Refresh check dialog opens and the check is performed (how long it takes depends on how many elements you selected).

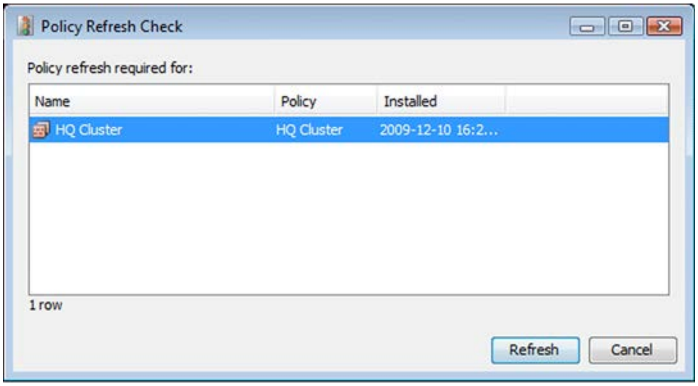

Policy refresh required for:

Name	Policy	Installed
HQ Cluster	HQ Cluster	2009-12-10 16:2...

1 row

3. If elements are shown, select the engines whose policies you want to refresh (all elements are selected by default).
4. Click **Refresh**. The policy upload for the selected elements begins.

Moving the Policy Under a Different Template

You can change the template of a Firewall, IPS, or Layer 2 Firewall policy. For example, if you created policy *A* based on template *X*, you can later reparent policy A to inherit rules from template *Y*.

▼ To switch the template of a policy

1. Select **Configuration→Configuration→Security Engine**. The Security Engine Configuration view opens.
2. Expand the **Policies** tree and select the appropriate policy type.
3. Right-click the correct policy and select **Properties**. The Properties dialog for the policy opens showing the policy tree.
4. Select the new parent template in the policy tree. The parent template must have insert points for the types of rules that the Policy contains.
5. Click **OK**. The policy is moved under your chosen template in the tree view, and starts using rules from the new template.

Deleting Policies, Templates, and Sub-Policies

Prerequisites: You must have created a policy, a policy template, or a sub-policy

To delete a template, there must be no policies that are based on that template. The deletion is not allowed until you either delete the policies or switch them to use a different template. See Moving the Policy Under a Different Template (page 656).

To delete a Sub-Policy, there must be no policies that have Jump rules that use the Sub-Policy. The deletion is not allowed until you edit or remove the Jump rules in the policies.

See Marking Elements as Obsolete (page 88) for general information on deleting elements.

Chapter 40

Editing Policies

The rules in Firewall, IPS, and Layer 2 Firewall Policies allow you to control how the engines inspect and filter network traffic, and how NAT (network address translation) is applied on Firewalls, Master Engines, and Virtual Firewalls.

The following sections are included:

Getting Started with Editing the Rules in Policies

Prerequisites: Creating a New Template Policy or a Policy

What Rules Do

Rules are instructions to the engines for handling traffic. There are five main types of rules.

- *Ethernet rules* (IPS and Layer 2 Firewall only) filter traffic based on MAC addresses and low-level network protocols. These rules can be used to segment a network.
- *Access rules* filter traffic based on IP addresses and IP-based protocols. These rules control access to resources. There are separate Access rules for IPv4 and IPv6 traffic.
- *NAT rules* (Firewall, Master Engine, and Virtual Firewall only) change source and/or destination IP addresses in traffic that is allowed to pass through the firewall. NAT (network address translation) can hide the network structure and allows several computers to use the same IP address on the Internet. There are separate NAT rules for IPv4 and IPv6 traffic.
- *Inspection* rules in Inspection Policies filter traffic based on patterns in any of the information that is transferred. These rules log complex traffic usage patterns and find network attacks, network worms, or other worrying or unwanted traffic like the use of peer-to-peer file transfer applications.
- *Exceptions* in Inspection Policies create detailed exceptions to the Inspection rules to eliminate false positives and to activate blacklisting or User Responses for specific traffic patterns.

The engines process the rules one type at a time in the order listed above. IPv4 and IPv6 traffic may be matched to both IPv4 and IPv6 Access rules in any order if traffic is tunneled (for example, IPv6 tunneled in IPv4 or IPv4 tunneled in IPv6), possibly several times.

Basic Rule Design Considerations

Rule tables are read from the top down, so the order of the rules is important. Make sure that the rules advance logically from specific rules at the top toward more general rules at the bottom whenever the matching criteria in rules overlap.

Example **A rule that denies access to your server from a particular network must be placed above a more general rule that allows access to that server from any source address.**

Any two rules that have completely overlapping matching criteria are redundant and should be merged. Automatic rule validation can be used to find such mistakes.

When rules are matched to traffic, the traffic is compared to each rule one by one until a match is found. What happens when the end of the rule table is reached without any matches varies by the component and the type of rules.

What Do I Need to Know Before I Begin?

There is no separate type of policy for Master Engines and Virtual Firewalls. Master Engines and Virtual Firewalls use Firewall Policies.

What's Next?

▶ Using the Policy Editing View (page 661)

Related Tasks

- Editing Ethernet Rules (page 668)
- Editing Access Rules (page 672)
- Editing Firewall NAT Rules (page 687)
- Editing Inspection Policies (page 699)
- Validating Rules Automatically (page 714)
- Changing Default Rules (page 718)

Using the Policy Editing View

Prerequisites: Creating a New Template Policy or a Policy

You can open a policy for editing through the right-click menu for the policy or through the preview mode. Only one administrator at a time can edit a policy. Save your changes and close the policy editing view when you are finished.

Illustration 40.1 Policy Editing View (IPv4 Access Tab)

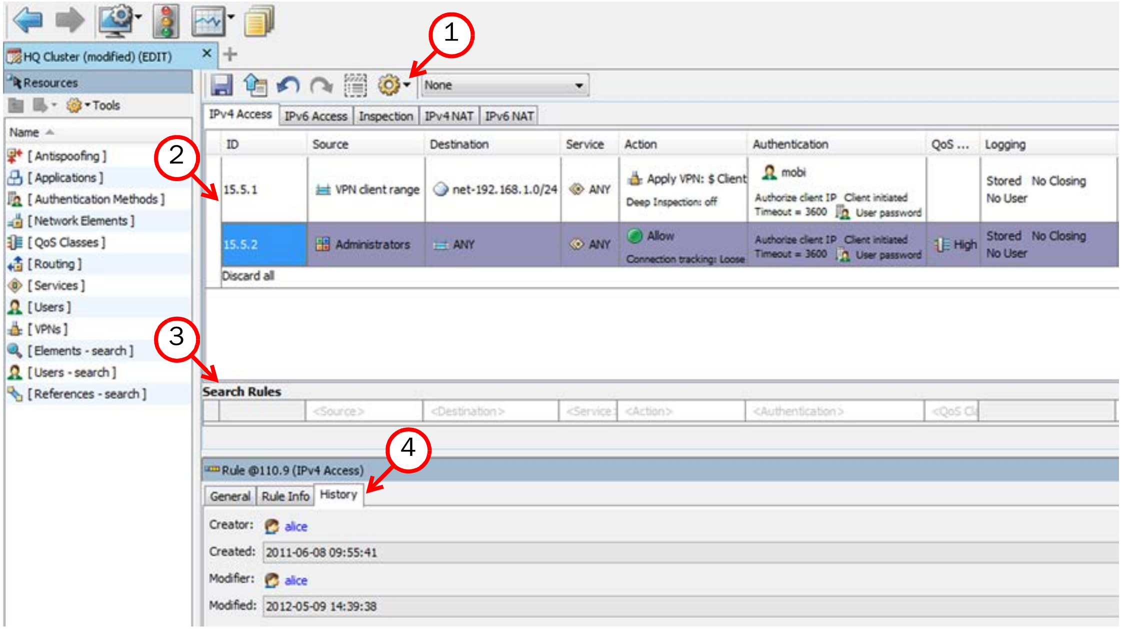

ID	Source	Destination	Service	Action	Authentication	QoS ...	Logging
15.5.1	VPN client range	net-192.168.1.0/24	ANY	Apply VPN: $ Client Deep Inspection: off	mobi Authorize client IP Client initiated Timeout = 3600 User password		Stored No Closing No User
15.5.2	Administrators	ANY	ANY	Allow Connection tracking: Loose	Authorize client IP Client initiated Timeout = 3600 User password	High	Stored No Closing No User
Discard all							

Creator: alice

Created: 2011-06-08 09:55:41

Modifier: alice

Modified: 2012-05-09 14:39:38

Legend: 1 - Policy toolbar. **2** - Rule table. **3** - Search tool. **4** - Change history for selected rule.

Illustration 40.2 Inspection Policy

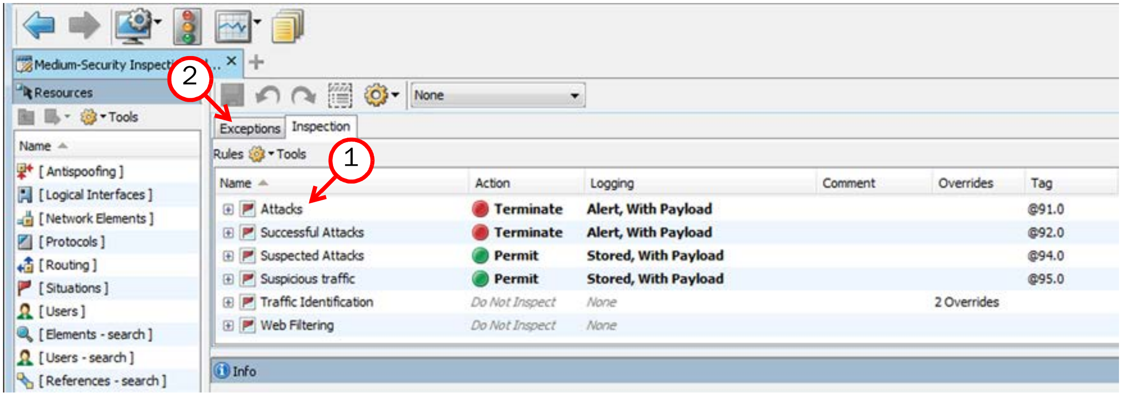

Name	Action	Logging	Comment	Overrides	Tag
Attacks	Terminate	Alert, With Payload			@91.0
Successful Attacks	Terminate	Alert, With Payload			@92.0
Suspected Attacks	Permit	Stored, With Payload			@94.0
Suspicious traffic	Permit	Stored, With Payload			@95.0
Traffic Identification	*Do Not Inspect*	*None*		2 Overrides	
Web Filtering	*Do Not Inspect*	*None*			

Legend: 1 - The main Rules tree. **2** - Detailed Exceptions to the main Rules.

Illustration 40.3 Policy Toolbar

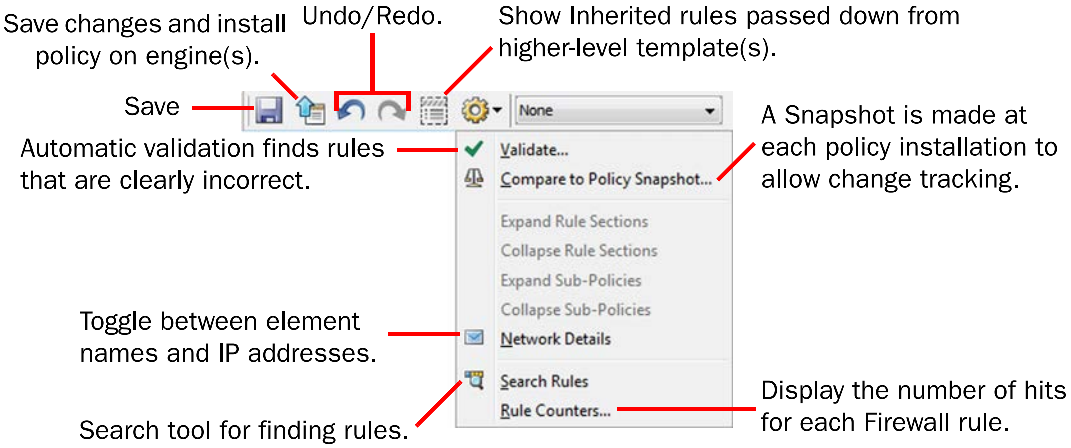

Save changes and install policy on engine(s).

Undo/Redo.

Show Inherited rules passed down from higher-level template(s).

Save

A Snapshot is made at each policy installation to allow change tracking.

Automatic validation finds rules that are clearly incorrect.

Toggle between element names and IP addresses.

Display the number of hits for each Firewall rule.

Search tool for finding rules.

Editing Rule Tables

You can edit the rule tables in the following ways:

- Use the actions in the right-click menu (add, cut, copy, paste, move etc.). If you right-click a cell that has cell-specific actions, the rule-specific actions are moved to the **Rule** submenu.
- Drag and drop (move) whole rules by the rule's ID or Tag number.

Illustration 40.4 Example of a Right-Click Menu for a Rule

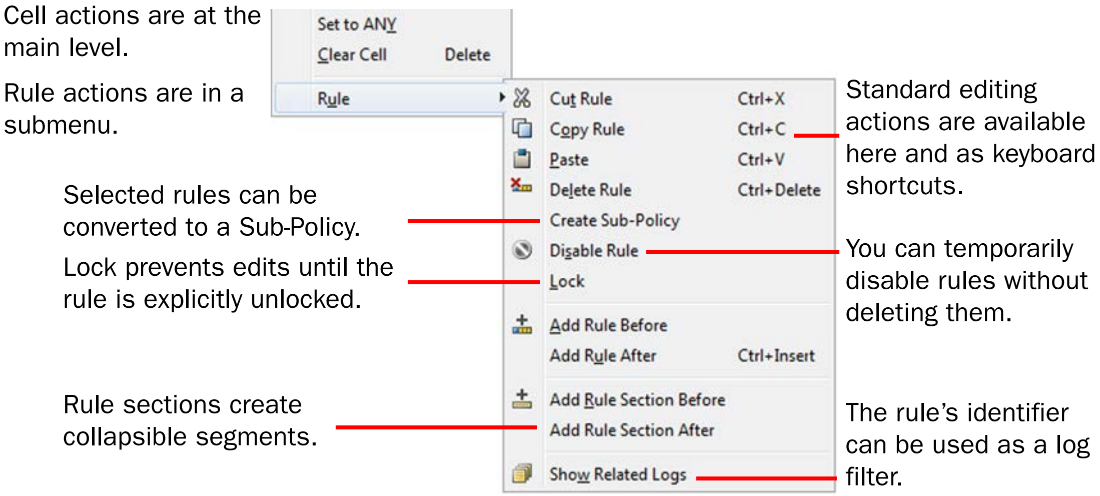

Cell actions are at the main level.

Rule actions are in a submenu.

Selected rules can be converted to a Sub-Policy.

Lock prevents edits until the rule is explicitly unlocked.

Rule sections create collapsible segments.

Standard editing actions are available here and as keyboard shortcuts.

You can temporarily disable rules without deleting them.

The rule's identifier can be used as a log filter.

Editing Rule Cells

Most rule cells require you to insert elements of specific types:

- When you click the cell, the Resources list on the left shows the type(s) of elements that you can insert in that cell.
- You can drag and drop the correct elements into the cell from the Resources list, from another cell, or even between tabs (the tab switches when you hover the pointer over a tab while dragging).
- You can define detailed sets of matching criteria in the Definitions dialog for the Source, Destination, and Service cells. See Defining Source, Destination, and Service Criteria (page 664).
- You can create new elements in the Resources list (of the type(s) that are currently displayed).

To modify rule cells that do not accept elements, right-click the cell and select an item from the right-click menu.

Illustration 40.5 Right-Click Menu for the Action Cell in a Firewall Access Rule

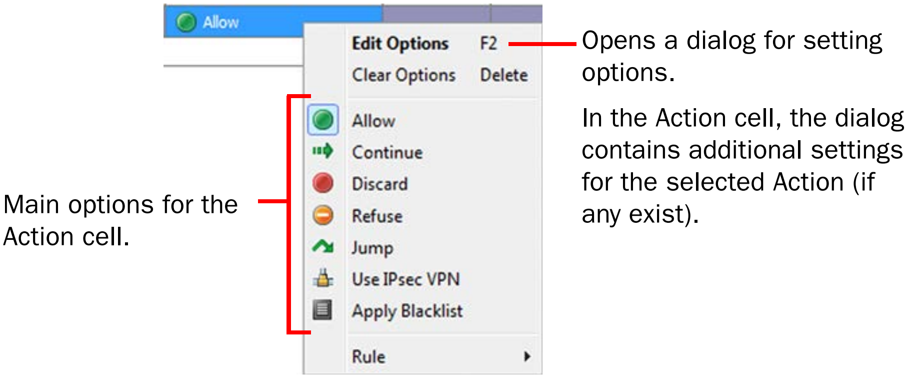

Opens a dialog for setting options.

In the Action cell, the dialog contains additional settings for the selected Action (if any exist).

Main options for the Action cell.

Defining Source, Destination, and Service Criteria

In addition to simply inserting elements into the Source, Destination, and Service cells, you can create detailed sets of matching criteria for the rule. You can create Source and Destination Definitions for all types of rules in Firewall, IPS, and Layer 2 Firewall Policies. You can create Service Definitions for IPv4 and IPv6 Access rules in Firewall, IPS, and Layer 2 Firewall Policies.

▼ To define source, destination, and service criteria

1. Right-click the **Source**, **Destination**, or **Service** cell and select **Edit Source**, **Edit Destination**, or **Edit Service**. The Definitions dialog for the selected cell opens.
2. Click **Add**. A new row is added to the list of matching criteria.

Note – All items on the same row must match the traffic for the row to match. You do not have to insert elements into all cells on the same row.

3. Drag and drop elements from the list on the left to the correct cell in the row. The following types of items can be used as matching criteria:

Table 40.1 Matching Criteria for Source and Destination Definitions

User	IP Address	Domain Name	Zone
User and User Group elements for users stored on an integrated Active Directory server with a User Identification Agent installed and configured. See Enabling Access Control by User (page 838) for more information.	Any element from the Network Elements branch that directly represents an IP address.	Domain Name elements. If DNS Server IP Addresses have been defined in the engine properties, the engine automatically resolves the Internet domain names to IP addresses. See Defining Domain Name Elements (page 723).	Zone elements for interface matching.

Note – VPN and NAT operations may change the routing of packets, potentially causing packets that no longer match the Destination Zone of an Access rule to be discarded. See Using Zones in the Destination of Access Rules (page 673) for more information.

Table 40.2 Matching Criteria for Service Definitions

URL Situation	Application	Service (Port)	TLS Match
URL Situation elements for web filtering. URL Situations must be used with a TLS Match element. We also recommend using an HTTP or HTTPS Service to improve inspection performance.	Application elements for application detection. See Getting Started With Applications (page 772) for more information. Applications are not supported on Virtual Firewalls.	TCP and UDP Service elements.	TLS Match elements for application detection. TLS Match elements must be used with a URL SItuation Element or with an Application element that contains a TLS Match. See Creating TLS Matches (page 773) for more information.

Note – You cannot use Application elements and Service elements on different rows of the Service Definition.

4. Repeat Step 2-Step 3 to add additional rows of matching criteria. Click **OK** when you are finished.

Adding Comments in Policies

You can add two types of comments in policies:

- The Comment cell in each rule allows you to write rule-specific comments, for example, to record why the rule was added. Note that the History in the Info panel shows when the rule was added and last changed and through which administrator account. Double-click the cell to edit the comment text.
- In rule tables, you can insert Rule Sections (through the right-click menu for any rule) to visually structure the policy under collapsible sections of rules that are preceded by a comment row. Double-click the row to edit the comment text. You can set each comment row's color through the **Colors** submenu in the comment row's right-click menu.

The maximum length of both types of comments is 249 characters.

Reading Rule Identifiers

Each rule has two (non-editable) identifiers:

- The **ID** cell shows the order of the rules. For example, the ID 14.1.2 shows that the rule is the second rule in the policy. It is in an insert point that is the first rule in the parent template. That insert point is the fourteenth rule in the top-level parent template. The number changes as you add, remove, and move rules.
- The **Tag** is the unique identifier of the rule in this policy. It contains a static part that does not change when rules are added, removed, or moved, and a changing part that indicated the version of the rule. For example, in Tag "@274.12", "274" is the unchanging part and "12" indicates that the rule is currently in its twelfth edit. The tag is used, for example, to provide links from logs to rules that created them.

Related Tasks

- ▶ Searching in Rules
- ▶ Finding Unused Rules in Firewall Policies (Hit Counters) (page 667)
- ▶ Adding Insert Points in Policy Templates (page 668)
- ▶ Editing Ethernet Rules (page 668)
- ▶ Editing Access Rules (page 672)
- ▶ Editing Inspection Policies (page 699)
- ▶ Editing Firewall NAT Rules (page 687)
- ▶ Validating Rules Automatically (page 714)

Searching in Rules

In rule tables, you can search rules based on most of the cells. Click the Tools icon in the Policy-specific toolbar and select **Search Rules** to display the search at the bottom of the rule table (Illustration 40.6).

In the Rules tree on the Inspection tab, you can search the Situations through type-ahead searching. When you type a part of the name of the Situation, the tree is filtered to contain only matching Situations (and their parent Situation Types, which make up the tree branches). The currently active type-ahead search is shown at the bottom of the tree panel.

The illustration below shows the rule search in rule tables.

Illustration 40.6 Rule Search for Rule Tables

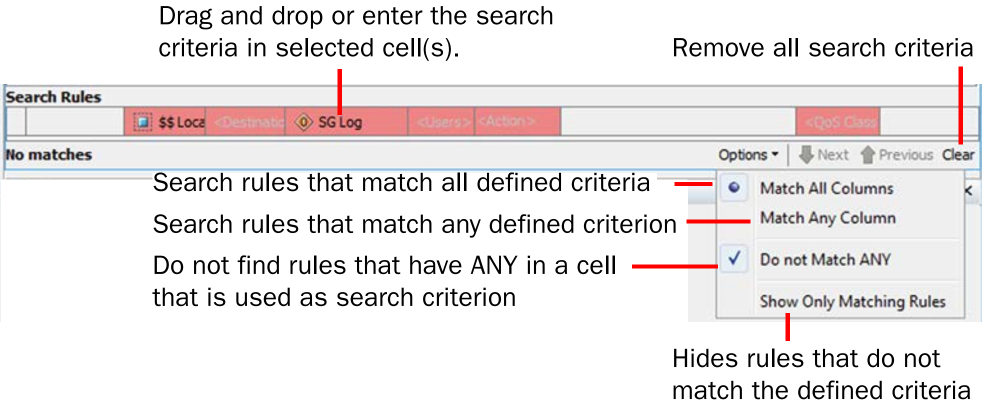

Drag and drop or enter the search criteria in selected cell(s).

Remove all search criteria

Search rules that match all defined criteria

Search rules that match any defined criterion

Do not find rules that have ANY in a cell that is used as search criterion

Hides rules that do not match the defined criteria

You can add values in different ways:

- Drag and drop elements from the rule table above, from different windows and tabs, or from the resource panel (shown in Edit mode).
- Right-click a cell and choose **Select** to browse for elements.
- In the **Source** and **Destination** search cells, you can manually type in IPv4 or IPv6 addresses, networks, or address ranges. Use standard notations (for example, 192.168.1.0/16, or 192.168.10.0 - 192.168.10.101 for IPv4 networks, or 2001:0db8:1234::/48, or 2001:0db8:1234:: - 2001:0db8:1234::100 for IPv6 networks.

Fill in the relevant cell(s). The cells you leave empty are ignored in the search. The first rule that matches your search is shown on a dark green background and all other matching rules are highlighted on a light green background. Click the **Next** or **Previous** arrow to move up or down from the currently selected rule to a matching rule above or below.

Related Tasks

- ▶ Finding Unused Rules in Firewall Policies (Hit Counters) (page 667)

Finding Unused Rules in Firewall Policies (Hit Counters)

Firewall Access rules and NAT rules contain a Hits cell that can show how many times each rule in your Firewall Policy has matched actual network traffic.

The main purpose of viewing the rule hits is to find valid rules that match traffic that the engine does not encounter in the network. This feature complements the rule validation checks, which can find rule design errors. Engines count rule hits automatically for all rules of supported types. The hits are stored as statistical counter data on the Log Server(s).

▼ To run a rule counter analysis

1. Open the policy for preview or editing and switch to the correct tab for the type of rules you want to examine.
2. Select an engine in the policy-specific toolbar.
3. Click the Tools icon in the policy-specific toolbar and select **Rule Counters**. The Rule Counter Analysis dialog opens.

Illustration 40.7 Rule Counter Analysis

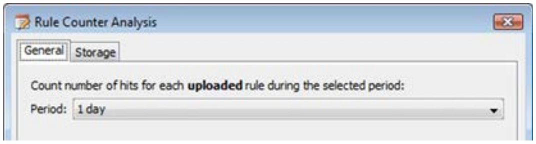

Count number of hits for each **uploaded** rule during the selected period:

Period: 1 day

4. Select the **Period** for which you want to check the rule matches; either one of the pre-set relative periods or **Custom** if you want to define the Period in detail.
5. (*Optional*) To select Management and/or Log Servers for this operation, or to include archived data, switch to the **Storage** tab and change the selection. Make sure you include the Log Servers and folders that contain data for the target engine and the period you selected.
6. Click **OK** to display the rule hits. The Hit information is displayed until you close the view.
   - The Hit cell in each rule is filled in with the number of connections that matched the rule during the chosen period.
   - If there is no statistical information about the rule with your selected criteria, the Hit cell shows “N/A” (for example, for rules added after the period you are analyzing).

Related Tasks

- ▶ Validating Rules Automatically (page 714)

Adding Insert Points in Policy Templates

Prerequisites: Creating a New Template Policy or a Policy

Insert Points mark the positions where rules can be added. When you edit a Template Policy, you must add at least one new yellow Insert Point on all tabs if you want the inheriting Policy or Policy Template to be editable on each tab. Green Insert Points are inherited insert points from the previous level, and they are not inherited further down in the hierarchy. They only show you where the higher-level template allows you to add rules and disappear as soon as you add a rule or a new (yellow) insert point in that position.

▼ To add an insert point in a Policy Template

1. Open the Policy template for editing.
2. Right-click the green insert point and select **Add Insert Point** or a rule that is editable in this Template and select **Rule→Add Insert Point Before** or **Add Insert Point After**.
3. Give the insert point a descriptive **Name** and click **OK**. An inheritable (yellow) insert point is added to the Template.

You can add as many insert points in the Template as your rule structure calls for.

Editing Ethernet Rules

Prerequisites: You must have privileges to edit the Policy element

Ethernet rules are used by IPS engines and Layer 2 Firewalls. These rules define whether Ethernet traffic is allowed to continue or whether the traffic is denied immediately. Inline interfaces of IPS engines and Layer 2 Firewalls can directly stop any traffic when the *Discard* action is used.

If you are using the IPS Template or the Layer 2 Firewall Template as the basis for your policies, the Ethernet rules direct all IPv4 and IPv6 traffic to the Inspection Policy for further inspection, and let ARP, RARP, and STP traffic through. You can use the first Insert Point in the template to make exceptions to this behavior for certain MAC addresses and/or Logical Interfaces. We recommend that you insert any other changes at the second insert point.

Make sure that your Ethernet rules direct IP traffic for inspection against Access rules (by applying the default IPv4 and IPv6 Services to traffic). When traffic does not match any Ethernet rule, the traffic is let through without further inspection.

▼ To add an Ethernet rule

1. Select **Configuration→Configuration→Security Engine**. The Security Engine Configuration view opens.
2. Expand the **Policies** branch and select **IPS Policies**.
3. Open a Template Policy or Policy for editing, and switch to the **Ethernet** tab.
4. Add the rule in one of the following ways:
   - Right-click the **ID** cell of an existing rule and select **Add Rule Before** or **Add Rule After**.
   - Copy and paste an existing rule.

5. Specify the portion of the traffic that you want to control and select what happens to traffic that matches the rule as explained in the table below.

Table 40.3 Matching Cells and Actions in Ethernet Rules

Cell	Option	Explanation
Logical Interface		If your engine has more than one Logical Interface defined, you can optionally add Logical interface elements in this cell to select which rules apply to which Logical Interfaces (network segments). The rules in the IPS Template and Layer 2 Firewall Template match any Logical Interface. Matches any Logical Interface by default.
Source		Match the rule to one or more MAC addresses. Does not match anything by default (whole rule is ignored).
Destination		
Service		Match the rule to an Ethernet Service. Does not match anything by default (whole rule is ignored).
Action (*Right-click to select*)	Allow	Matching traffic is allowed to pass. If the Service cell contains the default IPv4 or IPv6 Service element(s), the matching continues in the Access rules. Otherwise, the traffic is allowed to continue without further examination.
	Discard	If passing through inline interfaces, matching traffic is silently dropped. This action cannot be applied to traffic picked up through Capture interfaces.

Defining Logging Options for Ethernet Rules

Caution – Ethernet rules are matched to each packet. If logging is active for a rule, a new log entry is produced for each matching packet. Do not activate any logging for Ethernet rules that match sizeable portions of your network traffic.

▼ To create a log or alert entry when the rule matches

1. Double-click the **Logging** cell in the rule.

2. Define the options as explained in the table below.

Table 40.4 Logging Options in Ethernet Rules

Options		Explanation
Log Level	None	A matching packet creates no log entry. This is the recommended setting for most Ethernet rules.
	Transient	Each matching packet creates a log entry that is displayed in the Current Events mode in the Logs view (if someone is viewing it at the moment) but is not stored.
	Stored	Each matching packet creates a log entry that is stored on the Log Server.
	Essential	Creates a log entry that is shown in the Logs view and saved for further use. When the Log Server is unavailable, log entries are temporarily stored on the engine. When the engine is running out of space to store the log entries, it begins discarding log data in the order of importance. Monitoring data is discarded first, followed by log entries marked as Transient and Stored, and finally log entries marked as Essential. The Alert entries are the last log entries to be discarded. **Note!** The settings for storing the logs temporarily on the engine are defined in the log spooling policy. See Configuring Default Log Handling Settings (page 580).
	Alert	Each matching packet triggers the alert you add to the Alert field.
Alert		If Log Level is set to Alert, defines the Alert that is sent. Selecting different Alerts for different types of rules allows more fine-grained alert escalation policies.
Severity		If Log Level is set to Alert, allows you to override the severity defined in the Alert element.
Connection Closing	No log	No log entries are created when connections are closed.
	Normal log	Both connection opening and closing are logged, but no information is collected on the volume of traffic.
	Log Accounting Information	Both connection opening and closing are logged and information on the volume of traffic is collected. This option is not available for rules that issue Alerts. If you want to create reports that are based on traffic volume, you must select this option for all rules that allow traffic that you want to include in the reports.

Table 40.4 Logging Options in Ethernet Rules (Continued)

Options		Explanation
Recording	Log User Information	Stores information about Users when they are used as the Source or Destination of an Access rule. You must select this option if you want Users to be referenced by name in log entries, statistics, reports, and user monitoring. Otherwise, only the IP address associated with the User at the time the log was created is stored.
	Log Application Information	Stores information about Application use. You can log Application use even if you do not use Applications for access control. See Logging Application Use (page 776). Applications are not supported on Virtual Firewalls.
	Additional Payload	Stores packet payload extracted from the traffic. The collected payload provides information for some of the additional log fields listed in Log Fields Controlled by the Additional Payload Option (page 1252) depending on the type of traffic.

Defining a MAC Address for Ethernet Rules

The MAC Address element defines the MAC (Media Access Control) address of a network card. MAC Address elements are used to match a certain source or destination MAC address in the Ethernet rules.

▼ To define a MAC Address element

1. Select **Configuration→Configuration→Security Engine**. The Security Engine Configuration view opens.
2. Expand the **Other Elements** branch.
3. Right-click **MAC Addresses** and select **New MAC Address**.
4. **Name** the element.
5. Enter the **MAC Address**. You can enter any valid MAC address including, for example, the broadcast address (ff:ff:ff:ff:ff:ff).

Editing Access Rules

Prerequisites: You must have a custom Policy element and privileges to edit it

Access rules filter traffic by defining matching criteria and an action that is applied to packets that match all criteria defined in the rule. Access rules are used by Firewalls, IPS engines, Layer 2 Firewalls, Master Engines, and Virtual Firewalls:

- In Firewall and Layer 2 Firewall policies, the Access rules are the most important type of rules. The criteria you define in the Access rules determines which connections are allowed. By default, Firewall and Layer 2 Firewall Access rules stop traffic that you do not specifically allow.
- In IPS policies, Access rules can be used to optionally filter out some traffic and to exclude some traffic from further inspection. Only traffic on Inline interfaces can be filtered with Access rules. IPS engines allow all traffic that you do not specifically deny. If the policy is based on the IPS Template, all allowed traffic is inspected against the Inspection Policy by default.

▼ To add an Access rule

1. Select **Configuration→Configuration→Security Engine** from the menu. The Security Engine Configuration view opens.
2. Expand the **Policies** tree and select a type of policies (for example, **Firewall Policies**).
3. Open a Template Policy, Policy, or Sub-Policy for editing, and switch to the **IPv4 Access** tab or to the **IPv6 Access** tab.
4. Add the rule in one of the following ways:
   - Right-click the **ID** cell of an existing rule and select **Add Rule Before** or **Add Rule After**.
   - Copy and paste an existing rule.
5. Specify the portion of the traffic that you want to control as explained in Defining What Traffic an Access Rule Matches.
6. Define what happens to traffic that matches the rule as explained in Defining What Action an Access Rule Takes (page 675).
7. (*Optional*) Define options for triggering logs and alerts as explained in Defining Access Rule Logging Options (page 684).

Defining What Traffic an Access Rule Matches

IPv4 Access rules and IPv6 Access rules are both configured in the same way. Firewalls, IPS engines, Layer 2 Firewalls, Master Engines, and Virtual Firewalls use both types of Access rules.

In addition to more specific matching criteria, the matching cells can be set to two additional settings:

- ANY (available through right-click in a cell and selecting **Set to ANY**) matches all valid values for the cell, for example, all IPv4 addresses.
- NONE is the default value for mandatory traffic matching cells that have no matching criteria in them. As long as any cell in a rule contains NONE, the whole rule is invalid and is ignored.

Using Zones in the Destination of Access Rules

Due to the processing order of Access and NAT rules, the interface through which the packet will be sent out is not yet determined at the time Access and NAT rules are processed. During the matching against Access and NAT rules, the destination Zone is matched based on the current routing decision for the packet. NAT and VPN operations may change the route that is actually used when the packet is sent out. Because of this, the packet is checked against the Access rules again before being forwarded. If the changed destination Zone still matches, traffic is processed according to the original rule. If the changed destination Zone does not match the Access rule, the traffic is discarded. Carefully consider how the rules will be applied when using Zones in the Destination of Access rules when NAT and VPN operations may change the routing decision. See the *Firewall Reference Guide* for more information about how the engine processes the packets.

To define how an Access rule matches traffic, fill in the cells with elements as explained in the table below.

Table 40.5 Matching Cells in Access Rules

Cell	Explanation
Logical Interface (*IPS and Layer 2 Firewall only*)	If your engine has more than one Logical Interface defined, you can optionally add Logical interface elements in this cell to select which rules apply to which Logical Interfaces (network segments). The rules in the IPS Template and Layer 2 Firewall Template match any Logical Interface. For more information, see Defining Traffic Inspection Interfaces for IPS Engines (page 463). Matches any Logical Interface by default.
Source	A set of matching criteria that defines the IP addresses and interfaces that the rule matches. For more information, see Defining Source, Destination, and Service Criteria (page 664). Any elements in the Network Elements category, as well as User and User Group elements can be inserted into these cells. If you have both IPv4 and IPv6 networks, the elements must have the correct type of IP address for the type of rule or the element is ignored. Using User elements as the source or destination requires configuration of external Microsoft Active Directory server and a User Identifcation Agent. For more information, see Getting Started with Directory Servers (page 826). For more information about Network elements, see Getting Started with Defining IP Addresses (page 720). Does not match anything by default (whole rule is ignored).
Destination	

Table 40.5 Matching Cells in Access Rules (Continued)

Cell	Explanation
Service	A set of matching criteria that defines the network protocol or application the rule matches. The Service cell accepts Service and Service Group elements, URL Situations, Applications, and TLS matches. Applications are not supported on Virtual Firewalls. **Note!** You cannot insert a Service element and an Application element in the same Service cell. To use Service elements and Application elements together in the same rule, edit the Service Definition as explained in Defining Source, Destination, and Service Criteria (page 664). For more information, see Getting Started with Services (page 734). Does not match anything by default (whole rule is ignored).
Authentication (*Firewall and Virtual Firewall only, IPv4 only, Optional*)	If defined, the rule matches the specific Users or User Groups, and Authentication Methods you add to this cell. If the connection source is not currently authenticated, or the authentication is done using a method that is not included in this rule, the rule does not match and the matching continues to the next rule. Double-click the cell for more options, see Defining Firewall Access Rule Authentication Options (page 686). For more information, see Getting Started with Directory Servers (page 826) and Getting Started with User Authentication (page 856). Matches both authenticated and unauthenticated users by default (cell is empty).
Time (*Optional*)	Double-click to limit the rule's validity to a specific time period. During the specified time period, the rule matches. Outside the specified time period, the rule does not match and the matching continues to the next rule. The time is entered in the UTC time zone. For more information, see Limiting the Time when a Rule Is Active (page 713).
Source VPN (*Firewall and Virtual Firewall only, Optional*)	Matches the rule based on whether the traffic is received through a VPN. Double-click to specify that the rule matches only VPN traffic, only non-VPN traffic, only traffic from a specific VPN, or only IPsec VPN client traffic in any VPN (*IPv4 only*). This allows you, for example, to restrict the services VPN client users can access remotely when the IP addresses assigned to their laptops are the same both over the VPN and when connecting from within the local internal network. For more information, see Creating Rules for Policy-Based VPNs (page 958). Matches all traffic by default (cell is empty).

Defining What Action an Access Rule Takes

▼ To define what an Access rule does to matching traffic

➥ Right-click the Action cell and select the Action as explained in the table below.

Table 40.6 Access Rule Actions

Action	Options	Explanation
Allow	Yes	Traffic matching the rule is allowed to pass through the engine. The traffic may still be subject to various forms of deep packet inspection depending on the component and the options you set. Further options for traffic handling are available in the Action Options. See Defining Firewall Allow Action Options (page 678) or Defining IPS and Layer 2 Firewall Allow Action Options (page 683).
Continue	Yes	The options specified in a rule with this action are stored in memory while the matching process for the packet continues. The specified options are applied to any other rule that the same packet matches if the rules have no rule-specific definitions for the same options. Further options for traffic handling are available in the Action Options. See Defining Continue Action Options in Access Rules (page 681). For more information, see the *Firewall/VPN Reference Guide* or the *IPS and Layer 2 Firewall Reference Guide*.
Discard	Yes	The traffic is discarded without sending an ICMP error message or TCP reset to the source. Optionally, a response message can be shown to the end-user. See Defining Discard Action Options (page 677). This action can never be applied to traffic picked up through capture interfaces on an IPS engine or Layer 2 Firewall.
Refuse	Yes	The traffic is discarded and an ICMP error message or a TCP reset (for TCP connections) is sent in response to the source. Optionally, a response message can be shown to the end-user. See Defining Refuse Action Options (page 677). This action can never be applied to traffic picked up through Capture interfaces on an IPS engine or Layer 2 Firewall.
Jump	Yes	The processing is continued in a Sub-Policy (selected in the Action Options). If none of the rules in the Sub-Policy match, processing continues to the next rule in the main policy. Note that options that are set to specific values in a Jump rule override the corresponding settings in the sub-rules. See Defining Jump Action Options (page 678).
Use IPsec VPN (*Firewall and Virtual Firewall only*)	Yes	Includes a specific IPsec VPN configuration in the policy (selected in the Action options). Directs outgoing traffic into a specific VPN and/or allows incoming traffic from the specified VPN. See Defining Firewall Use VPN Action Options (page 682)

Table 40.6 Access Rule Actions (Continued)

Action	Options	Explanation
Apply Blacklist	Yes	The engine checks the traffic against its current blacklist entries. If the traffic matches a blacklist entry, it is discarded. If the traffic does not match any current entries on the blacklist, processing continues to the next rule in the main policy. There is only one blacklist for each Security Engine, but you can add several rules that reference this blacklist. You can limit the components that are allowed to add blacklist entries in the Action Options. See Defining Apply Blacklist Action Options.
Action Options	(N/A)	Opens a dialog that allows you to modify action-specific options if the selected action has options. See the explanations above for further information.

Defining Access Rule Action Options

Defining Apply Blacklist Action Options

The options for the Apply Blacklist action in Access rules affect the reception of blacklist entries on Security Engines. For information on setting up blacklisting and generating blacklist requests, see Getting Started with Blacklisting (page 816).

▼ To set the Action options for the Apply Blacklist action in Access rules

1. Right-click the **Action** cell in an IPv4 Access rule and select **Apply Blacklist**.
2. Double-click the **Action** cell. The action-specific options dialog opens.

Set the options as explained in the table below.

Table 40.7 Apply Blacklist Action Options

Tab	Option	Explanation
Blacklisting	Allowed blacklisters for the rule	Any: Blacklist entries are accepted from all components. Restricted: Blacklist entries are only accepted from the components you specify (and from the engine command line). Security Engines are always allowed to add entries to their own blacklists.
	Allowed Blacklisters	Add the specific components you want to allow to send blacklist requests.
Response	User Response *(HTTP only)*	Defines which automatic response is shown to the end-user when a connection is discarded. For more information about User Responses, see Getting Started with User Responses (page 778). User Responses are not supported on Virtual Firewalls.

Defining Discard Action Options

The Discard Action Options in Access rules allow you to define an automatic User Response to be shown to the user when an HTTP connection is discarded. User Responses are not supported on Virtual Firewalls.

▼ To set the Action options for the Discard action in Access rules

1. Right-click the **Action** cell in an Access rule and select **Discard**.
2. Double-click the **Action** cell. The action-specific options dialog opens.
3. Set the options as explained in the table below.

Table 40.8 Discard Action Options

Option	Explanation
Override Collected Values Set with “Continue” Rules	Select the option to activate the settings and override any options set in a preceding Continue rule that matches the same connections.
User Response (*HTTP only*)	Defines which automatic response is shown to the end-user when a connection is discarded. For more information about User Responses, see Getting Started with User Responses (page 778).

Defining Refuse Action Options

The Refuse Action Options in Access rules allow you to define an automatic User Response to be shown to the user when an HTTP connection is discarded. User Responses are not supported on Virtual Firewalls.

▼ To set the Action options for the Refuse action in Access rules

1. Right-click the **Action** cell in an Access rule and select **Refuse**.
2. Double-click the **Action** cell. The action-specific options dialog opens.
3. Set the options as explained in the table below.

Table 40.9 Discard Action Options

Option	Explanation
Override Collected Values Set with “Continue” Rules	Select the option to activate the settings and override any options set in a preceding Continue rule that matches the same connections.
User Response (*HTTP only*)	Defines which automatic response is shown to the end-user when a connection is discarded. For more information about User Responses, see Getting Started with User Responses (page 778). User Responses are not supported on Virtual Firewalls.

Defining Jump Action Options

The Jump action in Firewall, IPS, and Layer 2 Firewall IPv4 Access rules, and Firewall IPv6 Access rules adds a conditional section of rules to the policy. If a connection matches the Jump rule's matching criteria, it is matched against the rules in the Sub-Policy.

▼ To set the Action options for the Jump action in IPv4 Access rules

1. Right-click the **Action** cell in an IPv4 or IPv6 Access rule in a Firewall policy, or an IPv4 Access rule in an IPS policy and select **Jump**. The action-specific options dialog opens.
2. Set the options as explained in the table below.

Table 40.10 Jump Action Options

Tab	Options	Explanation
Jump	Sub-Policy	Select a Sub-Policy. Connections that match the Jump rule are matched against the selected Sub-Policy. If the Sub-Policy rules do not match, the processing is returned to the next rule in the main policy.
Connection Tracking (*Firewall and Virtual Firewall only*)	Standard connection tracking options	If you select specific connection tracking options in a Jump rule, the options you select here are used for all rules in the Sub-Policy regardless of the settings in those rules. See Defining Firewall Allow Action Options.

Defining Firewall Allow Action Options

The Connection Tracking Options control the following aspects of Firewall traffic handling:

- You can control stateful inspection by setting options for connection tracking, including idle timeouts, TCP segment size enforcement, and concurrent connection limits (DoS protection).
- (*License permitting*) For IPv4 traffic, you can activate the deep inspection to match traffic against the Inspection Policy. You can check the IPv4 traffic for viruses by setting deep inspection and anti-virus options. You can filter incoming IPv4 traffic for spam by activating deep inspection and anti-spam options.

Note – Anti-Virus and Anti-Spam are not supported on Virtual Firewalls.

▼ To set the Action options for the Allow action in firewall Access rules

1. Right-click the **Action** cell in a firewall Access rule and select **Allow**.
2. Double-click the **Action** cell. The action-specific options dialog opens.
3. Set the options as explained in the tables that follow.

Table 40.11 Firewall Allow Action Options

Option	Explanation
Override Connection Tracking Options Set with "Continue" Rules	Select the option to activate the Connection Tracking settings and override any Connection Tracking options set in a preceding Continue rule that matches the same connections.

Table 40.11 Firewall Allow Action Options (Continued)

Option	Explanation
On (Recommended)	Enables connection tracking. The firewall allows or discards packets according to the selected Connection Tracking mode. Reply packets are allowed as part of the allowed connection without an explicit Access rule. Protocols that use a dynamic port assignment must be allowed using a Service with the appropriate Protocol Agent for that protocol (in Access rules and NAT rules). The additional options available when connection tracking is on are explained in the next table.
Off	Disables connection tracking. The firewall operates as a simple packet filter, allowing packets based on their source, destination, and port used. You must add separate Access rules that explicitly allow the reply packets. NAT cannot be applied to traffic allowed without connection tracking. **Turn off logging** in the rule if you disable connection tracking. When connection tracking is off, a log entry is generated for each packet and this may considerably stress the firewall, network connections, and the Log Server.
Override Inspection Options Set with “Continue” Rules	Select the option to activate the Inspection settings and override any Inspection options set in a preceding Continue rule that matches the same connections.
Deep Inspection	Selects traffic that matches this rule for checking against the Inspection Policy referenced by this policy. Traffic is inspected as the Protocol that is attached to the Service element in this rule.
Anti-Virus (*IPv4 Only*)	Inspects matching traffic against the virus signature database. To use this option, anti-virus must be activated in the Firewall element’s properties as explained in Configuring Anti-Virus Settings (page 540) and the Service in the rule must include one of the Protocols supported for anti-virus inspection (HTTP, HTTPS, IMAP, POP3, or SMTP). Selecting this option also activates the Deep Inspection option. You can further adjust virus scanning in the Inspection Policy. See Defining Permit Action Options in Exception Rules (page 705) for details.
Anti-Spam (*IPv4 Only*)	Filters incoming spam e-mails. To use this option, anti-spam must be activated in the Firewall element’s properties as explained in Configuring Anti-Spam Settings (page 541) and the Service in the rule must include SMTP protocol. Selecting this option also activates Deep Inspection option. You can further adjust spam filtering in the Inspection Policy. See Defining Permit Action Options in Exception Rules (page 705).

Table 40.12 Additional Options When Connection Tracking Is Enabled

Option		Explanation
Connection Tracking Mode	Default	TCP traffic is handled as in the Normal mode or Strict mode depending on whether Strict Connection Tracking has been enabled in the firewall properties. ICMP and UDP traffic is handled as in the Normal mode.
	Normal	The firewall drops ICMP error messages related to connections that are not currently active in connection tracking (unless explicitly allowed by a rule in the policy). A valid, complete TCP handshake is required for TCP traffic. The firewall checks the traffic direction and the port parameters of UDP traffic. If the Service cell in the rule contains a Service that uses a Protocol Agent, the firewall also validates TCP and UDP traffic on the application layer. If a protocol violation occurs, the connection is dropped.
	Strict	Allows only TCP traffic that strictly adheres to the TCP standard as defined in RFC 793. The firewall also checks the sequence numbers of the packets in pre-connection establishment states and for RST and FIN packets, and drops packets that are out of sequence. If the Service cell in the rule contains a Service that uses a Protocol Agent, the firewall also validates the traffic on the application layer. If a protocol violation occurs, the connection is dropped.
	Loose	Allows some connection patterns and address translation operations that are not allowed in the Normal mode. Can be used, for example, if routing is asymmetric and cannot be corrected or if the use of dynamic routing protocols cause the firewall to receive non-standard traffic patterns.
Do not synchronize connections		Disables the synchronization of connection information between firewall cluster nodes. This reduces the traffic volume on the active heartbeat interface, but it also prevents transparent failover of connections to other nodes. This option is supported starting from Firewall/VPN version 5.2.1.
Idle Timeout		The timeout (in seconds) after which inactive connections are closed. This timeout concerns only idle connections. Connections are not cut because of timeouts while the hosts are still communicating. This setting overrides the default idle timeouts set in the Firewall properties. **Caution!** Do not set long timeouts for many connections. Each connection that is kept active consumes resources on the firewall. Setting excessive timeouts for a large number of connections may lead to serious performance problems. Generally, the idle timeout should not be more than a few minutes at most.

Table 40.12 Additional Options When Connection Tracking Is Enabled (Continued)

Option	Explanation
Enforce TCP MSS (*IPv4 Only*	Allows entering the Minimum and Maximum value for the MSS in bytes. Headers are not included in the MSS value; MSS concerns only the payload portion of the packet. In most cases, network equipment prefer to send packets at the Ethernet-standard MTU (maximum transmission unit) size of 1500 (including both payload and headers). Maximum value: If a TCP packet has an MSS value larger than the maximum, the firewall overwrites the packet's MSS with the maximum value you set here. Setting the maximum MSS size may be necessary to prevent fragmentation. Typically, the value you enter is lower than the standard Ethernet MTU (1500), with consideration for the packet headers that are added to the MSS. Minimum value: If a TCP packet has an MSS value smaller than the minimum you set here, the packet is dropped. The smaller the data content is, the less efficient the communications get due to the fixed-size headers. Limiting the minimum size may help alleviate certain types of network attacks. Typically, the minimum value you enter is smaller than the standard minimum MSS (536). If a TCP packet does not include a MSS value, the firewall does not add the MSS value to the packet, but enforces the minimum MSS.
Concurrent connection limit per source IP	Settings for DoS (denial of service) attack protection. Allows you to set a maximum limit for the number of total open connections from/to a single IP address at any one time. You can select between Discard (silent drop) and Refuse (with ICMP error message) as the Action that is applied to new connections if the limit is reached. These limits are enforced by rules that have their Action set to Allow, Continue, or Use IPsec VPN (all actions Apply/Enforce/Forward included). Be careful to apply the concurrent connection limits correctly for the types of communication that this rule handles to avoid cutting off connections unnecessarily.
Concurrent connection limit per destination IP	

Defining Continue Action Options in Access Rules

The specified options are applied to any other rule that the same packet matches if the rules have no rule-specific definitions for the same options.

▼ To set the Action options for the Continue action in firewall Access rules

1. Right-click the **Action** cell in an Access rule and select **Continue**.
2. Double-click the **Action** cell. The action-specific options dialog opens.
3. Set the options as explained in the table in Defining Firewall Allow Action Options (page 678) or Defining IPS and Layer 2 Firewall Allow Action Options (page 683).

Defining Firewall Use VPN Action Options

For more information on specific use cases, see Creating Rules for Policy-Based VPNs (page 958).

▼ To set the Action options for the Use VPN action in firewall Access rules

1. Right-click the **Action** cell in a firewall Access rule and select **Use IPsec VPN**. The action-specific options dialog opens.
2. Set the options as explained in the table below.

Table 40.13 Use IPsec VPN Action Options

Options	Explanation
Apply	*Incoming connections*: The traffic is allowed if it arrives through the specified VPN. Otherwise, the rule does not match and the matching process continues to the next rule. *Outgoing connections*: The traffic is sent through the specified VPN. If the connection is not allowed in the VPN configuration, it is discarded.
Enforce	*Incoming connections*: The traffic is allowed if the specified VPN is used. Otherwise, the connection is discarded. *Outgoing connections*: The traffic is sent through the specified VPN. If the traffic is not allowed in the VPN configuration, it is discarded.
Forward	*VPN traffic*: The engine forwards the traffic from one VPN to another. If the traffic is not allowed in the VPN configuration, it is discarded. For more information, see Redirecting Traffic Between VPN Tunnels (page 989). *Other traffic*: The traffic is sent through the specified VPN. If the traffic is not allowed in the VPN configuration, it is discarded.
The Selected IPsec VPN	The action is applied to a specific VPN.
$ Client-to-Gateway IPsec VPNs (*IPv4 only*)	The action is applied to IPsec VPN client traffic in any VPN.

3. (*Optional*) Switch to the *Connection Tracking* tab and set the options as explained in Defining Firewall Allow Action Options (page 678).

Defining IPS and Layer 2 Firewall Allow Action Options

You can activate/deactivate deep inspection for matching traffic in the IPS and Layer 2 Firewall Access rules. If deep inspection is disabled, the traffic is not checked against the Inspection Policy at all.

If you use the IPS Template or the Layer 2 Firewall Template as the basis for your policy, Deep Inspection is enabled by default for all supported protocols (with Continue rules), and can be disabled for a specific rule if necessary. Otherwise, make sure that your custom template policy directs all necessary Protocols to be inspected.

▼ To set the Action options for the Allow action in IPS Access rules

1. Right-click the **Action** cell in an IPS Access rule and select **Allow**.
2. Double-click the **Action** cell. The action-specific options dialog opens.
3. Set the options as explained in the table below.

Table 40.14 IPS Allow Action Options

Option	Explanation
Override Connection Options Set with “Continue” Rules	Select the option to activate the Connection settings and override any Connection options set in a preceding Continue rule that matches the same connections.
Concurrent connection limit per source IP	Settings for DoS (denial of service) attack protection. Allows you to set a maximum limit for the number of total open connections from/to a single IP address at any one time. You can select between Discard (silent drop) and Refuse (with ICMP error message) as the Action that is applied to new connections if the limit is reached. These limits are enforced by rules that have their Action set to Allow, Continue, or Use IPsec VPN (all actions Apply/Enforce/Forward included). Be careful to apply the concurrent connection limits correctly for the types of communication that this rule handles to avoid cutting off connections unnecessarily.
Concurrent connection limit per destination IP	
Override Inspection Options Set with “Continue” Rules	Select the option to activate the Inspection settings and override any Inspection options set in a preceding Continue rule that matches the same connections.
Deep Inspection	Selects traffic that matches this rule for checking against the Inspection Policy referenced by this policy. Traffic is inspected as the Protocol that is attached to the Service element in this rule.

Defining Access Rule Logging Options

Rules can create a log or alert entry each time they match. By default, logging options set in a previous rule with Continue as its action are used. If no such rule exists, Firewalls and Virtual Firewalls log the connections by default. IPS engines and Layer 2 Firewalls do not log the connections by default. Each individual rule can be set to override the default values.

Note – Log pruning may override the logging options by deleting any number of generated log entries when they are received at the Log Server.

▼ To set the Logging options for Access rules

1. Double-click the **Logging** cell. The logging options dialog opens.
2. Set the options as explained in the table below.

Table 40.15 Access Rule Logging Options

Option		Explanation
Override Collected Values Set with “Continue” Rules		Select the option to activate the settings and override any options set in a preceding Continue rule that matches the same connections.
Log Level	None	Does not create any log entry.
	Transient	Creates a log entry that is displayed in the Current Events mode in the Logs view (if someone is viewing it at the moment) but is not stored.
	Stored	Creates a log entry that is stored on the Log Server.
	Essential	Creates a log entry that is shown in the Logs view and saved for further use. When the Log Server is unavailable, log entries are temporarily stored on the engine. When the engine is running out of space to store the log entries, it begins discarding log data in the order of importance. Monitoring data is discarded first, followed by log entries marked as Transient and Stored, and finally log entries marked as Essential. The Alert entries are the last log entries to be discarded. **Note!** The settings for storing the logs temporarily on the engine are defined in the log spooling policy. See Configuring Default Log Handling Settings (page 580).
	Alert	Triggers the alert you add to the Alert field.
Alert		If Log Level is set to Alert, defines the Alert that is sent when the rule matches (the Default alert or a custom element created in your system). Selecting different Alerts for different types of rules allows more fine-grained alert escalation policies.
Severity		If Log Level is set to Alert, allows you to override the severity defined in the Alert element.

Table 40.15 Access Rule Logging Options (Continued)

Option		Explanation
Connection Closing	No log	No log entries are created when connections are closed.
	Normal log	Both connection opening and closing are logged, but no information is collected on the volume of traffic.
	Log Accounting Information	Both connection opening and closing are logged and information on the volume of traffic is collected. This option is not available for rules that issue Alerts. If you want to create reports that are based on traffic volume, you must select this option for all rules that allow traffic that you want to include in the reports.
Recording	Log User Information	Stores information about Users when they are used as the Source or Destination of an Access rule. You must select this option if you want Users to be referenced by name in log entries, statistics, reports, and user monitoring. Otherwise, only the IP address associated with the User at the time the log was created is stored.
	Log Application Information	Stores information about Application use. You can log Application use even if you do not use Applications for access control. See Logging Application Use (page 776). Applications are not supported on Virtual Firewalls.
	Additional Payload	Stores packet payload extracted from the traffic. The collected payload provides information for some of the additional log fields listed in Log Fields Controlled by the Additional Payload Option (page 1252) depending on the type of traffic.

Defining Firewall Access Rule Authentication Options

▼ To set the Authentication Options for Firewall IPv4 Access rules

1. Double-click the **Authentication** cell. The Authentication Parameters dialog opens.
2. Configure the settings on the Parameters tab as explained in the table below.

Table 40.16 Authentication Parameters

Option		Explanation
Method	Require Authentication	The users are required to authenticate themselves. The request is sent to the appropriate authentication service and evaluated.
Authorize	Connection	Authorization is granted for a single connection. Each connection requires that the users type in their credentials again. This option is useful for granting restricted access to a particular service that does not open many separate connections.
	Client IP	Permits connections from the client's IP address until reaching the set Timeout. Choose this option if you want to grant access to several services or a service that opens many connections during a single session.
	Timeout	Time until the Client IP authorization expires. The default Timeout is 3600 seconds.

3. Switch to the **Users** tab and select the Users and/or User Groups that this rule applies to.
4. Switch to the **Authentication Methods** tab and select the Authentication Methods to use.

Editing Firewall NAT Rules

Prerequisites: You must have a custom Policy element and privileges to edit it

Firewalls, Master Engines, and Virtual Firewalls can perform NAT (network address translation). NAT replaces the source and/or destination IP addresses in packets with other IP addresses. NAT rules are matched to allowed connections after Access rule matching. NAT is applied before a routing decision is made, so the address translation may affect how the traffic is routed. NAT can be applied to IPv4 and IPv6 traffic.

Available NAT Operations

You can define the following types of NAT:

- Static source NAT, typically used for translating the internal (“real”) IP address of an internal host to a different IP address in the external network.
- Static destination NAT, typically used for translating the public IP address of an internal host to the private IP address, so that the host (server) can receive new connections from external hosts. Allows IP address or port translation (PAT), or both.
- A combination of both of the above if you want to translate both the Source and Destination address in the same connection, for example, to allow internal hosts to access your organization’s public servers using the public IP addresses of both the client and the server.
- Dynamic source NAT, typically used to translate the internal IP addresses of several internal hosts to one or a few external IP addresses to hide the internal network structure from outsiders and to avoid acquiring a separate public IP address for each of the hosts.
- Dynamic destination NAT can be configured for IPv4 traffic as part of the Server Pool feature (not in NAT rules). See Inbound Traffic Management (page 617).

Guidelines

General guidelines for adding NAT rules:

- NAT rules only apply to connections that are handled statefully (Connection Tracking option is left selected in the Access rule that allows the connection).
- NAT rules are applied to whole connections; reverse NAT for reply packets is automatic, so you do not need to define rules for replies within a connection.
- The NAT rules are matched to connections with the same type of matching criteria as other types of rules. The first matching NAT rule is applied and any other NAT rules are ignored.
- To prevent a NAT rule from matching some connections, create a NAT rule for those connections that specifies no translation and place it above the rule that matches.
- By default, NAT rules are ignored for traffic that enters or leaves a VPN tunnel. To match such traffic to NAT rules, enable NAT in the Gateway element’s properties.
- Routing decisions are made after NAT, so keep in mind that translating the destination address may affect how the traffic is routed. If not included in existing definitions, you may need to add the translated addresses in the Routing view.
- If you install the Firewall Policy with the **Keep Previous Configuration Definitions** option selected, previous NAT rules are also be kept until all currently open connections that use those rules are closed. In some cases, the old and the new rules may conflict, preventing policy installation until the option is deactivated.

For more background information, see the *Firewall/VPN Reference Guide*.

Related Tasks

▶ NAT Rule Examples (page 695)

Adding a NAT Rule

▼ To define a NAT rule

1. Add the rule in one of the following ways:
   - Right-click the **ID** cell of an existing NAT rule and select **Add Rule Before** or **Add Rule After**.
   - Copy and paste some existing NAT rule.
   - Copy and paste an Access rule to match the rule to the same Source, Destination, and Service.
2. Match the rule to traffic as explained in Defining What Traffic a NAT Rule Matches.
3. Define the translation you want to apply as explained in the following sections:
   - Overwriting the Source Address in Packets (page 690).
   - Overwriting the Destination Address in Packets (page 693).

Defining What Traffic a NAT Rule Matches

NAT rules are matched based on IP addresses and services, but the each address translation operation places specific restrictions on what you can put in the cells (as explained in the table below). Consider the design of your NAT rules and Access rules separately, because attempting to match the different types of rules one-to-one is not effective in most cases, and not always possible.

To overwrite both the source and destination IP address in the same packet (for example, to achieve *hairpin NAT*), configure both address translations in the same NAT rule.

▼ To how a NAT rule matches traffic

➥ Fill in the cells as explained in the table below.

Table 40.17 Matching Cells in NAT Rules

Cell	Explanation
Source	A set of matching criteria that defines the IP addresses and interfaces that the rule matches. For more information, see Defining Source, Destination, and Service Criteria (page 664). A single continuous address space is required for static source IP address translation. ANY is invalid as a static address translation definition. Any combination of IP addresses is valid in dynamic source address translation. Only used for rule matching in Destination address translation rules, so any combination of IP addresses (including ANY) is valid in that role. Elements in IPv4 NAT rules must contain IPv4 addresses, and elements in IPv6 NAT rules must contain IPv6 addresses. Does not match anything by default (whole rule is ignored).

Table 40.17 Matching Cells in NAT Rules (Continued)

Cell	Explanation
Destination	A set of matching criteria that defines the IP addresses and interfaces that the rule matches. For more information, see Defining Source, Destination, and Service Criteria (page 664). A single continuous address space is required for destination IP address translation. ANY is invalid as an address translation definition. Only used for rule matching in Source address translation rules, so any combination of IP addresses (including ANY) is valid in that role. Elements in IPv4 NAT rules must contain IPv4 addresses, and elements in IPv6 NAT rules must contain IPv6 addresses. Does not match anything by default (whole rule is ignored).
Service	Defines the TCP or UDP port(s) that you want to translate for destination port translation. A single continuous port range is required for destination port translation. ANY is invalid as a port translation definition. Only used for rule matching in IP address translation rules, so any combination of ports (including ANY) is valid in that role. Protocols with complex connection models (such as FTP) require dynamic address allocation. To apply NAT to such protocols, use a Service that contains the correct Protocol Agent. The Protocol Agent translates the port or address information that is sent as packet payload between the client and the server (on the application layer). Does not match anything by default (whole rule is ignored).
Used On (*Optional*)	If the policy is shared by several engines, you can add Firewall, Master Engine, or Virtual Firewall element(s) in this cell to restrict the NAT rule's validity to the selected engine(s). By default, matches on all engines that receive the rule (cell is empty).

Overwriting the Source Address in Packets

▼ To translate source addresses

1. Double-click the **NAT** cell in the NAT rule. The Network Address Translation dialog opens.
2. Select the translation type as explained in the table below.

Table 40.18 NAT Rule Options - Source Translation Types

Translation Type	Explanation
None	Source addresses in matching connections are not translated. The packets are sent onwards with the source address intact.
Dynamic	Source addresses in matching connections are translated using a smaller pool of IP addresses than there are original source addresses included in the rule. Many hosts can use the same IP address, and the connections are distinguished by allocating a different TCP or UDP port for each connection. Also used for activating an Outbound Multi-Link configuration (*IPv4 only)*. Because ports are needed to keep track of connections, dynamic NAT only works with TCP and UDP protocols. If the protocol used in the communications is not transported on top of TCP or UDP, the communicating applications must encapsulate the packets in TCP or UDP (NAT traversal) to communicate through dynamic NAT.
Static	Source addresses in matching connections are translated using the same number of IP addresses as there are possible original source addresses. Each translated IP address corresponds to one original IP address.

3. If you selected an address translation operation, configure the additional options according to the type of operation:
   - Defining Static Source Translation Options (page 691).
   - Defining Dynamic Source Translation Options (page 692).

Defining Static Source Translation Options

Illustration 40.8 Network Address Translation - Source Tab With Static Translation Selected

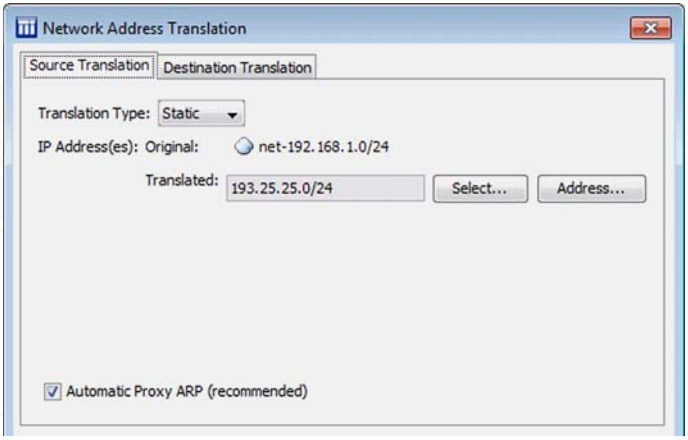

Translation Type: Static

IP Address(es): Original: net-192.168.1.0/24

Translated: 193.25.25.0/24

☑ Automatic Proxy ARP (recommended)

Table 40.19 NAT Rule Options - Static Source Translation Options

Option	Explanation
Original	The IP addresses you want to change with this address translation. These are defined in the Source cell of the NAT rule and shown here for your reference only; it is not possible to edit the Original addresses here.
Translated	The IP addresses you want the address translation to write in the packets. The Translated address space must have exactly the same number of IP addresses as there are in the Original address space, since each original address has a fixed pair in the translated address space. Click **Select** to use an existing network element to define the IP address(es). Click **Address** to manually type in the IP address or (sub)network you want to use for the address translation.
Automatic Proxy ARP (*IPv4 only*)	Allows the engine to answer address queries regarding the translated address(es). For this to work, the original IP address of all hosts whose IP address is translated must be included in the address definitions (for example, a Network element) under the correct interface in the Routing view. This option is required in most cases, but it must not be active for IP addresses that are used by any equipment in the directly connected networks.
Automatic Proxy Neighbor Discovery *(IPv6 only)*	Allows the engine to answer address queries regarding the translated address(es). For this to work, the original IP address of all hosts whose IP address is translated must be included in the address definitions (for example, a Network element) under the correct interface in the Routing view. There is a limit to the number of addresses that the engine can proxy for neighbor discovery.

Defining Dynamic Source Translation Options

Illustration 40.9 Network Address Translation - Source Tab With Dynamic Translation Selected

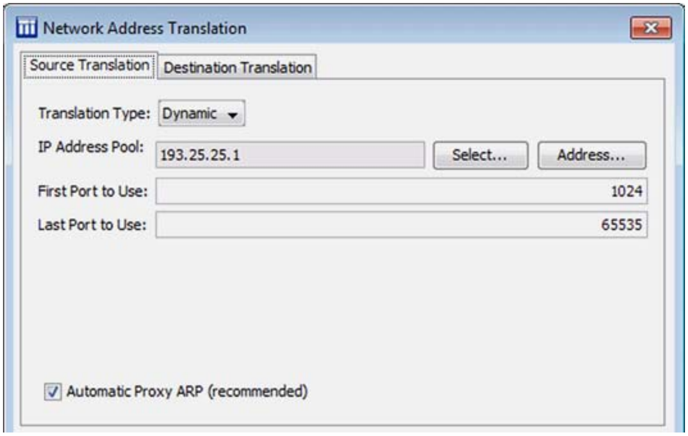

Network Address Translation

Source Translation | Destination Translation

Translation Type: Dynamic

IP Address Pool: 193.25.25.1

First Port to Use: 1024

Last Port to Use: 65535

☑ Automatic Proxy ARP (recommended)

Table 40.20 NAT Rule Options - Dynamic Source Translation Options

Option	Explanation
IP Address Pool	The IP address pool of IP address(es) that are used for the translation. The minimum size for the pool is one IP address. The number of IP addresses required depends on how many ports you allow the address translation to use and how many concurrent connections are handled by dynamic address translation at peak times. If the IP address/port pairs run out, new connections cannot be opened before existing connections are closed. The IP addresses used for NAT must not be in use in the network, as this will create an IP address conflict. However, the engine's own IP address (CVI on clusters) can be used for address translation if there are no free IP addresses available (make sure that your selected port range does not overlap with communications ports that the engine uses on this address). Click **Select** to use an existing network element to define the IP address(es). To use Multi-Link features for this traffic, select an Outbound Multi-Link element. Click **Address** to manually type in the IP address or (sub)network you want to use for the address translation.
First Port to Use	The start of the port range for source IP address translation. The default is the beginning of the "free" high port range, 1024.
Last Port to Use	The end of the port range for source IP address translation. The default is the highest possible port 65535.
Automatic Proxy ARP (*IPv4 only*)	Allows the engine to answer address queries regarding the translated address(es). For this to work, the original IP address of all hosts whose IP address is translated must be included in the address definitions (for example, a Network element) under the correct interface in the Routing view. This option is required in most cases, but it must not be active for IP addresses that are used by any equipment in the directly connected networks.

Table 40.20 NAT Rule Options - Dynamic Source Translation Options (Continued)

Option	Explanation
Automatic Proxy Neighbor Discovery *(IPv6 only)*	Allows the engine to answer address queries regarding the translated address(es). For this to work, the original IP address of all hosts whose IP address is translated must be included in the address definitions (for example, a Network element) under the correct interface in the Routing view. There is a limit to the number of addresses that the engine can proxy for neighbor discovery.

Overwriting the Destination Address in Packets

Note – Destination translation may change the routing of packets, potentially causing packets that no longer match the Destination Zone of an Access rule to be discarded. See Using Zones in the Destination of Access Rules (page 673) for more information.

▼ To translate Destination addresses

1. Double-click the **NAT** cell in the NAT rule. The Network Address Translation dialog opens.

Illustration 40.10 Network Address Translation - Destination Tab

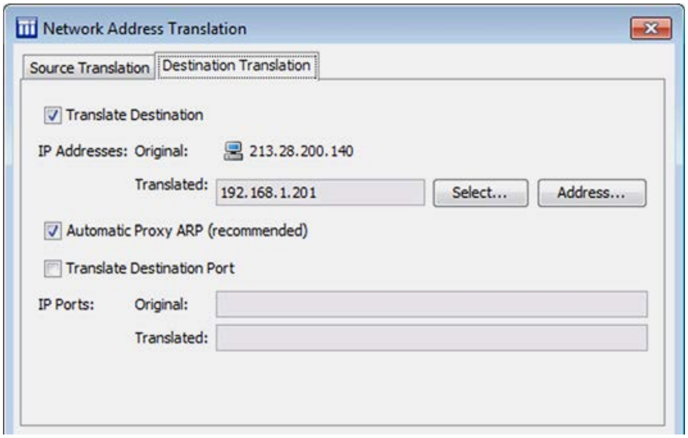

☑ Translate Destination

IP Addresses: Original: 213.28.200.140

Translated: 192.168.1.201

☑ Automatic Proxy ARP (recommended)

☐ Translate Destination Port

IP Ports: Original:

Translated:

2. Select the translation type as explained in the table below:

Table 40.21 NAT Rule Options - Destination Tab

Option	Explanation
Translate Destination *(Optional)*	Select this option if you want to translate destination IP addresses. If you do not select this option, IP addresses are not translated, so packets are sent onwards with the destination address intact.

Table 40.21 NAT Rule Options - Destination Tab (Continued)

Option		Explanation
IP addresses	Original	The IP addresses you want to change with this address translation. These are defined in the Destination cell of the NAT rule and shown here for your reference only; it is not possible to edit the Original addresses here.
	Translated	The IP addresses you want the address translation to write in the packets. The Translated address space must have exactly the same number of IP addresses as there are in the Original address space, since each original address has a fixed pair in the translated address space. Click **Select** to use an existing network element to define the IP address(es). Click **Address** to manually type in the IP address or (sub)network you want to use for the address translation.
Automatic Proxy ARP (*IPv4 only*)		Allows the engine to answer address queries regarding the translated address(es). For this to work, the original IP address of all hosts whose IP address is translated must be included in the address definitions (for example, a Network element) under the correct interface in the Routing view. This option is required in most cases, but it must not be active for IP addresses that are used by any equipment in the directly connected networks.
Automatic Proxy Neighbor Discovery *(IPv6 only)*		Allows the engine to answer address queries regarding the translated address(es). For this to work, the original IP address of all hosts whose IP address is translated must be included in the address definitions (for example, a Network element) under the correct interface in the Routing view. There is a limit to the number of addresses that the engine can proxy for neighbor discovery.
Translate Destination Port		Select if you want to translate destination port(s). If you do not select this option, ports are not translated, so packets are sent onwards with the destination port intact.
IP ports	Original	The port(s) you want to change with this address translation. These are defined in the Service element in the Service cell of the NAT rule and shown here for your reference only; it is not possible to edit the Original port(s) here.
	Translated	The port or port range you want the address translation to write in the packets. If you enter a port range, it must have exactly the same number of ports as there are in the Original ports, since each original port has a fixed pair in the translated address space (for example, 1-1023 could be translated to 50001-51023).

NAT Rule Examples

Example of a Static Source Translation Rule

This example shows a static address translation that translates the addresses in one network to IP addresses in another network. In this example, the NAT is done to access a particular server on the Internet.

Illustration 40.11 Example Scenario

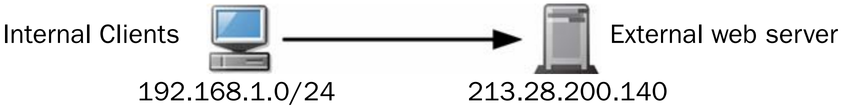

Table 40.22 Example NAT Rule Matching Cells

Source	Destination	Service
192.168.1.0/24	213.28.200.140	HTTP

Illustration 40.12 Example NAT Settings

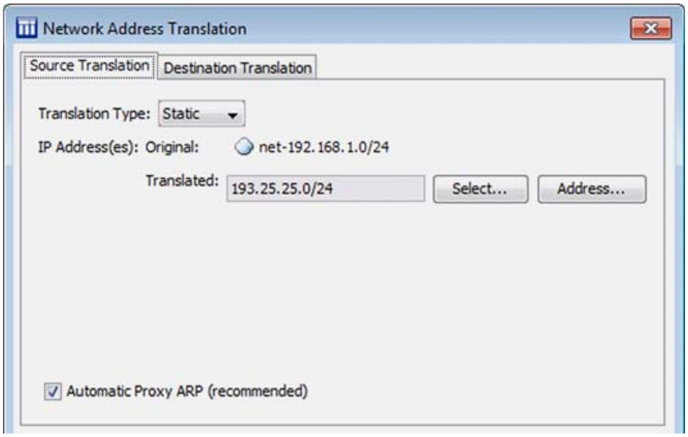

Translation Type: Static

IP Address(es): Original: net-192.168.1.0/24

Translated: 193.25.25.0/24

☑ Automatic Proxy ARP (recommended)

In static address translation using whole networks, each original source IP address has a static translated pair. For example here, the host 192.168.1.6 is always translated to 193.25.25.6 and host 192.168.1.11 is always translated to 193.25.25.11.

Example of a Dynamic Source Translation Rule

This example shows a dynamic address translation that translates the addresses in one internal network to a single external address for the purpose of general web browsing.

Illustration 40.13 Example Scenario

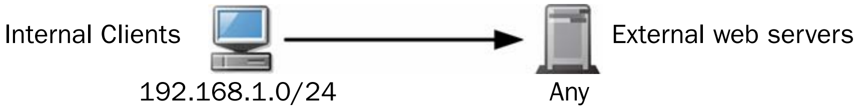

Table 40.23 Example NAT Rule Matching Cells

Source	Destination	Service
192.168.1.0/24	ANY	HTTP

Illustration 40.14 Example NAT Settings

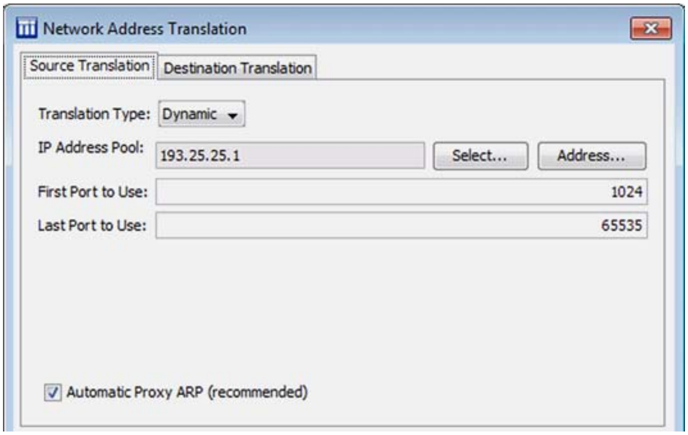

Translation Type: Dynamic

IP Address Pool: 193.25.25.1

First Port to Use: 1024

Last Port to Use: 65535

☑ Automatic Proxy ARP (recommended)

In dynamic address translation, several source IP addresses are translated using a smaller pool of translated addresses with the help of port translation. Each client connection uses a different port on an IP addresses that is shared between several different connections. Since each client reserves a port, the maximum number of simultaneous connections for a dynamic NAT operation can be calculated by multiplying the number of IP addresses by the number of ports in the range. Every port and IP address pair must be free from any other use (duplicate connections cannot successfully cross the firewall).

Example of a Destination Translation Rule

This example shows a static address translation that translates the external IP address (213.28.200.140) of a web server to the server's internal address (192.168.1.201).

Illustration 40.15 Example Scenario

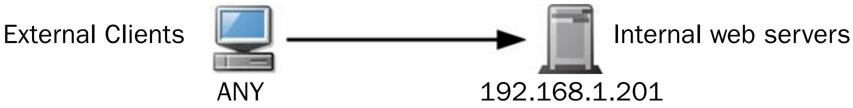

Table 40.24 Example NAT Rule Matching Cells

Source	Destination	Service
ANY	213.28.200.140	HTTP

Illustration 40.16 Example NAT Settings

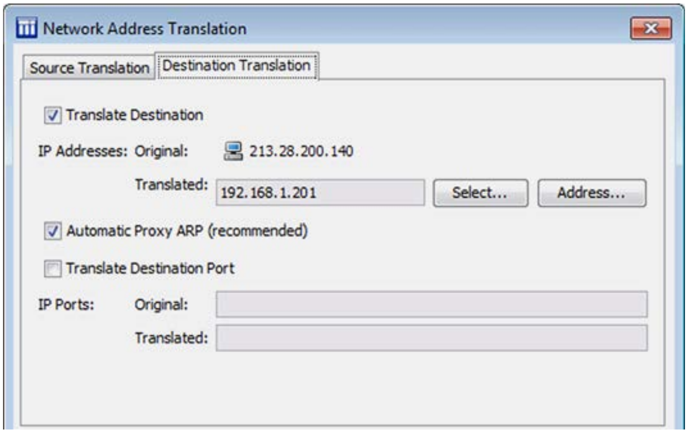

☑ Translate Destination

IP Addresses: Original: 213.28.200.140

Translated: 192.168.1.201

☑ Automatic Proxy ARP (recommended)

□ Translate Destination Port

IP Ports: Original:

Translated:

Example of a Combined Source And Destination Translation Rule

In this example, *hairpin NAT* is configured. Clients in the internal network (192.168.1.0/24) contact the organization's own public web server using the public IP address (213.28.200.140). The server's external address is translated to an internal address (192.168.1.201) that belongs to the same internal network address space as the contacting clients. Source address translation is used to prevent the server replies to the client's original IP address, since such replies would be routed directly within the local network instead of through the firewall, and the connections do not work without the reverse NAT that the firewall provides.

Illustration 40.17 Example Scenario

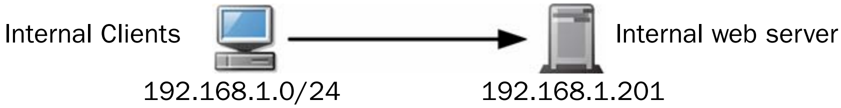

Table 40.25 Example NAT Rule Matching Cells

Source	Destination	Service
192.168.1.0/24	213.28.200.140	HTTP

Illustration 40.18 Example NAT Settings

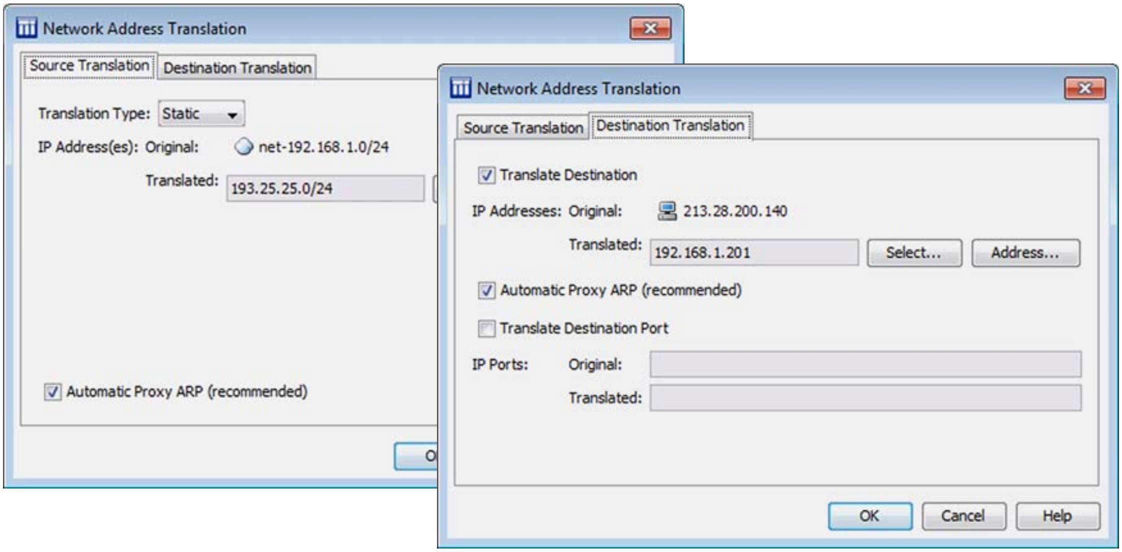
Network Address Translation

Source Translation | Destination Translation

Translation Type: Static

IP Address(es): Original: net-192.168.1.0/24

Translated: 193.25.25.0/24

☑ Automatic Proxy ARP (recommended)

Network Address Translation

Source Translation | Destination Translation

☑ Translate Destination

IP Addresses: Original: 213.28.200.140

Translated: 192.168.1.201

☑ Automatic Proxy ARP (recommended)

☐ Translate Destination Port

IP Ports: Original:

Translated:

The NAT settings on each tab are not any different than when you apply only source translation or only destination translation to matching connections. The key is that both definitions must be defined in the same NAT rule, because none of the other NAT rules are considered after the first match is found.

Editing Inspection Policies

Prerequisites: You must have a custom Policy element and privileges to edit it

Inspection policies activate checks for specific traffic patterns and define what action the system takes when a match is found. Inspection policies examine the packet payload throughout whole connections, and take action at the point when something threatening is discovered. Inspection policies are applied to connections that are selected for deep inspection in the Access rules. The IPS Template and Layer 2 Firewall Template enable deep inspection for all IP traffic. Deep inspection is not automatically enabled in the Firewall Template.

Security Engines examine IPv4 and IPv6 traffic against criteria defined in *Situation* elements. Security Engines and Log Servers process the detected events using *Correlation Situation* criteria.

Virtual Security Engines do not individually inspect traffic. One shared inspection process running on the Master Engine handles the inspection and correlation for all Virtual Security Engines associated with the Master Engine. To prevent excessive resource consumption on the Master Engine, take care when configuring Inspection policies for use on Virtual Security Engines.

You can optionally import Snort rules libraries to convert them into Inspection Policies. See Importing Snort Rules Libraries (page 709).

▼ To add an Inspection rule

1. Select **Configuration→Configuration→Security Engine** from the menu. The Security Engine Configuration view opens.
2. Browse to **Policies→Inspection Policies**.
3. Right-click a Template Inspection Policy or Inspection Policy element and select **Edit Inspection Policy**. The policy opens for editing on the Inspection tab.
4. Adjust the Rules that are applied to the majority of traffic. See Modifying the Inspection Rules Tree.
5. (*Optional*) Switch to the **Exceptions** tab and define detailed Exceptions to the main Rules. The main uses for exceptions are to eliminate false positives and to activate blacklisting or User Responses for specific traffic patterns. See Adding Exceptions to the Inspection Policy (page 702).

Modifying the Inspection Rules Tree

The Rules tree on the Inspection tab in Inspection Policies allows you to define what kind of action the system takes when Situation matches are found in traffic and how the match is logged. The Rules tree s the main tool for controlling the deep packet inspection. To edit these rules, you simply click and select an Action or Logging option. The definitions on the Exceptions tab are matched before the Rules tree on the Inspection tab.

In the tree, items that have subitems are Situation Type elements. The items that have no subitems are individual Situation and Correlation Situation elements. The Rules tree contains all Situation Types and Situations associated with them.

All levels of the tree are editable. Subitems inherit their values from their parent item by default. Any setting that is different from the default is considered an *override*. If you change a value in an item that has subitems, all subitems that are set to use the default value inherit this change. Any subitems that are set to an override continue to use that override.

Example **The parent item and ten of the subitems are set to the "Permit (Default)" action. Two of the subitems are set to the "Permit" action. You change the parent to the "Terminate" action. Ten subitems change to "Terminate (Default)". Two subitems continue to use "Permit".**

The illustration below shows how overrides are highlighted in the tree.

Illustration 40.19 Default Values and Overrides in the Tree

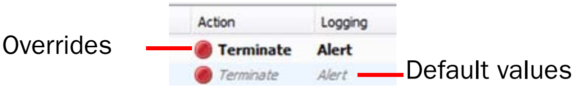

In lists of options, the default setting is also labeled, for example, "Permit (default)".

Regardless of the settings in the Rules tree in a higher-level template, it is still possible to change any Rules tree values in the inheriting policy. To add template rules that cannot be modified in the inheriting Policies, add the rules as Exceptions.

Table 40.26 Inspection Policy - Rules Panel - Action Column

Option	Explanation
Do Not Inspect	The definition is ignored. If the included Situations are not referred in the Exceptions either, the traffic patterns contained in those Situations are not transferred to the engines and therefore not matched to network traffic. This reduces the engine workload. This action is only available at the top level of the tree.
Permit	The connections are still allowed to continue even when an included pattern is found in traffic. This Action is useful if you want to log or record connections without stopping them. The inspection continues, so other Situations can still match the same connection and carry out a different action.

Table 40.26 Inspection Policy - Rules Panel - Action Column (Continued)

Option	Explanation
Terminate	On firewalls and Virtual Firewalls, stops the connection. On IPS engines and Layer 2 Firewalls, stops the connection if the traffic flows through an Inline interface pair. Traffic that is picked up through a Capture interface can be terminated by closing the connection through a Reset interface. When this action is triggered, the matching traffic is not inspected any further. IPS engines and Layer 2 Firewalls have a passive termination mode for testing purposes in which the connections are not stopped, but create an easily distinguishable log entry instead. This mode can be set as a default for the IPS engine or Layer 2 Firewall in the engine element or rule by rule in the options of Exceptions.

Table 40.27 Inspection Policy - Rules Panel - Logging Column

Option	Explanation
None	Does not create any log entry.
Transient	Creates a log entry that is displayed in the Current Events mode in the Logs view (if someone is viewing it at the moment) but is not stored.
Stored	Creates a log entry that is stored on the Log Server.
Essential	Creates a log entry that is shown in the Logs view and saved for further use. When the Log Server is unavailable, log entries are temporarily stored on the engine. When the engine is running out of space to store the log entries, it begins discarding log data in the order of importance. Monitoring data is discarded first, followed by log entries marked as Transient and Stored, and finally log entries marked as Essential. The Alert entries are the last log entries to be discarded. **Note!** The settings for storing the logs temporarily on the engine are defined in the log spooling policy. See Configuring Default Log Handling Settings (page 580).
Alert	Triggers an alert. By default, the Default Alert is generated, but you can select a different Alert through the Logging Options item.
Logging Options	Allows you to set up traffic recording and select a custom Alert. See Defining Logging Options for Inspection Rules and Exceptions (page 708).

Related Tasks

- Adding Situations to the Rules Tree (page 702)
- Removing Overrides From the Rules Tree (page 702)
- Adding Exceptions to the Inspection Policy (page 702)

Adding Situations to the Rules Tree

▼ To add your custom Situations to the tree

↳ Select the most appropriate Situation Type (these are all already included at some level of the tree) as the **Situation Type** in the properties of a custom Situation element.

Related Tasks

▶ Creating New Situation Elements (page 757)

Removing Overrides From the Rules Tree

▼ To remove overrides

↳ Right-click an item in the tree and select **Reset Branch**. The item you right-clicked and all of its subitems are reset to the Default value.

Adding Exceptions to the Inspection Policy

Inspection Exceptions allow you to make changes that are not applied to all connections and set some options (using the Continue action) for Exceptions and Rules that are processed later. Exceptions also contain some additional options that are not available in the Rules tree.

- You can match specific connections based on the IP addresses of the communicating hosts, the Service used, and the Logical Interfaces of IPS engines and Layer 2 Firewalls. For example, an Exception is needed to eliminate a false positive specifically in traffic between two internal hosts without disabling the inspection check completely.
- You can set additional responses to matches that are found. You can blacklist connections on a Security Engine, and you can add User Responses as notifications to some types of events.

▼ To define an Exception rule in the Inspection Policy

1. Add the rule in one of the following ways:
   - Right-click a generated log entry and select one of the options in the **Create Rule** submenu. The rule is created as an Exception with matching details from the log entry.
   - Right-click the **ID** cell in an existing Exception and select **Add Rule Before** or **Add Rule After**.
   - Copy and paste an existing Exception rule.
   - Copy and paste a rule from the Rules tree to the Exceptions table to match the same Situation(s) and options.
   - Copy and paste an Access rule to match the rule to the same Source, Destination, and Service.
2. Match the rule to traffic as explained in Defining What Traffic an Inspection Exception Rule Matches.
3. Define the exception you want to apply as explained in Defining What Action an Exception Rule Takes (page 704).
4. (*Optional*) Define options for triggering logs and alerts as explained in Defining Logging Options for Inspection Rules and Exceptions (page 708).

Related Tasks

▶ Importing Snort Rules Libraries (page 709)

Defining What Traffic an Inspection Exception Rule Matches

Inspection Exceptions are matched based on the patterns defined in Situation elements. The traffic is checked against all patterns in the policy, and when a match is found, the Situation element is used in looking up the rule that determines what happens to the traffic. If none of the Exceptions match, the matching continues in the Rules tree.

▼ To define how an Inspection rule matches traffic

➥ Fill in the cells as explained in the table below. The **Time** cell is optional; all other cells must always contain a value.

Table 40.28 Matching Cells in Inspection Exceptions

Cell	Explanation
Situation	The traffic patterns that you want the rule to match. The Situation cell accepts Situation, Situation Type, Tag, and Vulnerability elements. Situation Types, Tags, and Vulnerabilities are shown as branches in the Situations tree. Under the branches, you can view the individual Situations. Each Situation is usually listed more than once because there are alternative ways to categorize the Situations.
Severity	Matches the rule to only Situations with the specified Severity value(s). For example, if your rule is general and matches a wide range of Situations, you may want to create two similar rules: one for less severe Situations and one for more Severe situations. Useful in rules that contain Tags in the Situation cell.
Logical Interface (*IPS and Layer 2 Firewall only*)	If your engine has more than one Logical Interface defined, you can optionally add Logical Interface elements in this cell to select which rules apply to which Logical Interfaces (network segments). The rules in the IPS Template and Layer 2 Firewall Template match any Logical Interface. For more information, see Defining Traffic Inspection Interfaces for IPS Engines (page 463). Matches any Logical Interface by default.
Source	A set of matching criteria that defines the IP addresses and interfaces that the rule matches. For more information, see Defining Source, Destination, and Service Criteria (page 664). Any elements in the Network Elements category can be inserted into these cells. Both IPv4 and IPv6 addresses are valid in these cells. For more information, see Getting Started with Defining IP Addresses (page 720). Additionally, User and User Group elements can be inserted into these cells in the Inspection Policy. Using User elements as the source or destination requires configuration of external Microsoft Active Director server and a User Identifcation Agent. For more information, see Getting Started with Directory Servers (page 826). Does not match anything by default (whole rule is ignored).
Destination	
Protocol	You can optionally restrict the scope of Inspection rules by using the communications protocol as a matching criteria. The communications protocol information is attached to connections in the Access rules (Service cell). Useful in rules that contain Tags in the Situation cell.

Table 40.28 Matching Cells in Inspection Exceptions (Continued)

Cell	Explanation
Time (*Optional*)	Limits the rule's validity to the specified time period. During the specified time period, the rule matches. Outside the specified time period, the rule does not match and the matching continues to the next rule. The time is entered in UTC. For more information, see Limiting the Time when a Rule Is Active (page 713).

Defining What Action an Exception Rule Takes

▼ To define what an Exception rule does to matching traffic

➥ Right-click the **Action** cell and select one of the options explained in the table below.

Table 40.29 Inspection Exception Actions

Action	Explanation
Permit	The included Situations are disabled for the matching connections and the inspection continues. The Situations included in the Permit rule are ignored if they are included in further rules. Other Situations can still match the same connection and carry out a different action. Further options for traffic handling are available in the Action Options. See Defining Permit Action Options in Exception Rules (page 705).
Continue	The options specified in a rule with this action are stored in memory while the matching process continues. The specified options are applied to any other rule that the same traffic matches if the rules have no rule-specific definitions for the same options. All settings in the Options cell and the settings for the Terminate action in the Action cell can be set using Continue rules for further Exceptions and the Rules tree. However, the Rules tree ignores any logging options set with Continue rules, since the tree has a separate inheritance system for those options. Further options for traffic handling are available in the Action Options. See Defining Continue Action Options in Exception Rules (page 705).
Terminate	On Firewalls and Virtual Firewalls, stops the connection. On IPS engines and Layer 2 Firewalls, stops the connection if the traffic flows through an Inline interface pair. Traffic that is picked up through a Capture interface can be terminated by closing the connection through a Reset interface. When this action is triggered, the matching traffic is not inspected any further. IPS engines and Layer 2 Firewalls have a passive termination mode for testing purposes in which the connections are not stopped, but create an easily distinguishable log entry instead. This mode can be set for the whole engine or in the options of Exceptions. Further options for traffic handling are available in the Action Options. See Defining Terminate Action Options in Exception Rules (page 706).
Action Options	Opens a dialog that allows you to modify action-specific options. See the explanations above for further information.

Defining Continue Action Options in Exception Rules

The Continue action can set options for the Permit and Terminate actions in subsequent rules.

▼ To set the Action options for the Continue action in Exception rules

1. Right-click the **Action** cell and select **Continue**.
2. Double-click the **Action** cell. The action-specific options dialog opens.
3. Set the options as explained in Defining Permit Action Options in Exception Rules (page 705), and Defining Terminate Action Options in Exception Rules (page 706).

Defining Permit Action Options in Exception Rules

On the Firewall, the options for the Permit action in Exception rules allow you to control the inspection options in further detail and set a User Response for virus scanning or Situation matches.

Note – Anti-Virus, Anti-Spam, and User Responses are not supported on Virtual Firewalls.

On the IPS and the Layer 2 Firewall, the options for the Permit action in Exception rules allow you to blacklist traffic.

▼ To set the Action options for the Permit action in Exception rules

1. Right-click the **Action** cell and select **Permit**.
2. Double-click the **Action** cell. The action-specific options dialog opens.
3. Set the options as explained in the table below.

Table 40.30 Exception Rules - Permit Action Options

Tab	Option	Explanation
Anti-Virus (*Firewall only*)	Override collected values set with “Continue” rules	Select the option to activate the settings on this tab and override any options set in a preceding Continue rule that matches the same connections.
	Anti-Virus	Defines if traffic that matches the included Situation(s) is scanned for viruses. For example, you can deactivate virus scanning for a specific URL. This option only affects traffic that is already selected for virus scanning in the Access rules. By default, all traffic that is selected for virus scanning in the Access rules is scanned unless you make a specific exception.
Anti-Spam (*Firewall only*)	Override collected values set with “Continue” rules	Select the option to activate the settings on this tab and override any options set in a preceding Continue rule that matches the same connections.
	Anti-Spam	Defines if traffic that matches the included Situation(s) is filtered for spam. For example, you can deactivate spam filtering for traffic between your mail servers. This option only affects traffic that is already selected for spam filtering in the Access rules. By default, all traffic that is selected for spam filtering in the Access rules is filtered unless you make a specific exception.

Table 40.30 Exception Rules - Permit Action Options (Continued)

Tab	Option	Explanation
Blacklist Scope	Override collected values set with “Continue” rules	Select the option to activate the settings on this tab and override any options set in a preceding Continue rule that matches the same connections.
	Terminate the single connection	Creates entries that stop matching current connections, but which are not stored for any time.
	Block traffic between endpoints	Creates entries that stop matching connections and block traffic between the matching IP addresses for the set duration.
	Blacklist executors	Security Engines to which the blacklist requests are sent.

Defining Terminate Action Options in Exception Rules

The Terminate action options control connection termination, notifications, and the creation of blacklist entries that can be sent from one Security Engine to other Security Engines to block traffic dynamically for a specified time period.

Note – Virtual Security Engines cannot send blacklist requests to other Virtual Security Engines.

▼ To set the Action options for the Terminate action in Exception Rules

1. Right-click the **Action** cell and select **Terminate**.
2. Double-click the **Action** cell. The action-specific options dialog opens.
3. Set the options as explained in the tables below.

Table 40.31 Exception Rules - Terminate Action Options - Terminate Tab

Option	Explanation
Override collected values set with “Continue” rules	Select the option to activate the settings on this tab and override any options set in a preceding Continue rule that matches the same connections.
Terminates the connection	If you select **Yes**, traffic is stopped. If you select **No**, a special “Terminate (passive)” log entry is created, but the traffic is allowed to continue. This option is useful for testing purposes to make sure a new Terminate rule does not match traffic that you do not want to terminate. The log entry is generated regardless of the Logging options in the rule.
Notifies client and server with a reset	A TCP reset is sent to both communicating parties, so that they are notified that the connection did not succeed (like the Refuse action in Access rules). Further options can be set on the Reset tab for terminated traffic that is not a TCP connection

Table 40.32 Exception Rules - Terminate Action Options - Reset Tab

Option		Explanation
Override collected values set with “Continue” rules		Select the option to activate the settings on this tab and override any options set in a preceding Continue rule that matches the same connections.
Reset Action	Sends an ‘ICMP destination unreachable’ message if not a TCP connection	This option is only used if the Terminate options are set to notify communicating parties with a TCP Reset. When you select **Yes**, the engine sends an ICMP notification to communicating parties when non-TCP traffic is terminated. When you select **No**, no notification is sent for non-TCP traffic. Whether a notification is useful depends on the communicating application.

Table 40.33 Exception Rules - Terminate Action Options - Blacklist Scope Tab

Option	Explanation
Override collected values set with “Continue” rules	Select the option to activate the settings on this tab and override any options set in a preceding Continue rule that matches the same connections.
Terminate the single connection	Creates entries that stop matching current connections, but which are not stored for any time.
Block traffic between endpoints	Creates entries that stop matching connections and block traffic between the matching IP addresses for the set duration.
Blacklist executors	The Security Engines to which the blacklist requests are sent. By default, the engine that detected the event is automatically included in the Blacklist executors.

Table 40.34 Exception Rules - Terminate Action Options - Response Tab

Option	Explanation
Override collected values set with “Continue” rules	Select the option to activate the settings on this tab and override any options set in a preceding Continue rule that matches the same connections.
User Response (*not supported on Virtual Firewalls*)	Define which automatic response is sent to the user when an HTTP connection is terminated due to a Situation match. Response options are always specific to the Situation used. When the rule matches Situation “Anti-Virus_Virus-Found” you can define a response for “Virus found (Firewall)”. When you match the special URL filtering Situations, you can define a response for “URL not allowed”. When you match other HTTP Situations, you can define a response for “Connection terminated by inspection rule”. For more information, see Getting Started with User Responses (page 778).

Defining Logging Options for Inspection Rules and Exceptions

Rules and Exceptions can create a log or alert entry each time they match. By default, logging options set in a previous rule with Continue as its action are used. If no such rule exists, Firewalls and Virtual Firewall log the connections by default. IPS engines and Layer 2 Firewalls do not log the connections by default. Each individual rule can be set to override the default values.

▼ To define logging options for inspection rules and exceptions

1. Open the Logging Options dialog in one of the following ways:
   - (*Inspection tab*) Click the **Logging** setting and select **Logging Options**.
   - (*Exceptions tab*) Double-click the **Logging** cell.
2. Set the options as explained in the table below.

Table 40.35 Inspection Policy - Logging Options

Option		Explanation
Override default value		Activates the settings below and sets the Rule to override mode.
Logging	None	Does not create any log entry. If this option is selected, the Recording options are not available.
	Transient	Creates a log entry that is displayed in the Current Events mode in the Logs view (if someone is viewing it at the moment) but is not stored.
	Stored	Creates a log entry that is stored on the Log Server.
	Essential	Creates a log entry that is shown in the Logs view and saved for further use. When the Log Server is unavailable, log entries are temporarily stored on the engine. When the engine is running out of space to store the log entries, it begins discarding log data in the order of importance. Monitoring data is discarded first, followed by log entries marked as Transient and Stored, and finally log entries marked as Essential. The Alert entries are the last log entries to be discarded. **Note!** The settings for storing the logs temporarily on the engine are defined in the log spooling policy. See Configuring Default Log Handling Settings (page 580).
	Alert	Triggers the alert selected below.
Alert	Select	Click to select a custom Alert for alert escalation. Selecting different alerts for different types of rules allows more fine-grained alert escalation policies.

Table 40.35 Inspection Policy - Logging Options (Continued)

Option		Explanation
Recording	Excerpt	Stores an excerpt of the packet that matched. The maximum recorded excerpt size is 4 KB. This allows quick viewing of the payload in the Logs view.
	Additional Payload	Stores packet payload extracted from the traffic. The collected payload provides information for some of the additional log fields listed in Log Fields Controlled by the Additional Payload Option (page 1252) depending on the type of traffic.
	Record	Records the traffic up to the limit you set in the Record Length field. This allows storing more data than the Excerpt option.
	Record Length	Sets the length of the recording for the Record option in bytes.

Note – Storing or viewing the packets' payload may be illegal in some jurisdictions due to laws related to the privacy of communications.

Importing Snort Rules Libraries

You can import rule definitions from Snort rules library (`.rules`) files. Importing a Snort rules library creates a new Inspection Policy. Each Snort rule is converted into a Situation element and an Exception rule in the Inspection Policy:

- The Action and Source/Destination Network parameters in the Snort rules are used to define the Exception rule.
- The Snort rule options are used to define the Situation element. The Situation element is used in the corresponding Exception rule.
  - The original Snort rule is included as comment in the regular expression for the Situation. For more information about working with Situations elements, see Getting Started With Situations (page 756).
  - If the Situation being imported already exists in the SMC, you must change the original Snort rule (for example, the revision information in the metadata) if you want a new regular expression to be generated.

Snort rules to be imported must be defined using the following Snort Rule syntax:

<action> <protocol> <source_network> <source_port> -> <destination_network> <destination_port> (<option1>[: <arguments>]; ...)

Example **alert tcp $HOME_NET 6789 -> $EXTERNAL_NET any (msg:"BACKDOOR Doly 2.0 access"; flow:established,from_server; content:"Wtzup Use"; depth:32; reference:arachnids,312; classtype:misc-activity; sid:119; rev:5;)**

Snort rules that do not use this syntax may not be imported correctly. Not all Snort rule parameters and options are supported. See the tables below for lists of the supported Snort rule parameters and rule options for importing rules.

Table 40.36 Supported Snort Rule Parameters

Snort Parameter		Description
Action		Defines the Action for the Exception in the Inspection Policy. See Defining What Action an Exception Rule Takes (page 704) for information about changing the Action.
	alert	Converted to Permit, Log Level Alert in the Inspection Policy.
	log	Converted to Permit, Log Level Stored in the Inspection Policy.
	pass	Converted Permit, Log Level None in the Inspection Policy.
	activate	Converted to Permit, Log Level Alert in the Inspection Policy.
	dynamic	Converted to Permit, Log Level Stored in the Inspection Policy.
	drop	Converted to Terminate in the Inspection Policy.
	reject	Converted to Terminate, Active with Reset in the Inspection Policy.
	sdrop	Converted to Terminate, Log Level None in the Inspection Policy.
Protocol		The Context for the Situation element is automatically selected based on the Protocol in the Snort rule. The Situation Context may be further refined based on the Port parameter or the Content-Matching rule options.
	TCP UDP	
Source/Destination Network		Negation of any of the parameters below is also supported. **Example**: !$HOME_NET
	IPv4 Address	A /32 network is assumed. **Example**: 192.168.1.1
	IPv4 Network	In CIDR notation. **Example**: 192.168.1.0/24
	Address List	A comma-separated list of IP address without spaces. **Example**: 192.168.1.1,92.168.1.2
	$alias	The following default Snort aliases are recognized: $HOME_NET $EXTERNAL_NET $HTTP_SERVERS $DNS_SERVERS $SMTP_SERVERS $SQL_SERVERS The corresponding Alias elements are predefined in the SMC. You must define the correct translation values for the Aliases before installing the policy. If you use any custom Snort aliases in rules that you want to import, you must create the corresponding Alias elements in the SMC before importing the Snort rules. See Defining Alias Translation Values (page 534).

Table 40.36 Supported Snort Rule Parameters (Continued)

Snort Parameter	Description
Port	Can be defined as a single number or a port range. If a port range is used, the port information is not used in defining the Context for the Situation element. Negation is also supported. **Examples**: 20 :1024 8000:8010 !1025:
Direction Operator	Only -> and <> are supported.

Table 40.37 Supported Snort Rule Options

Snort Option		Description
Metadata		All of the metadata from the Snort rule is included in the Description in the Situation element.
	sid	Situation elements created from Snort rules are named according to the sid option in the Snort rule (for example, Snort_1450). A unique sid option must be included in every imported rule.
	reference rev msg classtype metadata	These options describe the Snort rule. They do not affect the Snort rule import, and are not used in inspection.
Content-Matching		These options are used to create Situation elements. Some options are used in the regular expression for the Situation. Other options are used to select the correct Context for the Situation. Only the options listed below are supported. Any other options are ignored and a warning is added as comment in the regular expression for the Situation.
	content	Used to generate the regular expression for the Situation.
	uricontent	Used to generate the regular expression for the Situation. Selects HTTP as the Context.
	pcre	Some features of pcre, such as back-reference, are not supported. Other options may cause the DFA generated for the Situation to become very large. Carefully check any imported Snort rules that contain this option.
	byte_jump	Used to generate the regular expression for the Situation.
	byte_test	Used to generate the regular expression for the Situation.
	flow	The flow option is used mainly to select the correct Context for Situations generated from imported Snort rules. It is not used to construct the regular expression for the Situation.

Table 40.37 Supported Snort Rule Options (Continued)

Snort Option		Description
Content-Matching (*cont.*)		
	flowbits	Used to generate the regular expression for the Situation.
	isdataat	Used to generate the regular expression for the Situation.
	offset	Used to generate the regular expression for the Situation. **Note!** Snort detection is packet-based, while Security Engines analyze whole stream. The offset option works as expected only within the first packet of a stream (in each direction)
	within	Used to generate the regular expression for the Situation.
	distance	Used to generate the regular expression for the Situation.
	depth	Used to generate the regular expression for the Situation. **Note!** Snort detection is packet-based, while Security Engines analyze whole stream. The depth option works as expected only within the first packet of a stream (in each direction).
	nocase	Used to generate the regular expression for the Situation.
	http_uri	Selects HTTP as the Context for the Situation.

▼ To import Snort rules

1. Select **File→Import→Import Elements** from the menu. The Import File dialog opens.
2. Select **Snort Rules file (*.rules)** as the file type.
3. Select the file you want to import and click Import. A new tab opens.
4. If the file to be imported contains elements that already exist in the system, select the **Action** to resolve the element *Conflicts*:
   - **Do Not Import**: the element is not imported.
   - **Import**: the element that already exists in the system is overwritten with the element in the import file.
   - **Rename**: the element in the import file is imported as a new element in the system and renamed.
5. If the file to be imported contains *New Elements* that do not conflict with the existing elements, select the **Action**:
   - **Do Not Import**: the element is not imported.
   - **Import**: the element that already exists in the system is overwritten with the element in the import file.
   - For *Identical Elements* that conflict with existing elements, but do not cause changes, the **Do Not Import** option is automatically selected.
6. If there are no similar elements in the system or when you have selected the **Action** for conflicting elements, click **Continue** to start the import.

7. Check that the imported rules match the intended traffic and use the intended responses and options.
   - To edit the imported rules, see Adding Exceptions to the Inspection Policy (page 702).
   - To edit the Context matching options in the Situation element, see Defining Context Options for Situations (page 759).

Note – We recommend validating the policy before installing it in a production environment. See Validating Rules Automatically (page 714)

8. Save the Inspection Policy.

Limiting the Time when a Rule Is Active

Prerequisites: None

You can specify whether Access rules and inspection Exception rules are effective during certain days or times of day only. If you do not specify when a rule is active (the **Time** cell is left empty), the rule is always active.

▼ To specify the validity time of a rule

1. Double-click the **Time** cell in a rule. The Rule Validity Time dialog opens.

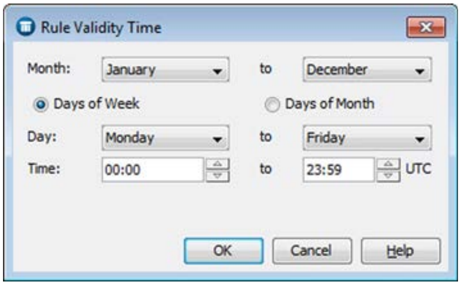

Rule Validity Time

Month: January to December

● Days of Week ○ Days of Month

Day: Monday to Friday

Time: 00:00 to 23:59 UTC

2. Select the **Month** range during which the rule is effective.
3. Select one of the following:
   - **Days of Week**: creates a weekly schedule.
   - **Days of Month**: creates a monthly schedule.
4. In **Day**, select the beginning and end of the period during which the rule is enforced.
5. In **Time**, specify the times of day during which the rule is enforced.

Note – The time is entered in UTC, so you must calculate which UTC day and time corresponds to the time that you want to set. Also remember that UTC does not adjust for Daylight Saving Time (summer time).

6. Click **OK**.

Validating Rules Automatically

Prerequisites: None

You can automatically validate the rules in a policy at any time when you view or edit a policy in the Policy Editing view. You can also validate the policy when you install or refresh the policy on an engine. In both cases, you can also select which issues are checked in the policy.

▼ To validate a policy

1. Start the policy validation in one of the following ways:
   - If the policy is open in the Policy Editing view, click the Tools icon in the policy toolbar and select **Validate**. The Validate Policy dialog opens.
   - If you are installing or refreshing a policy and the Task Properties dialog is open, make sure that the **Validate Policy before Upload** option is selected and click **Select Settings**.

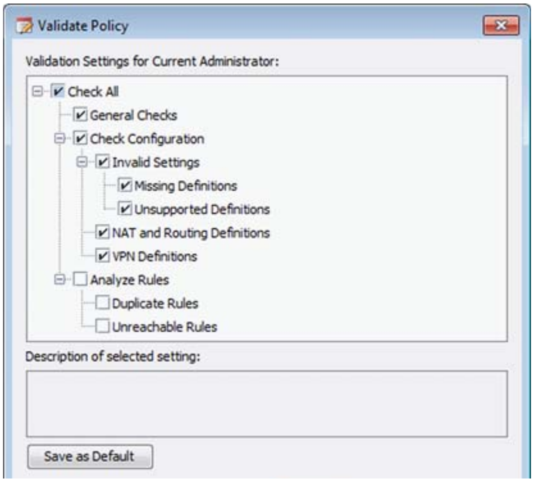

Validate Policy

Validation Settings for Current Administrator:

- ☑ Check All
  - ☑ General Checks
  - ☑ Check Configuration
    - ☑ Invalid Settings
      - ☑ Missing Definitions
      - ☑ Unsupported Definitions
    - ☑ NAT and Routing Definitions
    - ☑ VPN Definitions
  - ☐ Analyze Rules
    - ☐ Duplicate Rules
    - ☐ Unreachable Rules

Description of selected setting:

2. (*Optional; available in the Policy Editing view*) Select the **Target** engine on which you want to install the policy to get more accurate results.
   - The Target engine selection is used to resolve Alias elements when a policy is validated.
   - If no Target engine is selected, all the issues related to the engine configuration cannot be checked (for example, parts of the VPN configuration).
3. (*Optional*) Modify the **Validation Settings** (the types of issues that are checked) as described below:

Table 40.38 Rule Validation Options

Option		Explanation
General Checks		Finds combinations of general settings that are not valid.
Check Configuration	Invalid Settings	Finds incorrect or incomplete definitions.
	Missing Definitions	Checks that there is a definition in all the mandatory matching cells (Source, Destination, and Service cells).
	Routing Definitions (*IPS only*)	Checks routing definitions.

Table 40.38 Rule Validation Options (Continued)

Option		Explanation
Check Configuration (*cont.*)	NAT and Routing Definitions (*Firewall only*)	Checks the NAT and routing definitions.
	VPN Definitions (*Firewall only*)	Checks the VPN configurations in the rule.
Analyze Rules	Duplicate Rules	Checks if there are any identical rules in the policy.
	Unreachable Rules	Finds rules that are in a position in which they can never match.

4. (*Optional*) Click **Save as Default** if you want so save the selected settings as the default set for future policy validations.
5. Click **OK**.

Related Tasks

▶ Finding Unused Rules in Firewall Policies (Hit Counters) (page 667)

Overriding Default Validation Options for Rules

The rule properties allow you to view some rule-specific information and select the settings that are applied to the selected rule when the policy is validated. The rule-specific settings override the default validation options for your administrator account.

▼ To select the Validation Settings for a rule

1. Double-click the rule's **ID** cell. The Rule Properties dialog opens.
2. Switch to the **Validate** tab

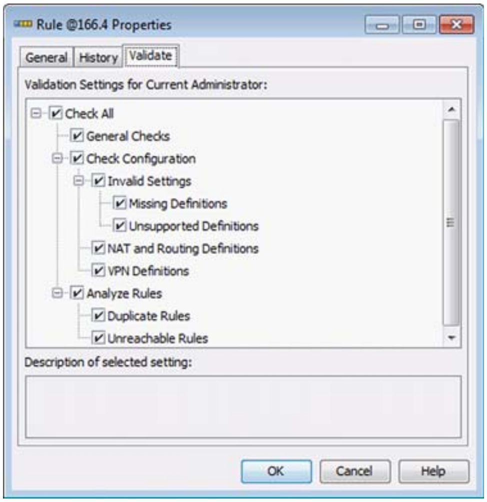

Validation Settings for Current Administrator:

- ☑ Check All
  - ☑ General Checks
  - ☑ Check Configuration
    - ☑ Invalid Settings
      - ☑ Missing Definitions
      - ☑ Unsupported Definitions
    - ☑ NAT and Routing Definitions
    - ☑ VPN Definitions
  - ☑ Analyze Rules
    - ☑ Duplicate Rules
    - ☑ Unreachable Rules

Description of selected setting:

3. Modify the **Validation Settings** (the types of issues that are checked) as described below:

Table 40.39 Rule Validation Options

Option		Explanation
General Checks		Finds combinations of general settings that are not valid.
Check Configuration	Invalid Settings	Finds incorrect or incomplete definitions.
	Missing Definitions	Checks that there is a definition in all the mandatory matching cells (Source, Destination, and Service cells).
	Routing Definitions (*IPS only*)	Checks routing definitions.
	NAT and Routing Definitions (*Firewall only*)	Checks the NAT and routing definitions.
	VPN Definitions (*Firewall only*)	Checks the VPN configurations in the rule.
Analyze Rules	Duplicate Rules	Checks if there are any identical rules in the policy.
	Unreachable Rules	Finds rules that are in a position in which they can never match.

4. Click **OK**.

The selected Validation Settings are now applied to this rule when you next validate the policy. A green check mark is added to the rule's **ID** cell in the rule table to indicate that the Validation Settings of the rule are different from the Validation Settings of the whole policy.

Related Tasks

▶ Validating Rules Automatically (page 714)

Viewing Policy Validation Issues

If policy validation finds something, the issues are displayed in the Issues panel in the Policy Editing view or on the tab that shows the progress of the policy installation.

Illustration 40.20 Issues Panel

Revalidate policy with the same validation settings

Number of found issues

Validate policy and select validation properties

Descriptions of found issue

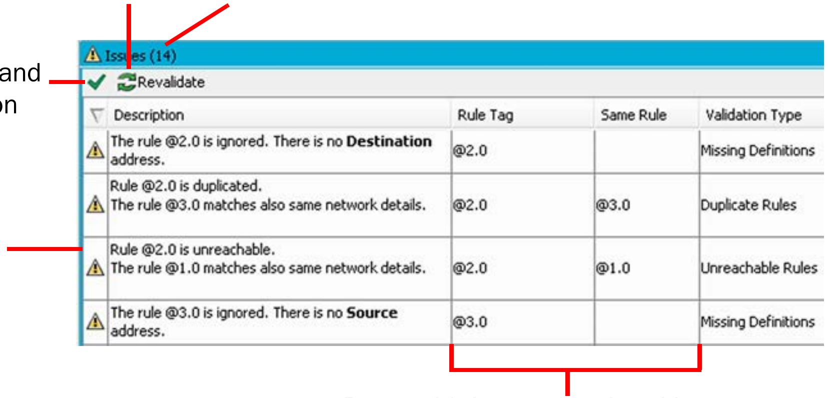

Description	Rule Tag	Same Rule	Validation Type
The rule @2.0 is ignored. There is no **Destination** address.	@2.0		Missing Definitions
Rule @2.0 is duplicated. The rule @3.0 matches also same network details.	@2.0	@3.0	Duplicate Rules
Rule @2.0 is unreachable. The rule @1.0 matches also same network details.	@2.0	@1.0	Unreachable Rules
The rule @3.0 is ignored. There is no **Source** address.	@3.0		Missing Definitions

Rules with issues are listed by rule tag.

▼ To view a configuration in which a validation issue was found

➥ Double-click the issue. The relevant part of the configuration is shown.

In rules, the **ID** cell shows the status of validation issues in the rule. If issues are found for a rule, the rule's **ID** cell has a yellow triangle with an exclamation mark. If a rule's Validation Settings override the Validation Settings of the whole policy, the rule's **ID** cell has a green check mark (see Overriding Default Validation Options for Rules (page 715) for more information).

Illustration 40.21 Rules with Validation Issues or Rule-Specific Validation Settings

Rule in which a validation issue was found

Rule with modified Validation Settings

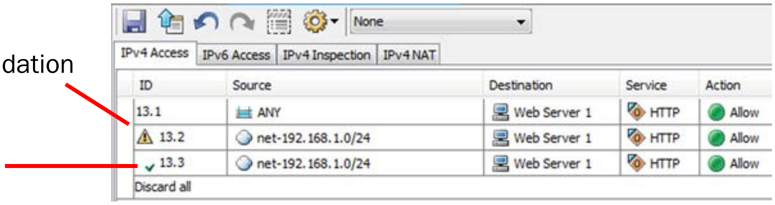

ID	Source	Destination	Service	Action
13.1	ANY	Web Server 1	HTTP	Allow
⚠ 13.2	net-192.168.1.0/24	Web Server 1	HTTP	Allow
✓ 13.3	net-192.168.1.0/24	Web Server 1	HTTP	Allow

What's Next?

- ▶ Fix the issues that are indicated.

Related Tasks

- ▶ Disabling a Validation Warning for a Rule (page 718).
- ▶ Excluding Rules from Policy Validation (page 718).
- ▶ Overriding Default Validation Options for Rules (page 715).

Disabling a Validation Warning for a Rule

You can optionally set an issue listed in the Issues panel to be ignored for a specific rule.

▼ To exclude a validation issue for a rule

➥ Right-click the rule in the Issues panel and select **Ignore Issue Type for Rule**.

This issue type is no longer checked for this rule. A green check mark is added to the rule's **ID** cell in the rule table to show that the rule's Validation Settings are not the same as the policy's Validation Settings.

You can change the overall selection of validation issues for a rule in the rules's properties. See Overriding Default Validation Options for Rules (page 715).

Excluding Rules from Policy Validation

If you do not want to validate a certain rule in the policy, you can exclude the rule from policy validation.

▼ To exclude a rule from policy validation

➥ Right-click the rule in the Issues panel and select **Disable All Issues for the Rule**.

The rule is no longer checked when you validate the policy. A green check mark is added to the rule's **ID** cell in the rule table to show that the rule's Validation Settings are not the same as the policy's Validation Settings.

Changing Default Rules

Prerequisites: None

In most cases, templates and policies are inherited from the default Template policies in the system. However, it is not possible to modify these system elements. If you must modify the default templates, you can create a copy.

Version upgrades and dynamic updates may require changes to the default Template policies. These changes are not applied to any copies of the templates. You must manually modify the copies to ensure that the system communications continue to be allowed and all necessary inspection continues to be performed.

▼ To create a custom version of a default Template policy

1. Right-click the template and select **New→Duplicate**.
2. Save the copy of the template under a different name.
3. Right-click existing policies that should use the modified template and select **Properties**.
4. Select your copy of the template in the **Template** panel and click **OK**.

Caution – The operation of the engines can be seriously disturbed by incorrect modifications to the default Template policies.

CHAPTER 41

DEFINING IP ADDRESSES

IP addresses are defined as elements in the system to allow reuse of the same definitions in any number of configurations for any types of components in the system. Any element that represents at least one IP address is a *Network Element* and can be used to define IP addresses in policies (with some restrictions).

The following sections are included:

Getting Started with Defining IP Addresses

Prerequisites: None

The elements that you can use for defining IP addresses are called *Network Elements* (not to be confused with the *Network* element, which defines an IP network). Each element can be inserted in several places in the Access, Inspection, and/or NAT rules (as source and destination of traffic) and in many other places where you need to define IP addresses (for example, in routing and log filtering).

There are elements in the system whose only role is to define an IP address. But also elements created for configuring a feature in the system can also be used in policies (with some limitations) if they represent an IP address.

To access network elements, select **Configuration→Configuration→Security Engine** from the menu and expand the **Network Elements** tree.

To create/modify elements:

- Click the New icon in a view that is currently displaying Network Elements.
- Right-click the Network Elements category or a sub-category in a tree view and select **New** from the menu that opens.
- To create a copy of an existing element, right-click it and select **New→Duplicate**.
- To modify an existing element, right-click an element where it appears and select **Properties**. Note that default system elements cannot be edited.

What's Next?

- ▶ The primary tools for defining IP addresses are elements that only exist to represent one or more IP addresses. These are explained in Defining IP Addresses as Elements (page 721).
- ▶ Many elements that are configured as parts of some feature can also be used to represent IP addresses. The special issues related to these elements are discussed in Using Feature-Specific Elements in Policies (page 731).

Defining IP Addresses as Elements

Prerequisites: None

Different types of elements allow you to flexibly define any set of IP addresses:

- *Address Ranges* allow you to define any continuous range of IP addresses. See Defining Address Range Elements.
- *Aliases* represent a variable value in policies. The IP address value is filled in based on the engine on which the policy is installed. Aliases make using the same policy on several engines practical. See Defining Alias Elements (page 722).
- *Domain Names* represent all the IP addresses that belong to a particular domain. See Defining Domain Name Elements (page 723).
- *Expressions* allow you to define any set of IP addresses in a single element. They are especially well suited for excluding some IP addresses from otherwise continuous ranges. See Defining Expression Elements (page 724).
- *Groups* allow you to combine different types of elements into a single element. See Defining Group Elements (page 726).
- *Hosts* represent a single device in the network. Each Host can represent one or more individual IP addresses in policies. See Defining Host Elements (page 727).
- *Networks* represent complete network segments. See Defining Network Elements (page 728).
- *Routers* represent a gateway device in the network and are primarily meant for configuring routing. Each Router can represent one or more IP addresses in policies. See Defining Router Elements (page 728).
- *Zones* allow you to combine engines’ network interfaces into a single element. See Defining Zone Elements (page 730).

Some elements that are primarily meant for configuring a certain feature also define an IP address. Such elements can be used in policies, but there are some additional considerations. See Using Feature-Specific Elements in Policies (page 731).

Defining Address Range Elements

An *Address Range* element can specify any continuous range of IP addresses.

Table 41.1 Address Range Properties

Property	Description
Name	The element’s name in the Management Center. Used, for example, to resolve IP addresses to names in the logs.
IP Range	Either two valid IPv4 addresses or two valid IPv6 addresses. You must enter the same type of address in both fields. The address value on the left (start of range) must be lower than the address value on the right (end of range).
Category (*Optional*)	Categories allow you to flexibly filter your Management Client view. See Using Categories (page 76).
Comment (*Optional*)	Free-form comment for your reference. The comment is displayed in the element’s Properties dialog and in the Info view.

Defining Alias Elements

Aliases are like variables: they can be given different IP address values in the policy depending on the engine on which the policy is installed. This makes it possible to create rules that are valid on several engines without including all IP addresses in the policies of all elements. Alias elements are especially useful in Template Policies.

Tip – You can view the IP addresses that the Alias represents in the policy of each component: switch the policy view to the "network details" mode through the Tools icon in the view-specific toolbar and then select a component from the list that is added to the toolbar.

Some of the default system Aliases always receive their values directly from other parts of each engine's configuration and cannot be modified; these Aliases start with two $$ symbols. There are also some default Aliases in the system that either do not receive any value without your action, or allow you to add to and change the default values. These Aliases start with one $ symbol, as do all Alias elements you create yourself. For descriptions of default aliases, see Predefined Aliases (page 1175).

Table 41.2 Alias Properties

Property	Description
Name	The element's name in the Management Center. Used, for example, to resolve IP addresses to names in the logs.
Category (*Optional*)	Categories allow you to flexibly filter your Management Client view. See Using Categories (page 76).
Comment (*Optional*)	Free-form comment for your reference. The comment is displayed in the element's Properties dialog and in the Info view.
Default Value (*Optional*)	The value that is used when the policy is installed on a component that does not have a more specific value for the Alias. If there is no default value, falling back to the default translation returns an empty value. If the Alias is the only element in the rule cell, the rule is invalid and ignored when the policy is installed (a warning is displayed if this happens).
Component-Specific Values (*Optional*)	The values that the Alias is replaced with when used in the policies of the components listed. If there is no specific translation value for a component, the default value is used instead.

Defining Domain Name Elements

A *Domain Name* element represents all the IP addresses that belong to a particular domain. If you have entered the IP addresses of one or more DNS servers in the engine properties, the Firewall, IPS, and Layer 2 Firewall engines periodically queries the DNS server to automatically resolve domain names to IP addresses (see Creating New Engine Elements (page 359) for more information). This makes it possible to create rules that are valid even if new addresses are added to the domain or the domain's IP addresses change. If the DNS server returns multiple IP addresses for the same domain name, the engine associates all of the IP addresses with the domain name. However, if there are a very large number IP addresses associated with the same domain name, the DNS server may only reply with a few of the IP addresses at a time. In this case, the engine may need to make additional queries to the DNS server to resolve all of the IP addresses for the domain name. By default, the engine queries the DNS server every six minutes. Resolved IP addresses are kept in the engine's DNS cache for a maximum of one hour by default.

Note – The DNS cache is not synchronized between nodes of a cluster. Each node separately queries the DNS server using the node's NDI address. It is possible that the DNS cache may be different on different nodes of a cluster.

Table 41.3 Domain Name Properties

Property	Description
Name	The name of the domain. Used, for example, to resolve domain names to IP addresses in policies and logs. The domain names must be in a format that the Firewall, IPS and Layer 2 Firewall engines can resolve (for example, www.example.com).
Category (*Optional*)	Categories allow you to flexibly filter your Management Client view. See Using Categories (page 76).
Comment (*Optional*)	Free-form comment for your reference. The comment is displayed in the element's Properties dialog and in the Info view.

Defining Expression Elements

Expression elements allow you to combine other Network Elements with logical operators to represent complex sets of IP addresses.

If you are unfamiliar with the concepts presented here, we recommend reading the chapter on expressions in the *Management Center Reference Guide* before you start creating Expression elements.

The main parts of the Expression properties are explained in the illustration below.

Illustration 41.1 Expression Element Properties

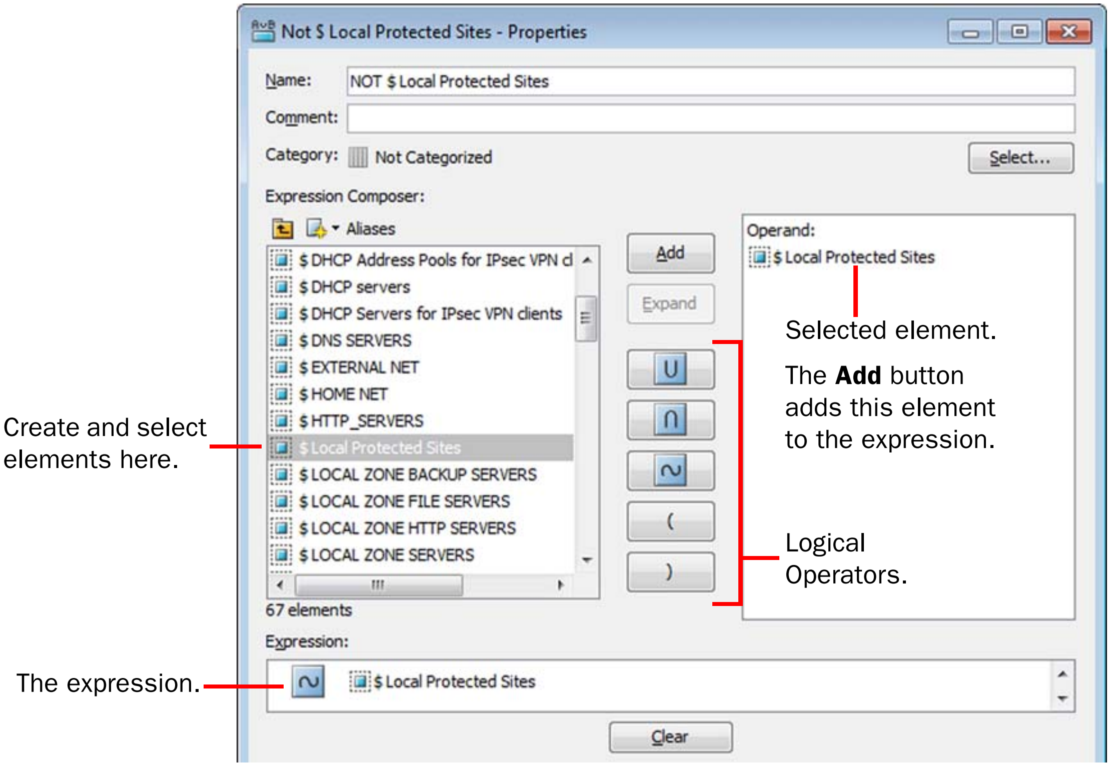

Not $ Local Protected Sites - Properties

Name: NOT $ Local Protected Sites

Comment:

Category: Not Categorized

Expression Composer:

Aliases

- $ DHCP Address Pools for IPsec VPN cl
- $ DHCP servers
- $ DHCP Servers for IPsec VPN clients
- $ DNS SERVERS
- $ EXTERNAL NET
- $ HOME NET
- $ HTTP_SERVERS
- $ Local Protected Sites
- $ LOCAL ZONE BACKUP SERVERS
- $ LOCAL ZONE FILE SERVERS
- $ LOCAL ZONE HTTP SERVERS
- $ LOCAL ZONE SERVERS

67 elements

Operand: $ Local Protected Sites

Expression: ∼ $ Local Protected Sites

Create and select elements here.

Selected element.

The **Add** button adds this element to the expression.

Logical Operators.

The expression.

Logical Operators in Expressions are resolved in the following order (however, items inside parentheses are evaluated before any items outside parentheses are considered):

1. Negations.
2. Intersections.
3. Unions.

▼ To edit the expression

1. Select the first element in the left panel.
2. Click **Add**.
3. Select a logical operator.
4. Select the second element in the left panel.
5. Click **Add**.

Repeat the steps to create a more complex expression. Expressions can also start with a negation operation or a section in parenthesis. AND and OR logical operators must have an element (or a section in parenthesis) on both sides of the operator.

You can also:

- Click the expression to view the cursor. Click again or use the arrow keys to move the cursor.
- When the cursor is visible, you can press the Delete or Backspace key to remove an item (no undo).

Tip – You can view the actual IP addresses that the element adds to a policy: insert the expression in a rule and activate the "network details" mode through the Tools icon in the view-specific toolbar.

Table 41.4 Expression Properties

Property	Description
Name	The element's name in the Management Center. Used, for example, to resolve IP addresses to names in the logs.
Comment (*Optional*)	Free-form comment for your reference. The comment is displayed in the element's Properties dialog and in the Info view.
Category (*Optional*)	Categories allow you to flexibly filter your Management Client view. See Using Categories (page 76).
Add	Adds the item in the top right panel into the expression at the bottom at the cursor location.
Expand/ Compact	If the selected item in the right panel is an Expression element, the button switches the view between showing the element and showing the expression defined in the element.
∪	Union. Includes all IP addresses defined in both elements that the union combines.
∩	Intersection. Includes only those IP addresses that the two elements that the intersection combines have in common.
~	Negation. Includes all IP addresses except those that are defined in the element that follows the negation.
( )	Parentheses group parts of the expression together. Items inside parenthesis are evaluated before they are compared to items outside the parenthesis.

Example Expression defining any other network but one internal network (172.16.2.0/24): “[negation] net-172.16.2.0/24”.

Example Expression for selecting Hosts in the Network 172.17.1.0/24 from a Group element which contains Hosts from many different networks: “net-172.17.1.0/24 [intersection] Intranet Hosts Group”.

Defining Group Elements

The *Group* element can be used to combine any number of previously defined elements into a single element. The elements can be of different types. You can use Group elements in policies to make the policies clearer to read and to simplify the editing of configurations where the same elements always appear together. You can also use Groups to add monitored elements that are displayed in the System Status view.

Example Host elements for file servers could be gathered together in a Group element, which is then used in several rules in Firewall, IPS, and Layer 2 Firewall policies. When a new file server is introduced, it is simply added to the Group. The change is propagated to all rules in all policies in which the Group is used.

Table 41.5 Group Properties

Property	Description
Name	The element’s name in the Management Center. Used, for example, to resolve IP addresses to names in the logs.
Comment (*Optional*)	Free-form comment for your reference. The comment is displayed in the element’s Properties dialog and in the Info view.
Category (*Optional*)	Categories allow you to flexibly filter your Management Client view. See Using Categories (page 76).
Resources	Select elements that you want to add to the Group in this panel.
Content	The elements that are included in the Group.
Use for Status Monitoring (*Optional*)	When selected, adds the Group to the System Status tree if the Group contains elements that can be monitored.

Defining Host Elements

A *Host* element represents the IP address(es) of any single device. Host elements are used to represent individual devices that have no additional special role in the system configuration (such as being a next-hop router, an external authentication server etc.).

Table 41.6 Host Properties - General Tab

Property	Description
Name	The element's name in the Management Center. Used, for example, to resolve IP addresses to names in the logs.
IPv4 Address (*Optional**)	Single valid IPv4 or IPv6 address. You can enter either one or both addresses. (*Either an IPv4 or IPv6 address is mandatory.)
IPv6 Address (*Optional**)	
Secondary IP Addresses (*Optional*)	If the device has additional IP addresses, you can enter them here instead of creating additional elements for the other IP addresses. The secondary IP addresses are valid in policies and in routing and antispoofing. You can add several IPv4 and IPv6 addresses (one by one).
Category (*Optional*)	Categories allow you to flexibly filter your Management Client view. See Using Categories (page 76).
Tools Profile (*Optional*)	Allows you to add commands to the element's right-click menu. See Adding Custom Commands to Element Menus (page 55).
Comment (*Optional*)	Free-form comment for your reference. The comment is displayed in the element's Properties dialog and in the Info view.

Table 41.7 Host Properties - Monitoring Tab

Property	Description
Log Server	Select the Log Server that monitors this device (third-party monitoring). You must select a Log Server to activate the other options. See Getting Started with Third-Party Device Monitoring (page 132).
Status Monitoring	Activates status monitoring for this device. You must also select the **Probing Profile** that contains the definitions for the monitoring. When you select this option, the element is added to the tree in the System Status view. See Creating a Probing Profile (page 143).
Log Forwarding	Activates syslog forwarding from this device. You must select the **Logging Profile** that contains the definitions for converting the syslog entries to Stonesoft log entries. See Converting Logs From External Devices (page 133). You must also select the **Time Zone** in which this device is located. By default, the local time zone of the computer you are using is selected.
SNMP Trap Forwarding	Activates SNMP trap forwarding from this device. The data that the device sends must be formatted according to the MIB definitions currently active in the system. See Importing MIBs (page 142).

Table 41.7 Host Properties - Monitoring Tab (Continued)

Property	Description
NetFlow Forwarding	Activates NetFlow data forwarding from this device to the Log Server.

Defining Network Elements

A Network element represents the IP address space of a complete network or subnetwork.

Table 41.8 Network Properties

Property	Description
Name	The element's name in the Management Center. Used, for example, to resolve IP addresses to names in the logs.
IPv4 Address and IPv4 Netmask (*Optional**)	A valid IPv4 address and netmask and/or a valid IPv6 address and prefix length. You can enter either one or both addresses. (*Either an IPv4 address and netmask or an IPv6 address and prefix length are mandatory.)
IPv6 Address and Prefix Length (*Optional**)	
Category (*Optional*)	Categories allow you to flexibly filter your Management Client view. See Using Categories (page 76).
Comment (*Optional*)	Free-form comment for your reference. The comment is displayed in the element's Properties dialog and in the Info view.
Broadcast and Network Addresses Included	Includes the broadcast address and the network address in the definition. **This option is only used in the Source and Destination cells in rules**. In other uses, the broadcast and network addresses are always included in the definition regardless of this option.

Defining Router Elements

A Router element represents a next-hop gateway's IP address in routing configurations when the Router has a fixed IP address. The element can also be used to represent IP address(es) in rules and other configurations as needed.

Note – If the interface towards the next-hop gateway has a dynamic IP address, a special *Gateway (DHCP Assigned)* element must be added directly through the right-click menu of the automatically added *Network (DHCP assigned)* element in the Routing view. The Gateway (DHCP Assigned) element is not valid in policies; use a corresponding Alias element instead, see Predefined Aliases (page 1175).

Table 41.9 Router Properties - General Tab

Property	Description
Name	The element's name in the Management Center. Used, for example, to resolve IP addresses to names in the logs.
IPv4 Address (*Optional**)	A valid IPv4 address and/or a valid IPv6 address. You can enter either one or both addresses. (*Either an IPv4 address or an IPv6 address is mandatory.)
IPv6 Address (*Optional**)	
Secondary IP Addresses (*Optional*)	If the device has additional IP addresses, you can enter them here instead of creating additional elements for the other IP addresses. The secondary IP addresses are valid in policies and in routing and antispoofing. You can add several IPv4 addresses (one by one).
Category (*Optional*)	Categories allow you to flexibly filter your Management Client view. See Using Categories (page 76).
Tools Profile (*Optional*)	Allows you to add commands to the element's right-click menu. See Adding Custom Commands to Element Menus (page 55).
Comment (*Optional*)	Free-form comment for your reference. The comment is displayed in the element's Properties dialog and in the Info view.

Table 41.10 Router Properties - Monitoring Tab

Property	Description
Log Server	Select the Log Server that monitors this device (third-party monitoring). You must select a Log Server to activate the other options. See Getting Started with Third-Party Device Monitoring (page 132).
Status Monitoring	Activates status monitoring for this device. You must also select the **Probing Profile** that contains the definitions for the monitoring. When you select this option, the element is added to the tree in the System Status view. See Creating a Probing Profile (page 143).
Log Forwarding	Activates log forwarding from this device. You must select the **Logging Profile** that contains the definitions for converting the syslog entries to Stonesoft log entries. See Converting Logs From External Devices (page 133). You must also select the **Time Zone** in which this device is located. By default, the local time zone of the computer you are using is selected.
SNMP Trap Forwarding	Activates SNMP trap forwarding from this device. The data that the device sends must be formatted according to the MIB definitions currently active in the system. See Importing MIBs (page 142).
NetFlow Forwarding	Activates NetFlow data forwarding from this device to the Log Server.

Related Tasks

- Configuring Routing (page 587)

Defining Zone Elements

Zone elements allow you to group together network interfaces of Firewall, IPS, and Layer 2 Firewall engines. You can use Zones to specify the receiving or sending interfaces in policies. The Zone element represents all the interfaces that belong to the Zone. All the rules that include a Zone element also apply to any new interfaces that you associate with the same Zone.

There are several predefined System Zones available:

- **DMZ**: interfaces connected to DMZ networks.
- **External**: interfaces connected to the Internet or other external networks.
- **Guest**: interfaces connected to guest networks.
- **Internal**: interfaces connected to internal networks.
- **Node-internal**: Firewall, IPS, and Layer 2 Firewall nodes themselves. This Zone is automatically assigned to interfaces through which traffic to or from the engine node travels. It cannot be assigned to other interfaces, but it can be used in policies.

You can also create your own Zones as described below:

Table 41.11 Zone Properties

Property	Description
Name	The element's name in the Management Center. Used to represent network interface identifiers in Firewall, IPS, and Layer 2 Firewall policies.
Category (*Optional*)	Categories allow you to flexibly filter your Management Client view. See Using Categories (page 76).
Comment (*Optional*)	Free-form comment for your reference. The comment is displayed in the element's Properties dialog and in the Info view.

Related Tasks

- ▶ Network Interface Configuration (page 419)
- ▶ Defining Source, Destination, and Service Criteria (page 664)
- ▶ Using Zones in the Destination of Access Rules (page 673)

Using Feature-Specific Elements in Policies

Prerequisites: None

Many elements are created as part of configuring a particular feature. When such elements define an IP address for a device, the element can also be used to represent the IP address in policies. However, there are some special issues that may have to be considered depending on the element type.

Tip – To view the actual IP addresses that the element adds to a policy, insert the element in a rule and activate the "network details" mode through the Tools icon in the view-specific toolbar.

System Components

Using elements that represent system components as a source or destination IP address in policies may produce unexpected results. Be careful especially when you use engine elements as source or destination IP addresses in policies:

- Firewall, Single IPS, IPS Cluster, Single Layer 2 Firewall, and Layer 2 Firewall Cluster: These elements represent all static IP addresses defined for all interfaces. Create separate Host elements to represent individual IP addresses.
- Firewall Cluster: Represent all CVI IP addresses of all interfaces, but not the NDI addresses. Create separate Host elements to represent individual CVI addresses and NDI addresses.
- Firewalls with dynamic IP addresses: The Firewall element does not represent any of the dynamic IP addresses. There are default Aliases that can be used to represent the firewall's own dynamic IP addresses in the firewall's own policy. Fixed IP address definitions are needed for the dynamically assigned IP addresses when they need to be defined in the policies of any other components.
- SSL VPN Gateway: the element represents the single IP address that is the source address for centralized monitoring and logging information.
- Management Center servers: represent the single primary IP address defined for the element.
- Contact addresses are not considered when the element is used in a policy. Consider which IP address has to be added to the rule and create separate Host elements for the contact addresses as necessary.

External Servers

Several types of external servers can be integrated with the SMC when configuring different features. In general, each server element simply represents the single primary IP address defined in the element when used in a policy. Some elements have additional considerations when used in policies:

- Secondary IP addresses: Many server elements can contain one or more secondary IP addresses in addition to the primary address displayed for the element. The secondary addresses are equally valid in policies.
- Contact addresses: Some server elements can have a contact address. Contact addresses are not considered when the element is used in a policy. Consider which IP address has to be added to the rule and create separate Host elements for the contact addresses as necessary.
- Server Pools: The Server Pool element represents the external address(es) that the clients contact. Use the Server Pool in rules that allow clients' traffic to the servers whenever you want to use the Server Pool features. Elements that represent the individual members of the pool can be used to allow connections to individual pool members (for example, to allow remote administration of each server).

Traffic Handlers

Traffic handlers are necessary in policies when configuring Multi-Link for Firewall. They can be used in rules as follows:

- In Source and Destination cells: A NetLink element represents the whole network address space that is associated with the NetLink element. An Outbound Multi-Link element represents the network address spaces of all NetLinks included. Traffic handlers are seldom used this way.
- In the NAT cell in NAT rules: When the source address is translated using the Outbound Multi-Link element as the address space, the traffic is balanced between the included NetLinks according to the options selected for the Outbound Multi-Link.

Related Tasks

- ▶ Creating an Outbound Multi-Link Element (page 612)
- ▶ Defining a Server Pool (page 619)

Chapter 42

Defining Network Services

Service elements match traffic based on protocol or port and set options for advanced inspection of traffic. Services are used in Firewall Policies, IPS Policies, and Layer 2 Firewall Policies.

The following sections are included:

Getting Started with Services

Prerequisites: None

What Services Do

Service elements specify a network protocol, as well as source and/or destination ports for TCP and UDP traffic.

- Services can be used to match rules to traffic in Ethernet rules (Ethernet Services), Access rules, and NAT rules.
- Services can refer to *Protocol* elements, which activate further inspection checks and advanced traffic handling. Some Protocols have additional options that you can set in the Service element's properties.

What Do I Need to Know Before I Begin?

Most of the time, you can use the default Services already available in the system for standard protocols and ports. However, you may need to occasionally create a new Service:

- If none of the default Services match the type of traffic you want to allow, for example, if some TCP or UDP service in your network uses a non-standard port.
- If you want to set options for advanced traffic handling, for example, Access rules that disallow the use of either the active or passive FTP connection mode, Firewall Access rules for CIS (content inspection server) redirection, or Firewall Access rules and Firewall NAT rules for protocols that assign ports dynamically inside the packet payload.

Configuration Overview

Illustration 42.1 Elements in the Services Configuration

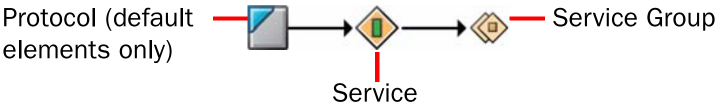

1. Create a new Service that matches the correct protocol number and port number (if applicable), see Defining Services (page 735).
2. (*Optional*) Select one of the default Protocol elements if you want the traffic to be inspected further or if you want to use a Protocol Agent. See Using Protocol Elements (page 738) and Defining Protocol Parameters (page 739).
3. (*Optional*) Add the Service to a Service Group to make it easier to insert several related Services into configurations, see Grouping Services (page 738).

Defining Services

Prerequisites: None

There are predefined Service elements in the system that correspond to reserved and commonly used protocols and ports. In addition to these, you may need to add your own customized Service elements to the system for any non-standard ports in use or if you want to define options for advanced traffic handling (Protocol Agents).

What's Next?

- ▶ To define a Service for Access, Inspection, or NAT rules, see Defining a New IP-Based Service.
- ▶ To define a Service for Ethernet rules, see Defining a New Ethernet Service (page 737).

Defining a New IP-Based Service

IP-based services are used in Access rules and NAT rules. Make sure that you are clear on which underlying protocol the traffic you want to allow uses, and be aware of whether you must define a protocol number or a port number. Usually, the Services you define yourself will be TCP- or UDP-based and are identified by the port number they use. However, there are many common protocols that are not TCP- or UDP-based (for example, ICMP and RPC) and are identified by other information.

Example **For example, the GRE protocol is transported directly over IP as protocol number 47 - on the same layer as TCP (#6) and UDP (#17). Therefore, any custom Services created for TCP and UDP ports 47 do not allow GRE to pass the Firewall.**

▼ To define a new custom service

1. Select **Configuration→Configuration→Security Engine** from the menu. The Security Engine Configuration view opens.
2. Expand the **Other Elements** tree and select **Services**.
3. Create the new Service element:
   - To create a new element with no settings predefined, right-click the branch for the type of Service you want to create and select **New**→[Service type] **Service**. The Service Properties dialog opens without any options set.
   - To create a new Service based on some other Service element, right-click the existing Service and select **New→Duplicate**. The Service Properties dialog opens with options copied from the existing Service.

Illustration 42.2 TCP Service Properties

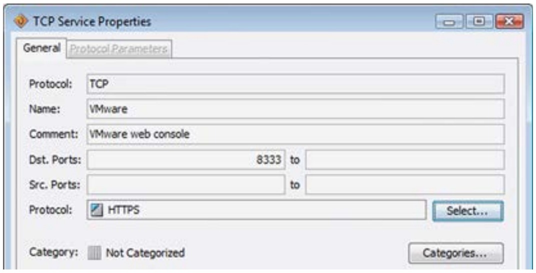

TCP Service Properties

General | Protocol Parameters

Protocol: TCP

Name: VMware

Comment: VMware web console

Dst. Ports: 8333 to

Src. Ports: to

Protocol: HTTPS

Category: Not Categorized

4. Give the new Service a unique **Name** and write an optional **Comment**.

5. Identify the traffic that this Service element matches depending on the Service type as explained in the table below (IANA assigns the protocol codes; visit www.iana.org for a list):

Table 42.1 Traffic Identification Options in IP-Based Service Elements

Protocol	Option	Explanation
TCP and UDP	Dst. Ports (*Optional**)	A destination port or port range. To match a single port, enter it in the first field and leave the other field empty. To enter a range, enter a value in both fields. (*Either source or destination port is mandatory.)
	Src. Ports (*Optional**)	A source port or port range. To match a single port, enter it in the first field and leave the other field empty. To enter a range, enter a value in both fields. (*Either source or destination port is mandatory.)
ICMP	Type	The ICMP type number that the traffic uses.
	Code (*Optional*)	The ICMP code that the traffic uses. If the Code field is empty, the Service matches traffic regardless of the ICMP code. If you enter 0 as the code, the Service matches only packets that contain no ICMP code.
SUN RPC	Program Number	The program number.
	Version (*Optional*)	The remote program version number. If you do not enter a program version, the element matches traffic of any version.
	Allow TCP (*Optional*)	Allows the RPC message when transported over TCP.
	Allow UDP (*Optional*)	Allows the RPC message when transported over UDP.
IP	Code	The numerical protocol code that the traffic uses.

6. (*Optional*) If you want to associate the Service with a Protocol, click the **Select** button for the **Protocol** field and select the correct Protocol.
   - Selecting the Protocol is mandatory if the Service is inserted in an Access rule that directs packets to deep inspection against the inspection rules.
   - Some types of traffic may require a Protocol of the type Protocol Agent.

What's Next?

▶ Some Protocols have more options that you can set on the **Protocol Parameters** tab, see Using Protocol Elements (page 738).

▶ Otherwise, click **OK** to create the new Service element.

Related Tasks

▶ Grouping Services (page 738).

▶ Defining What Traffic an Access Rule Matches (page 672).

▶ Defining What Traffic a NAT Rule Matches (page 688)

Defining a New Ethernet Service

There are predefined Ethernet Service elements in the system that correspond to commonly used Ethernet services. The predefined Ethernet Services should cover most needs you have, but you can add additional Ethernet Services to the system if you want to filter other Ethernet-level traffic, use different names, or add your own comment for the Services.

Caution – Match any IP traffic you allow in Ethernet rules to the default IPv4 and IPv6 Services. These Services match the traffic using the correct Protocol element. Only IP traffic matched to the correct Protocol element is inspected further against the Access rules. Non-IP traffic is never inspected any further.

Ethernet Services are used with IPS engines and Layer 2 Firewalls.

▼ To define a new Ethernet Service

1. Select **Configuration→Configuration→Security Engine** from the menu. The Security Engine Configuration view opens.
2. Expand the **Other Elements** tree.
3. Right-click **Ethernet Services** and select **New→Ethernet Service**. The Ethernet Service Properties dialog opens.
4. Give the new Ethernet Service a unique **Name**.
5. Select the Ethernet-level protocol and enter the details as explained in the table below (IEEE assigns the protocol codes).

Table 42.2 Traffic Identification Options in Ethernet Service Elements

Protocol	Option	Explanation
Ethernet 2 (DIX)	MAC type	Enter the EtherType code of the protocol that the traffic uses.
Raw IPX (Novell)	(none)	Matches IPX (internetwork packet exchange) traffic.
LLC	SSAP	The SSAP (source service access point) address that the traffic uses.
	DSAP	The DSAP (destination service access point) address that the traffic uses.
SNAP	Vendor	The OUI (organizational unique identifier) that the traffic uses.
	Type	The type that the traffic uses.
Protocol	N/A	Not editable. Present in the default IPv4 or IPv6 services to direct traffic to further filtering and inspection. Not present in any custom Service elements you create.

6. Click **OK**.

What's Next?

▶ To group similar Ethernet Services together, see Grouping Services (page 738).

Related Tasks

- Editing Ethernet Rules (page 668)

Grouping Services

Grouping several Services into Service Group elements simplifies your policies, since you can insert the Group instead of several individual Service elements. You can group both default system Services as well as your custom-made elements.

▼ To create a new Service Group

1. Select **Configuration→Configuration→Security Engine** from the menu. The Security Engine Configuration view opens.
2. Expand the **Other Elements** tree.
3. Right-click the **Services** or (*IPS only*) **Ethernet Services** branch according to the type of service you want to create and select either:
   - **New**→[Service type] **Service Group** to create a Service Group under the branch for the selected type of services (for example, TCP).
   - or select **New→Service Group** to create the Service Group under the Group branch (item not available if you right-clicked Ethernet Services).
4. Give the new Service Group a unique **Name**.
5. Select the required Services in the list on the left and click **Add**.
6. Click **OK**.

Using Protocol Elements

Prerequisites: Defining a New IP-Based Service

Protocol elements can be inserted directly in Inspection rule exceptions. In Access rules, the Protocol elements are always contained in a Service element, which can then be inserted into the Service cell in rules. Some Protocols add options that you can adjust to custom Service elements that you create. You cannot add or modify the Protocol elements directly.

All Protocol Elements identify traffic as being of a certain network protocol. A Protocol element in Access rules identifies the protocol for inspection against Inspection rules. In Inspection rules, the Protocol can be used to limit the scope of exception rules according to the Protocol (as identified in the Access rules) in rules that otherwise match a large number of Protocols. Additionally, the Protocols may activate some additional software modules on the engines. This depends on the type of the Protocol element:

- **Protocol Tag**: a Protocol element that does not activate additional modules.
- **Protocol Agent**: a Protocol element that activates an additional module on the engines to provide advanced application layer features.

Some Protocol elements have additional parameters. See Defining Protocol Parameters (page 739).

Defining Protocol Parameters

Many Protocols provide options for you to set in the Service that uses them. Some of the parameters are only used by a specific type of component. To set options for a Protocol, you must attach it to a custom Service element: either open the properties of a custom Service you have created previously or create a new Service.

See the following sections for more information:

- Defining DNS Protocol Parameters.
- Defining FTP Protocol Parameters (page 740)
- Defining GRE Protocol Parameters (page 741)
- Defining H323 Protocol Parameters (page 742)
- Defining HTTP/HTTPS Protocol Parameters (page 742)
- Defining IPv4 Encapsulation Protocol Parameters (page 744)
- Defining MSRPC Protocol Parameters (page 745)
- Defining NetBIOS Protocol Parameters (page 746)
- Defining Oracle Protocol Parameters (page 746)
- Defining Shell (RSH) Protocol Parameters (page 747)
- Defining SIP Protocol Parameters (page 748)
- Defining SMTP Protocol Parameters (page 749)
- Defining SSH Protocol Parameters (page 749)
- Defining SunRPC Protocol Options (page 750)
- Defining TCP Proxy Protocol Parameters (page 751)
- Defining TFTP Protocol Parameters (page 752)

Defining DNS Protocol Parameters

The DNS protocol parameters control the DNS protocol enforcement on the IPS engines and Layer 2 Firewalls. When you activate this feature, the engines can determine if traffic on the DNS port is actual DNS traffic or some other application that misuses this commonly open port (for example, peer-to-peer file transfer applications may use this port).

▼ To configure DNS Protocol options

1. In the properties of a custom Service you have created, click the **Select** button for the **Protocol** field and select **DNS**.
2. Switch to the **Protocol Parameters** tab and set the parameters for the Protocol:

Table 42.3 DNS Protocol Parameters

Option	Explanation
Deny DDNS updates	When set to **On**, the engine terminates DDNS update messages. When set to **Off**, the engine allows DDNS update messages to pass.
Deny DNS zone transfers	When set to **On**, the engine terminates DNS zone transfer messages. When set to **Off**, the engine allows DNS zone transfer messages to pass.

Table 42.3 DNS Protocol Parameters (Continued)

Option	Explanation
Enforce DNS protocol usage	When set to **On**, the engine terminates traffic that is not actually using the DNS protocol. When set to **Off**, the engine allows traffic to pass even if the traffic is not actually DNS-related.

3. Click **OK**.

Defining FTP Protocol Parameters

A File Transfer Protocol (FTP) session starts with a control connection (by default, TCP port 21), and the communications continue using a dynamically allocated port. The FTP Protocol Agent keeps track of the actual ports used, for example, to open them as needed so that the whole range of possible dynamic ports does not need to be allowed in the policy.

▼ To configure FTP Protocol options

1. In the properties of a custom Service you have created, click the **Select** button for the **Protocol** field and select **FTP**.
2. Switch to the **Protocol Parameters** tab.
3. Set the parameters for the Protocol. Options marked as **Firewall Only** should not be changed from their default values when used on IPS engines or Layer 2 Firewalls:

Table 42.4 FTP Protocol Parameters

Option		Explanation
Allow active mode	Yes	Server is allowed to open data connections to the client (according to information exchanged in the control connection).
	No	Server-initiated data connections are forbidden.
Allow passive mode	Yes	Client is allowed to open data connections to the server (according to information exchanged in the control connection).
	No	Client-initiated data connections are forbidden.
Allow related connections	On	Allows control connections to be opened with the data connection.
	Off	Disables the Protocol Agent.
Control data inspection mode (*Firewall only*)	Strict	If commands that do not comply with the RFC 959 FTP standard are used, the connection is dropped.
	Loose	The Protocol Agent tries to identify information for opening the data connection even if the communications do not strictly follow the FTP standards. Sometimes needed with non-standard FTP configurations.

Table 42.4 FTP Protocol Parameters (Continued)

Option	Explanation
Highest allowed source port for active data connection (*Firewall only*)	Enter a port value to limit the range of allowed source ports of data connection on the server. Value 0 for the lowest port means that the server always uses the port number immediately preceding the destination port. If the server uses a standard port, both the lowest and highest port number must be 0.
Lowest allowed source port for active data connection (*Firewall only*)	
Redirect connections to CIS (*Firewall only*)	Selects the Content Inspection Server (CIS) to which the connections are redirected. For more information, see External Content Inspection (page 809).

4. Click **OK**.

Defining GRE Protocol Parameters

The Generic Routing Encapsulation (GRE) protocol is a tunneling protocol that allows the encapsulation of network layer packets inside IP tunneling packets.

▼ To configure GRE Protocol options

1. In the properties of a custom Service you have created, click the **Select** button for the **Protocol** field and select **GRE**.
2. Switch to the **Protocol Parameters** tab.
3. Set the parameters for the Protocol:

Table 42.5 GRE Protocol Parameters

Option		Explanation
Apply Tunnel Rematch	On	Rematches the encapsulated payload inside the tunneling packet until the maximum rematch count defined in the engine properties is reached
	Off	Encapsulated packets are not rematched.
Tunnel IPv4 Protocol	On	Allows tunneling over IPv4.
	Off	Stops connections that are tunneled over IPv4.
Tunnel IPv6 Protocol	On	Allows tunneling over IPv6.
	Off	Stops connections that are tunneled over IPv6.

4. Click **OK**.

Defining H323 Protocol Parameters

H.323 consists of a series of different types of standards relating to video and audio services, real-time transport, control channels, security, etc. The H323 Protocol Agent tracks the H.323 connections, for example, to allow H.323 traffic through a firewall when NAT is used.

> **Note – T.120 connections, used for instance for file transfer and whiteboard drawing, are not allowed by the H.323 Protocol Agent. To allow T.120, use the H.323 Service Group or the T.120 Service element.**

▼ To configure the H323 Protocol options

1. In the properties of a custom Service you have created, click the **Select** button for the **Protocol** field and select **H323**.
2. Switch to the **Protocol Parameters** tab.
3. Set the parameters for the Protocol:

Table 42.6 H323 Protocol Parameters

Option		Explanation
Allow related connections	On	The Protocol Agent monitors the H.323 connection and allows the related connections in Access and NAT rules.
	Off	Disables the Protocol Agent.
Allow special logical channels through (No NAT)	Yes	Allows H.323 clients to open a special logical channel for audio and video without NAT.
	No	Special logical channels are not allowed.

4. Click **OK**.

Defining HTTP/HTTPS Protocol Parameters

You can configure parameters for the following two HTTP Protocol elements:

- The *HTTP* Protocol Agent can be used for redirecting traffic to an external content inspection server and to log the URLs from HTTP requests.
- The *HTTPS* Protocol Agent is for SSL encrypted HTTP.

▼ To configure the HTTP or HTTPS Protocol options

1. In the properties of a custom Service you have created, click the **Select** button for the **Protocol** field and select **HTTP** or **HTTPS**.
2. Switch to the **Protocol Parameters** tab.

3. Set the parameters for the Protocol (see the two tables below):

Table 42.7 HTTP Protocol Parameters

Option		Explanation
Logging of Accessed URLs	Yes	The URLs of sites that users access are included in generated log entries.
	No	URLs are not included in generated log entries.
Optimized server stream fingerprinting	Yes	When matching connections to the Inspection rules, the server stream matching is done only for patterns that are valid for the client's browser type and version.
	No	All server stream patterns are matched.
Redirect Connections to CIS (*Firewall Only*)		Selects the Content Inspection Server (CIS) to which the connections are redirected. See Getting Started with External Content Inspection (page 810).

Table 42.8 HTTPS Protocol Parameters

Option		Explanation
HTTPS Inspection Exceptions		Specifies the HTTPS Inspection Exceptions according to which traffic is decrypted and inspected or allowed to pass without decryption. See Getting Started with TLS inspection (page 794).
HTTPS Decryption and Inspection	Yes	Enables HTTPS decryption and inspection.
	No	HTTPS traffic is not decrypted for inspection.
Logging of Accessed URLs	Yes	The URLs of sites that users access are included in generated log entries. With HTTPS traffic, requires that the traffic is decrypted (see Getting Started with TLS inspection (page 794)).
	No	URLs are not included in generated log entries.
Optimized server stream fingerprinting	Yes	When matching connections to the Inspection rules, the server stream matching is done only for patterns that are valid for the client's browser type and version.
	No	All server stream patterns are matched.

4. Click **OK**.

Defining IPv4 Encapsulation Protocol Parameters

The IPv4 Encapsulation Agent provides protocol inspection for tunneled IPv4 traffic. The parameters define if IPv4 packets encapsulated in IPv6 packets are rematched to the policy.

▼ To configure the IPv4 Encapsulation Protocol options

1. In the properties of a custom Service you have created, click the **Select** button for the **Protocol** field and select **IPv4 Encapsulation**.
2. Switch to the **Protocol Parameters** tab.
3. Set the parameters for the Protocol:

Table 42.9 IPv4 Encapsulation Protocol Parameters

Option		Explanation
Apply Tunnel Rematch	On	Rematches the encapsulated IPv4 payload inside the IPv6 tunneling packet until the maximum rematch count defined in the engine properties is reached
	Off	Does not rematch encapsulated payload.
Next Ethernet Type		For information only. Shows the Ethernet frame type used for examining the encapsulated packet.

4. Click **OK**.

Defining IPv6 Encapsulation Protocol Parameters

The IPv6 Encapsulation Agent provides protocol inspection for tunneled IPv6 traffic. The parameters define if IPv6 packets encapsulated in IPv4 packets are rematched to the policy.

▼ To configure the IPv6 Encapsulation Protocol options

1. In the properties of a custom Service you have created, click the **Select** button for the **Protocol** field and select **IPv6 Encapsulation**.
2. Switch to the **Protocol Parameters** tab.
3. Set the parameters for the Protocol:

Table 42.10 Protocol Parameters

Option		Explanation
Apply Tunnel Rematch	On	Rematches the encapsulated IPv6 payload inside the IPv4 tunneling packet until the maximum rematch count defined in the engine properties is reached
	Off	Does not rematch encapsulated payload.
Next Ethernet Type		For information only. Shows the Ethernet frame type used for examining the encapsulated packet.

4. Click **OK**.

Defining MSRPC Protocol Parameters

The MSRPC (Microsoft RPC) Protocol Agent allows related connections for MSRPC applications, and handles also NAT modifications for communications between Microsoft Outlook clients and Microsoft Exchange servers.

The supported end-point mapper (EPM) connection method is TCP. By default, the Microsoft RPC/EPM service is available at port 135/tcp and the communications continue using a dynamically allocated port. The Protocol Agent keeps track of the actual ports used to dynamically allow the connections based on the port allocation. This removes the need to allow the full range of ports.

▼ To configure the MSRPC Protocol options

1. In the properties of a custom Service you have created, click the **Select** button for the **Protocol** field and select **MSRPC**.
2. (*Firewall only*) Switch to the **Protocol Parameters** tab and set the parameters for the Protocol:

Table 42.11 MSRPC Protocol Parameters

Option		Explanation
Allow MS Exchange Remote administration service	Yes	Allows remote administration of the Microsoft Exchange server through the Exchange System Attendant service.
	No	Prevents remote administration.
Allow MS Exchange user services	Yes	Allows the normal use of the Microsoft Outlook client; the Protocol Agent allows the use of Exchange Database service, Directory service, Information Store service, MTA service, and Store service.
	No	Prevents end-user services.
Allow any UUID in endpoint mapping	Yes	Allows other MSRPC requests in addition to Outlook/Exchange.
	No	The Service allows only Outlook/Exchange traffic.
Allow related connections	On	Allows responses sent by the end-point mapper (EPM) service.
	Off	Disables the Protocol Agent.
Allow other RPC traffic	Yes	Allows message types that are not supported by the Protocol Agent to bypass the control connection.
	No	Allows only supported message types (bind, bind ack, request, and response).

3. Click **OK**.

Defining NetBIOS Protocol Parameters

This Protocol Agent can be used to make NAT modifications in IP addresses transported in the payload of Windows NetBIOS Datagram Service connections through the Firewall, or for deep inspection on IPS engines and Layer 2 Firewalls.

▼ To configure the NetBIOS Protocol options

1. In the properties of a custom UDP Service you have created, click the **Select** button for the **Protocol** field and select **NetBIOS (UDP)**.
2. Switch to the **Protocol Parameters** tab.
3. Set the parameters for the Protocol Agent:

Table 42.12 NetBIOS Protocol Parameters

Option		Explanation
Make corresponding NAT modifications to payload	On	If inserted in a NAT rule, the addresses relayed in the NetBIOS communications are translated according to the NAT rule.
	Off	Only the IP addresses in packet headers are translated if inserted in a NAT rule.

4. Click **OK**.

Defining Oracle Protocol Parameters

Caution – The Oracle Protocol Agent is meant for cases in which TCP port 1521 is used only for negotiating the port number for Oracle database connections, and the port number for the actual connection is assigned dynamically. It must not be used in any other cases.

This Protocol Agent handles Oracle Transparent Network Substrate (TNS) protocol-based SQL*Net, Net7, and Net8 connections. It is meant only for non-SSL connections where the port number is assigned dynamically. If TCP port 1521 is used for the actual database connection, do not use a Service that contains this Protocol element, because this may consume excessive resources on the firewall and lead to performance problems (instead, use a Service that matches TCP port 1521 without any Protocol element).

If you plan to use NAT for Oracle connections, you must configure the *Oracle listener* so that the listener tells the client its original non-NATed IP address. This, and the Protocol Agent itself, is necessary only if the database is located on a different computer than the Oracle listener. The Oracle Protocol Agent does not modify payload data because the database service connections could go through a different route than the listener connection.

▼ To configure the Oracle Protocol options

1. In the properties of a custom Service you have created, click the **Select** button for the **Protocol** field and select **Oracle**.
2. Switch to the **Protocol Parameters** tab.

3. Set the parameters for the Protocol:

Table 42.13 Oracle Protocol Parameters

Option		Explanation
Allow related connections	On	Allows database connection based on information in the listener connection.
	Off	Disables the Protocol Agent.
Max. length allowed for one TNS packet		Enter the maximum amount of TCP payload data that each Oracle TNS packet is allowed to carry.
Netmask for allowed server addresses		Enter a netmask for limiting the allowed traffic. The value 255.255.255.255 allows the database connection only to the address in which the Oracle Listener service is located. The value 0.0.0.0 allows database connections to all addresses.
Set checksum to zero for modified TNS packets	Yes	Resets the header and packet checksums to zero when the Protocol Agent modifies the packet payload data.
	No	Checksums remain even when the packet is changed.

4. Click **OK**.

Defining Shell (RSH) Protocol Parameters

Remote Shell (RSH) is a widely used remote management protocol. This Protocol Agent manages Remote Shell connections and allows NAT modifications to the standard output (stdout) stream. It also manages RExec connections.

▼ To configure the Shell Protocol options

1. In the properties of a custom Service you have created, click the **Select** button for the **Protocol** field and select **Shell**.
2. Switch to the **Protocol Parameters** tab.
3. Set the parameters for the Protocol:

Table 42.14 Shell (RSH) Protocol Agent Parameters

Option		Explanation
Allow Related Connections	On	Standard error (stderr) stream is allowed through the firewall as a response to an RSH command.
	Off	Protocol Agent is disabled.

4. Click **OK**.

Defining SIP Protocol Parameters

The Session Initiation Protocol (SIP) agent can be used to handle multimedia connections that use SIP as their transfer protocol.

Using the agent allows SIP to be used across a firewall that uses NAT. SIP uses TCP or UDP port 5060 to initiate the connection, after which the traffic is allocated a dynamically assigned port. The Protocol Agent keeps track of the actual ports used, so that the range of dynamic ports does not need to be allowed in the firewall policy.

The SIP agent can be configured to force the client and/or server address used within the SIP transport layer to be used also for the media stream carried over SIP (by default set on for both client and server).

Note – Connections used for file transfer and whiteboard drawing are not allowed by the SIP Protocol Agent. Allow them in a different rule as necessary.

▼ To configure the SIP Protocol options

1. In the properties of a custom Service you have created, click the **Select** button for the **Protocol** field and select **SIP**.
2. Switch to the **Protocol Parameters** tab.
3. Set the parameters for the Protocol:

Table 42.15 SIP Protocol Parameters

Option		Explanation
Allow Related Connections (*Firewall Only*)	On	Allows SIP media connections based on the signalling connection.
	Off	Protocol Agent is disabled.
Enforce client side media	Yes	Requires that the media stream uses the same client-side address as the transport layer.
	No	Media stream can use any address.
Enforce server side media	Yes	Requires that the media stream uses the same server-side address as the transport layer.
	No	Media stream can use any address.
Maximum number of calls		The maximum number of calls allowed by the Access rule. If the value is 0, no limit is set for the number of calls.

4. Click **OK**.

Defining SMTP Protocol Parameters

On Firewalls, the SMTP Protocol Agent redirects connections to an external content inspection server (CIS) for screening. The SMTP client and the external content inspection server have to be in different networks to ensure that they always communicate through the firewall.

On IPS engines and Layer 2 Firewalls, this agent provides protocol inspection and deep inspection.

▼ To configure the SMTP Redirection options

1. In the properties of a custom Service you have created, click the **Select** button for the **Protocol** field and select **SMTP**.
2. (*For firewalls only*) Switch to the **Protocol Parameters** tab and set the parameters for the Protocol Agent:

Table 42.16 SMTP Protocol Agent Parameters

Option	Explanation
Redirect Connections to CIS	Selects the Content Inspection Server (CIS) to which the connections are redirected. See Getting Started with External Content Inspection (page 810).

3. Click **OK**.

Defining SSH Protocol Parameters

Secure Shell (SSH) is an encrypted remote use protocol. This Protocol Agent validates the communications to make sure the protocol used really is SSH. You can create custom SSH agents with different settings, if required. The SSH Agent validates SSHv1 only. This Protocol Agent is available on Firewalls, IPS engines, and Layer 2 Firewalls.

▼ To configure the SSH Protocol options

1. In the properties of a custom Service you have created, click the **Select** button for the **Protocol** field and select **SSH**.
2. (*For firewalls only*) Switch to the **Protocol Parameters** tab and set the parameters for the Protocol:

Table 42.17 SSH Protocol Agent Parameters

Option	Explanation
Bytes allowed from client before Server ID	Amount of data that the client is allowed to send to the server before the server sends its own identification string.
Bytes allowed from server before Client ID	Amount of data that the server can send to the client before the client sends its own identification string.
Bytes allowed from server before Server ID	Amount of data that the server can send to the client before the server sends its own identification string.

Table 42.17 SSH Protocol Agent Parameters (Continued)

Option		Explanation
Make protocol validation	On	Validates the SSH transfers according to the parameters defined in this dialog.
	Off	Protocol Agent is disabled.

3. Click **OK**.

Defining SunRPC Protocol Options

There are both UDP and TCP based Protocol Agents for Sun Remote Procedure Call (RPC) protocol. On the firewall, these agents only assist the firewall in *Portmapper* connections. They make the handling of RPC program numbers used in the Access rules more rapid. On IPS engines and Layer 2 Firewalls, these protocol agents provide deep inspection.

Note – The Protocol Agent is meant only for Portmapper connections. Allow other RPC services using Service elements without the Protocol Agent.

The Portmapper Protocol Agents collect information about RPC services by interpreting the GET PORT and DUMP PORTS requests and their respective answers. All information it collects is stored in the Portmapper cache.

When the packet filter needs to evaluate RPC matches, it consults the Portmapper cache to check if the destination of the packet has the appropriate service defined in the rule. If the cache does not have the requested information available, the packet under evaluation is not let through and a query is sent to the destination host for RPC information. The information received is stored in cache.

When a security policy contains a rule referring to RPC program number, it is essential to pay attention to the structure of the security policy. We recommend you to follow these precautions with the RPC protocol:

- Attach Portmapper Protocol Agent only to Portmapper connections passing through the firewall.
- Allow the firewall engine to send RPC queries as explained below.
- Optimize the structure of your security policy (see https://my.stonesoft.com/support/document.do?product=StoneGate&docid=1313 for more information).

RPC queries are sent from the firewall to TCP port 111 of the external host. You can use either the sunrpc services configured for TCP and UDP, or the Portmapper combined service with both sunrpc services. We recommend you to insert the following rule that allows connections to allow such connection without the Protocol Agent above any other Portmapper rules:

Table 42.18 Rule for RPC Queries

Source	Destination	Service	Action
Firewall engine IP address (NDIs on clusters)	Any	SunRPC (TCP) SunRPC (UDP)	Allow

▼ To configure the SunRPC Proxy options

1. In the properties of a custom Service you have created, click the **Select** button for the **Protocol** field and select **SunRPC (**[TCP|UDP]**)**.
2. (*Firewall only*) Switch to the **Protocol Parameters** tab and set the parameters for the Protocol Agent:

Table 42.19 SunRPC Protocol Parameters

Option		Explanation
Learn RPC program number to port mapping for future RPC service matches	Yes	Protocol Agent is enabled.
	No	Protocol Agent is disabled.

3. Click **OK**.

Defining TCP Proxy Protocol Parameters

The TCP Protocol Agent is a proxy agent that allows closing connections after a set amount of idle time.

Certain TCP based applications do not properly handle closing of connections but instead leave them open for a long period of time, unnecessarily consuming resources. For such situations, the TCP proxy agent can be used to actively close the connections after a certain idle time. The TCP Proxy Agent can also close a connection if connection closing initiated by one of the communicating parties does not complete in a timely manner.

Apply this Protocol Agent strictly to those TCP connections that require this feature. Proxy operations use more resources than normal TCP connection handling and therefore increase the firewall's load.

▼ To configure the TCP Proxy options

1. In the properties of a custom Service you have created, click the **Select** button for the **Protocol** field and select **TCP Proxy**.
2. (*Firewall only*) Switch to the **Protocol Parameters** tab and set the parameters for the Protocol Agent:

Table 42.20 TCP Proxy Protocol Parameters

Option	Explanation
Abort on close	Timeout in seconds for aborting a connection counted from when the connection closing is initiated by one of the communicating parties. The connection is aborted by sending TCP Reset packets to the unresponsive endpoint. Setting this value to 0 disables this timeout (connections are left open).

Table 42.20 TCP Proxy Protocol Parameters (Continued)

Option		Explanation
Idle Timeout		Timeout in seconds for closing a connection after the latest transmission. Setting this value to 0 disables this timeout (connections are left open).
Use Proxy	On	Protocol Agent is enabled.
	Off	Protocol Agent is disabled.

3. Click **OK**.

Defining TFTP Protocol Parameters

Trivial File Transfer Protocol (TFTP) Protocol Agent transfers data using dynamically assigned ports. The TFTP protocol (RFC 1350) does not limit the port range that can be used. This Protocol Agent is available on Firewalls, IPS engines, and Layer 2 Firewalls. The TFTP Protocol Agent supports NAT operations (Firewall only).

A TFTP Agent is attached to a UDP connection established between the client and the server. The client opens the control connection from a dynamically selected source port to the fixed destination port 69/udp on the server. A separate UDP data connection is established between the client and the server after the client has sent a “read” of “write” command to the server. The server opens a data connection from a dynamic source port to the client’s destination port, which is same as the one used as the source port of the control connection.

The parameters in this protocol are for firewalls only.

▼ To configure the TFTP Protocol options

1. In the properties of a custom Service you have created, click the **Select** button for the **Protocol** field and select **TFTP**.
2. (*Firewall only*) Switch to the **Protocol Parameters** tab and set the parameters for the Protocol:

Table 42.21 TFTP Protocol Parameters

Option		Explanation
Allow ‘read’	Yes	Allows file transfer from server to client (downloads)
	No	Downloads are not allowed.
Allow Related Connections	On	Allows control connections to be opened with the data connection.
	Off	Protocol Agent is disabled.
Allow ‘write’	Yes	Allows file transfer from client to server (uploads).
	No	Uploads are not allowed.

Table 42.21 TFTP Protocol Parameters (Continued)

Option		Explanation
Log filenames and paths	Yes	Names of transferred files and their paths are included in generated log entries.
	No	File and path information is not available in logs.

3. Click **OK**.

CHAPTER 43

DEFINING SITUATIONS

Situations are the central element in Inspection Policies. Situations define the patterns that recognize interesting events in the traffic.

The following sections are included:

Getting Started With Situations

Prerequisites: None

Situations are elements that define deep inspection patterns.

What Situations Do

Situations define a pattern in traffic that the system looks for. he patterns and events are defined by selecting a *Context* for the Situation. The Context contains the information on the traffic to be matched, and the options you can set for the matching process.

The Inspection Policy defines how the Situations are matched to traffic and what kind of a action the system takes when a match is found.

Correlation Situations are Situations that group together event data to find patterns in that data.

Situations also provide a description that is shown in the logs, and a link to relevant external information (CVE/BID/MS/TA) in the form of a *Vulnerability* element attached to the Situation.

Situations use a special grouping system called *Tags*. The Tag elements are shown as branches in the Situations tree and they can be used in policies to represent all Situations that are associated with that Tag.

Example **Using the Tag “Windows” in a rule means that the rule matches to all Situations in the system that are classified as concerning Windows systems.**

Associating a Situation with a *Situation Type* includes the Situation in the Rules tree in the Inspection Policy, which is grouped according to the Situation Type.

Limitations

Depending on the Usage Context properties of a Correlation Situation, correlation may be done only on the Security Engine, only on the Log Server, or on both the Security Engine and the Log Server. When correlation is done only on the Security Engine, the Correlation Situation only matches when all correlated events are detected by the same Security Engine. The following table lists the Usage Contexts for predefined Correlation Situations:

Table 43.1 Usage Contexts for Predefined Correlation Situations

Correlation Context	Usage Context
Compress	Engine Only
Count	Engine Only
Group	Engine Only
Match	Engine Only
Sequence	Log Server Only

By default, correlation is done on both the Security Engine and the Log Server for custom Correlation Situations. For more information about selecting the Usage Context in custom Correlation Situations, see Defining Context Options for Correlation Situations (page 762).

Configuration Overview

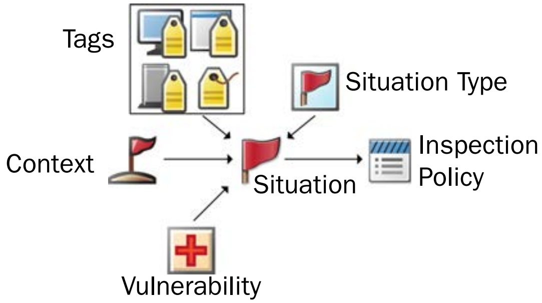

Illustration 43.1 Elements in the Configuration

1. Create a new Situation element as explained in Creating New Situation Elements.
2. Give the Situation a Context, and fill in the context information according to the patterns in traffic that you want to match as explained in Defining Context Options for Situations (page 759) or Defining Context Options for Correlation Situations (page 762).
3. *(Optional)* Associate the Situation with the relevant Tags as explained in Defining Tags for Situations (page 767).
4. *(Optional)* Associate the custom situation description with a relevant Vulnerability as explained in Working With Vulnerabilities (page 769).
5. Use the Situation in the Inspection Policy. See Editing Inspection Policies (page 699).

Creating New Situation Elements

Prerequisites: None

The predefined Situation elements in your system are updated through dynamic update packages. You can also create new Situation elements to fine-tune the patterns that the Security Engines look for in the traffic. There are two ways to create new Situation elements:

- You can import existing Snort rules to automatically create Situation elements and an Inspection Policy as explained in Importing Snort Rules Libraries (page 709).
- You can manually define a custom Situation element.

Note – Creating custom Situations requires you to have a basic understanding of the protocols involved and a clear picture of the patterns of traffic that you want to look for.

▼ To create a new Situation

1. Select **Configuration→Configuration→Security Engine** from the menu. The Security Engine Configuration view opens.
2. Expand the **Other Elements** branch and right-click the **Situations** branch in the tree view.

3. Select either **New→Situation** or **New→Correlation Situation**. The Situation Properties dialog opens.

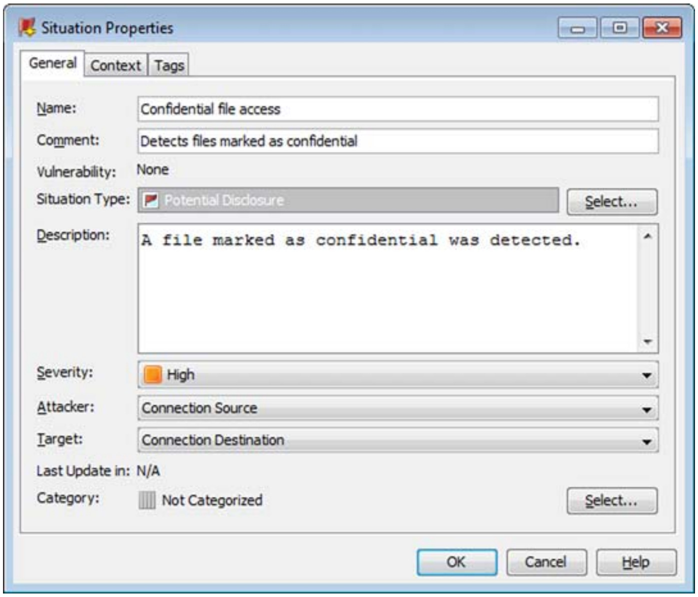

Name: Confidential file access

Comment: Detects files marked as confidential

Vulnerability: None

Situation Type: Potential Disclosure

Description: A file marked as confidential was detected.

Severity: High

Attacker: Connection Source

Target: Connection Destination

Last Update in: N/A

Category: Not Categorized

4. Enter a unique **Name** and optionally a **Comment** to describe the element for your reference.
   - With the exception of whitelisted URLS in URL Filtering, Situations are identified in the system only by the element name. Creating multiple Situations that match the same pattern can make the policy difficult to read and manage.
   - The comment is only shown in the element properties and in the Info panel.
5. (*Optional*) Click **Select** and select the Situation Type with which to associate this Situation. You can only select on Situation Type for each Situation. The Situation Type specifies the branch of the Rules tree under which the Situation is included.
6. Use the **Description** field to describe the traffic pattern that the Situation represents. This description is shown, for example, in log entries.
7. Select a **Severity** for the Situation. The Severity is shown in the logs and can be used in Alert Policies as a criterion for alert escalation.
8. (*Optional*) Select how the **Attacker** and **Target** are determined when the Situation matches. This information is used for blacklisting and in log entries.
   - **None** does not define the Attacker and Target information, so blacklisting entries cannot be created using the **Attacker** and **Target** options.
   - **IP Source** and **IP Destination** are the IP addresses of the (last) packet that triggers the Situation. Since the packet can be a request or a reply, make sure to select the option correctly based on the pattern that the situation detects to avoid reversed labeling.
   - **Connection Source** and **Connection Destination** IP addresses depend on which host opened the connection and provide a constant point of reference to the client and server in the communications.
9. (*Optional*) Click **Select** to select a category for the correlation situation.

What's Next?

▶ Depending on the type of Situation element, proceed to Defining Context Options for Situations (page 759) or Defining Context Options for Correlation Situations (page 762).

Defining Context Options for Situations

Prerequisites: Creating New Situation Elements

The Context information gives the Situation the information on which kind of patterns you want it to match in the traffic.

Example **You want to look for a certain character sequence in an HTTP stream from the client to the server.**

The Context gives you a set of options or a field for entering a regular expression that you can use to define the pattern you want to look for in the traffic. See Regular Expression Syntax (page 1179).

> **Note – With the exception of whitelisted URLs in URL Filtering, Situations are identified in the system only by the element name. Avoid matching the same pattern in different Situation elements. Situations with duplicate patterns can make the policy difficult to read and manage.**

These instructions are for Situations. For Correlation Situations, see Defining Context Options for Correlation Situations (page 762).

▼ To define the Context options for a Situation

1. In the Situation properties, switch to the **Context** tab.
2. Click the **Select** button. The available Context categories are shown.
3. **Select** the Context you want to associate with this Situation. The options for the selected Context are added to the Situation Properties.
   - If you are editing a Situation that was automatically created while importing Snort rules, the Context is automatically selected based on the original Snort rule. You can optionally change the Context.

> **Note – The details related to the Contexts in your system may be different from what is described here because the Contexts may have been updated through dynamic update packages after this guide was published. Read the release notes of each update package you import to see which elements are affected.**

What's Next?

- ▶ To fill in HTTP URL Filter options, see Defining HTTP URL Filter Options (page 760).
- ▶ To fill in scan detection options, see Defining Port/Host Scan Detection Options (page 760).
- ▶ For many Contexts, you must type in a regular expression, see Regular Expression Syntax (page 1179).
- ▶ In other cases, open the Properties dialog for the Context element for more information. When you have configured the necessary context information, proceed to Adding Tags to One Situation at a Time (page 768).

Related Tasks

- ▶ Defining Context Options for Correlation Situations (page 762)

Defining HTTP URL Filter Options

HTTP URL Filtering Situations allow you to define lists of URLs that blacklist (block) access to the specified URLs. When HTTP URL Filtering Situations are used in combination with category-based web filtering, they can be used to whitelist (allow) individual URLs that are included in an otherwise blocked category. URL whitelisting only affects URL-based filtering, and it does not exclude the traffic from other inspection checks. To be able to filter URLs in HTTPS traffic, you must first decrypt the traffic. See Getting Started with TLS inspection (page 794).

▼ To configure website access control

1. Create a new Situation element with the appropriate basic properties. See Creating New Situation Elements (page 757).
2. Switch to the **Context** tab.
3. Click the **Select** button. The Select Context dialog opens.
4. Browse to **Application Protocols→HTTP→HTTP URL Filter** and click **Select**.
5. Click the **Add** button below the entry panel. You can add up to 20 URLs per Situation.
6. Double-click the newly added row and enter the URL without specifying the protocol, for example, `www.example.com`.
   - Each URL can be a maximum of 64 characters long.
   - To block or allow access to an exact URL, type: `www.example.com/index.html`
   - To block or allow access to a domain, type: `example.com`
   - An asterisk (*) can be used at the beginning of the URL. For example: `*.example.com` (matches requests for `www.example.com` and `extranet.example.com` but not the shorter variant `example.com`).
7. Add Tags as necessary to organize this Situation and click **OK**.

Defining Port/Host Scan Detection Options

The Scan Started Event Context options can be used to define the bounds for normal traffic, allowing action to be taken when the values are exceeded. The table below explains the options for the Scan Started Event Context.

Table 43.2 Scan Detection Parameters

Parameter	Description
Port scan start period (seconds)	Port scan is reported if any of the thresholds is exceeded within this time limit.
Port scan idle timeout (seconds)	Port scan is assumed finished if the originator makes no scan attempts within this time limit.
Port scan status delay (seconds)	Defines how often an interim status of ongoing port scans is reported.

Table 43.2 Scan Detection Parameters (Continued)

Parameter	Description
Maximum normal TCP connections	Defines how many TCP connections that proceed normally according to the protocol definitions are allowed before action is taken.
Maximum allowed open TCP ports	Number of SYN+ACK replies that are allowed during the tracking before action is taken.
Maximum unreplied TCP connections	Number of unreplied TCP connections that are allowed during the tracking before action is taken.
Maximum allowed closed TCP ports	Number of RST replies that are allowed during the tracking before action is taken.
Maximum TCP segments with no SYN or ACK	Number of TCP segments with no SYN or ACK flag that are allowed before action is taken.
Wait time for TCP connections	Delay for judging if a TCP connection is a successful port scan or an unreplied connection attempt.
Maximum UDP packet destinations	Number of UDP destinations that are allowed per host during the tracking before action is taken.
Maximum bidirectional UDP transfers	Number of bidirectional UDP transfers that are allowed per host during the tracking before action is taken.
Maximum unidirectional UDP transfers	Allowed number of UDP destinations that have not replied or have replied with ICMP error during the tracking period before action is taken.
Maximum allowed closed UDP ports	Number of ICMP Port Unreachable replies that are allowed per host during the tracking before action is taken.
Maximum ICMP requests per host	Number of ICMP request destinations that are allowed per host during the tracking before action is taken.
Maximum unreplied ICMP request destinations	Number of ICMP request destinations that have not replied during the tracking that are allowed per host before action is taken.
Maximum ICMP Echo Request destinations	Number of ICMP Echo Request (ping) destinations that are allowed per host during the tracking before action is taken.
Maximum unreplied ICMP Echo Requests	Number of ICMP Echo Request (ping) destinations that have not replied during the tracking that are allowed per host before action is taken.
Maximum ICMP Timestamp Request destinations	Number of ICMP Timestamp Request destinations that are allowed per host during the tracking before action is taken.
Maximum unreplied ICMP Timestamp Requests	Number of ICMP Timestamp Request destinations that have not replied during the tracking that are allowed per host before action is taken.
Maximum ICMP Netmask Request destinations	Number of ICMP Netmask Request destinations that are allowed per host during the tracking before action is taken.

Table 43.2 Scan Detection Parameters (Continued)

Parameter	Description
Maximum unreplied ICMP Netmask Requests	Number of ICMP Netmask Request destinations that have not replied during the tracking that are allowed per host before action is taken.

Defining Context Options for Correlation Situations

Prerequisites: Creating New Situation Elements

Correlation Situations are used by Security Engines and Log Servers to conduct further analysis of detected events. Correlation Situations do not handle traffic directly. Instead they analyze the events generated by matches to Situations found in traffic.

Depending on the Usage Context properties of the Correlation Situation, correlation may be done only on the Security Engine, only on the Log Server, or on both the Security Engine and the Log Server. By default, correlation is done on both the Security Engine and the Log Server for custom Correlation Situations. For more information about selecting the Usage Context, see the Context-specific sections below.

▼ To set Context information for Correlation Situations

1. Switch to the **Context** tab.
2. Click the **Select** button. The Select Context dialog opens.
3. Select the Context you want to associate to this Correlation Situation:

Context	Description
Compress	Combines repeated similar events into the same log entry, reducing the amount of data that is sent to the Log Server. If you select this context, continue in Configuring Compress Contexts (page 763).
Count	Finds recurring patterns in traffic by counting the times certain Situations occur within a defined time period. The Situation matches if the threshold values you set are exceeded. If you select this context, continue in Configuring Count Contexts (page 764).
Group	Finds patterns in traffic by following if all events in the defined set of Situations match at least once in any order within the defined time period. If you select this context, continue in Configuring Group Contexts (page 765).
Match	Allows filtering event data using filters. If you select this context, continue in Configuring Match Contexts (page 766).
Sequence	Finds patterns in traffic by following if all events in the defined set of Situations match in a specific order within a defined time period. If you select this context, continue in Configuring Sequence Contexts (page 766).

Configuring Compress Contexts

Caution – Be careful when defining the Compress Context options. You must make sure that all event data you compress are truly parts of the same event. Otherwise you risk losing valuable event information.

▼ To configure the Compress context parameters

1. Browse to the Situation(s) you want to compress in the left panel of the dialog and drag and drop them into the **Correlated Situations** field.
2. Enter the **Time Window Size** in seconds. The matches to the Situation(s) selected are combined to a common log entry when they are triggered within the defined time period.
3. Enter the **Events per Window**. This defines the maximum number of events that are forwarded within the Time Window defined.
4. Select one of the following options for **Log Fields Enabled**:

Option	Description
Select	Includes matching items. Events triggered by the selected Situations are considered the same when the values are identical in the Log Fields you place in **Log names**.
Ignore	Excludes matching items. Events triggered by the selected Situations are considered the same except when the values are identical in the Log Fields you place in **Log names**.

5. Click the **Log names** field and drag and drop the log field(s) that are used by the matching option you selected in the previous step.
6. Select a **Location** to determine the execution order of this Compress operation in relation to other Compress operations. Operations that share the same Location are executed in parallel; each compress operation receives the same events as the other compress operations in the same Location. The next Location receives only the events that are left after the compression. The possible Locations are:

Location	Description
Very Early	Executes the operation before “Early”. Using Very Early can improve performance if the engine load is constantly high.
Early	Executes the operation before any other correlation (non-compress) operations. After this Location, remaining events are processed by other correlation operations before applying further Compression operations. Using Early can improve performance if the engine load is constantly high.
Late	Executes the operation after all other correlation (non-compress) operations.
Very Late	Executes the operation after “Late”.

Caution – Be careful when using the Early or Very Early Locations. The compression may affect the other types of correlation tasks.

7. Click **Edit** and define a **Compress filter** in the Local Filter Properties dialog as instructed in Creating and Editing Local Filters (page 187). The filter is used for filtering data to be included in the compression.
8. Make sure **Engine Only** is selected as the **Usage Context**.

> **Note – The purpose of the Compress context is to reduce the amount of data that is sent to the Log Server. Including the Log Server in the Usage Context of a Compress context actually increases the amount of data that is sent to the Log Server.**

Configuring Count Contexts

▼ To configure the Count context parameters

1. Browse to the Situation(s) you want to count in the left panel and drag and drop them to the **Correlated Situations** field.
2. Enter the **Time Window Size** in seconds. All events must occur during this length of time for the Correlation Situation to match.
3. Enter the **Alarm Threshold** number. This is the number of times that the event must occur for the Correlation Situation to match.
4. Select one of the following options for **Log Fields Enabled**:

Option	Description
Select	Includes matching items. Events triggered by the selected Situations are considered the same when the values are identical in the Log Fields you place in **Log names**.
Ignore	Excludes matching items. Events triggered by the selected Situations are considered the same except when the values are identical in the Log Fields you place in **Log names**.

5. Click the **Log names** field and drag and drop the log field(s) that are used by the matching option you selected in the previous step.
6. (*Optional*) Select the Usage Context to define where correlation is done.

> **Note – If you select a Usage Context that does not include the Log Server, events only match if they are all detected by the same Security Engine or Security Engine Cluster.**

Configuring Group Contexts

The Group context has a table that allows you to define local filters and log fields for selecting which details are considered when events are grouped. In this context, the order in which the events occur is not considered. If you would like the order of the events to matter, use the Sequence context instead. See Configuring Sequence Contexts (page 766).

▼ To configure the Group context parameters

1. Double-click the **Event Match** cells for each **Member** cell and define a local filter as instructed in Creating and Editing Filter Elements (page 190). The local filter selects the events for examination.
   - You can add and remove members using the buttons to the right (to remove a member, first select a cell within that member's column).
2. Double-click the **Needed Number** cell of each member and enter the number of occurrences of the Event Match that are required for the events to be grouped.
3. Drag and drop a log field from the left panel to the **Binding** cell. You can add and remove Bindings by using the buttons to the right.
   - You must select log fields that receive the same type of value for each Binding row. The grouping is done per Binding row if the selected events have identical values in the Log Fields you select.

Example A grouping is possible if Member #1 includes a source IP address field for Binding #1 and Member #2 includes a destination IP address field for Binding #1, since both fields receive an IP address as their value.

4. Drag and drop the relevant Situation(s) to the **Correlated Situations** field.
5. Select whether you want to **Keep and Forward Events**:
   - **Yes**: the Situation examines the events and triggers the desired response defined in the Inspection Exception but does not combine the matching events into a single event. All the individual events are still available for further inspection, even though they have already triggered a response.
   - **No**: the Situation combines the matching events together, so that only the response defined for the one matching Inspection Exception is triggered, and no further processing is done on the individual events.
6. Enter the **Time Window Size** in seconds. All events must occur during this length of time for the Correlation Situation to match.
7. Select whether you want to trigger **Continuous Responses**:
   - **Yes**: the Security Engine responds as defined in the Inspection Exception to each occurrence of the defined event within the selected Time Window, or several times if the events occur frequently.
   - **No**: the Security Engine responds only to the first occurrence of the defined event within the selected Time Window (only once for each time period, even if the events occur more frequently).
8. (*Optional*) Select the Usage Context to define where correlation is done.

Note – If you select a Usage Context that does not include the Log Server, events only match if they are all detected by the same Security Engine or Security Engine Cluster.

Configuring Match Contexts

▼ To configure the Match context parameters

1. Browse to the Situation(s) you want to count in the left panel of the dialog and drag and drop them into the **Correlated Situations** field.
2. Click **Edit** and define a local filter as instructed in Creating and Editing Local Filters (page 187).
3. (*Optional*) Select the Usage Context to define where correlation is done.

Note – If you select a Usage Context that does not include the Log Server, events only match if they are all detected by the same Security Engine or Security Engine Cluster.

Configuring Sequence Contexts

The Sequence context has a table you can use to define, in order from left to right, the events that comprise a sequence. This allows detecting events such as when a request from a client triggers one pattern and the response from a server triggers a second pattern.

The table has gray and white cells; white cells must be filled, gray cells are left empty. If you would like to match events regardless of the order in which they occur, use the Group context instead. See Configuring Group Contexts (page 765).

▼ To configure the Sequence context parameters

1. Double-click the **Entry to #1** cell and define a local filter as instructed in Creating and Editing Local Filters (page 187).
2. Drag and drop a log field from the left panel to the **Binding #1** cell in the **Exit from #1** column. You can add and remove Bindings using the buttons to the right.
   - You must select log fields that receive the same type of value for each Binding row. The sequencing is done per Binding row if the selected events have identical values in the Log Field(s) you define.

Example **A sequence can be detected if Exit from #1 includes a source IP address field for Binding #1 and Entry to #2 includes a destination IP address field for Binding #1, since both fields may receive an IP address as their value.**

3. After adding the necessary number of Binding rows, define the Event Match field for Entry to #2 in the same way as for Entry to #1.
4. Define the **Binding** rows for the Entry to #2 column. As explained above, the type of value this field may contain must be the same as in other columns on the same Binding row.
5. (*Optional*) To make a more complex sequence, add a stage using the buttons to the right. A stage adds the necessary Exit column for the previous stage and an Entry column for a new event you can add to the sequence. The further sequences are filled in the same way as described above.
6. When you are finished defining the sequence, drag and drop the relevant Situation(s) in the **Correlated Situations** field below the table.
7. Select whether you want to **Keep and Forward Events**:
   - **Yes**: the Situation examines the events and triggers the desired response defined in the Inspection Exception but does not combine the matching events into a single event. All

the individual events are still available for further inspection, even though they have already triggered a response.

- **No**: the Situation combines the matching events together, so that only the response defined for the one matching Inspection Exception is triggered, and no further processing is done on the individual events.

8. Enter the **Time Window Size** in seconds. All events must occur during this length of time for the Correlation Situation to match.
9. Select one of the following options as the **Usage Context**:
   - Engine and Log Server.
   - Log Server Only.

Note – If you select a Usage Context that does not include the Log Server, events only match if they are all detected by the same Security Engine or Security Engine Cluster.

Defining Tags for Situations

Prerequisites: Creating New Situation Elements

Proceed to one of the following topics:

- Creating a New Tag
- Adding Tags to One Situation at a Time (page 768)
- Adding Tags to Several Situations at Once (page 768)
- Removing Tags from Situations (page 769)

Creating a New Tag

Tag elements collect together several Situations that have something in common (for example, Situations that detect threats against a particular Operating System). Tags are shown as branches in the Situations tree. Tag elements help you organize the tree, and you can use the Tags in the Inspection Policy to easily match the rule to all Situations that reference the Tag.

▼ To Create a new Tag

1. Select **Configuration→Configuration→Security Engine** from the menu. The Security Engine Configuration view opens.
2. Expand the **Other Elements→Situations** branch of the element tree.
3. Right-click the **By Tag** branch under Situations and select one of the following from the menu that opens:
   - **New→Application**.
   - **New→Hardware**.
   - **New→Operating System**.
   - **New→Situation Tag**.
4. Enter a **Name** and optionally a **Comment** for the new Tag and click **OK**. The new Tag is added to the tree under the main category you selected in the previous step.

What's Next?

▶ Proceed to Adding Tags to One Situation at a Time (page 768) or Adding Tags to Several Situations at Once (page 768).

Adding Tags to One Situation at a Time

When you are editing the Situation properties, you can add tags as explained below. If you want to add the same tags to several Situations at once, see Adding Tags to Several Situations at Once.

▼ To add Tags in the Situation properties

1. In the Situation properties, switch to the **Tags** tab.
2. Click **Add Tags** and select a Tag type from the list that opens. The Add dialog opens.
3. Select the Tags you want to use with this Situation and click **Select**.
4. Repeat the steps to add more Tags of any Type.
5. Click **OK** to confirm the Situation properties change.

If you just created a new Situation, you may want to associate appropriate Vulnerability information to it, see Working With Vulnerabilities (page 769).

Adding Tags to Several Situations at Once

You can add any number of Tags to each Situation.

▼ To add tags from the element tree

1. Select **Configuration→Configuration→Security Engine** from the menu. The Security Engine Configuration view opens.
2. Expand the **Other Elements→Situations** branch of the element tree and browse to the Situations you want to tag.
3. Ctrl-select or Shift-select the Situations that you want to tag.
4. Right-click one of the selected Situations and select **Add Tag**. The tag types are shown in a submenu.

Illustration 43.2 Adding Tags to Several Situations at Once

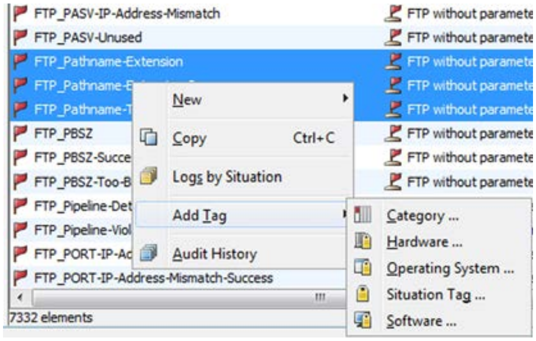

5. Select the type of Tag you want to add. The Add dialog that opens
6. Select the Tags you want to attach to the Situation(s).

Removing Tags from Situations

The default Tags in System Situations (provided in Update Packages) cannot be removed. You can only remove Tags that administrators have added in your system.

▼ To remove a Tag association from a Situation

1. Select **Configuration→Configuration→Security Engine** from the menu. The Security Engine Configuration view opens.
2. Expand the **Other Elements→Situations** branch of the element tree and browse to the Situations you want to edit.
3. Right-click the Situation and select **Properties**.
4. Switch to the **Tags** tab.
5. Right-click the Tag you want to remove and select **Remove**.
6. Click **OK**.

Working With Vulnerabilities

Prerequisites: None

Vulnerabilities provide a short description of the event that has matched and a reference to external vulnerability information (CVE/BID/MS/TA). When a Situation element refers to a Vulnerability, the vulnerability information is included in the log entries generated when the Situation matches traffic. You can also use the vulnerability IDs in the element search to find Situation and Vulnerability elements in the system that refer to a particular vulnerability.

Vulnerability information is included in dynamic update packages, so Situations provided by Stonesoft are already linked to a Vulnerability when appropriate. You can associate Situations with an existing Vulnerability or add a custom Vulnerability element.

Related Tasks

- ▶ Creating New Vulnerability Elements
- ▶ Associating Vulnerabilities With Situations (page 770).

Creating New Vulnerability Elements

▼ To create a new Vulnerability element

1. Select **Configuration→Configuration→Security Engine** from the menu. The Security Engine Configuration view opens.
2. Expand the **Other Elements→Situations** branch of the element tree.
3. Right-click the **By Vulnerability** branch in the tree view and select **New→Vulnerability** from the menu. The Vulnerability Properties dialog opens.
4. Give the Vulnerability a descriptive name and optionally a comment.
   - The Comment is not shown in the Logs view. Use the **Description** field to enter information to be shown in the logs.

5. To create a reference to external vulnerability information, select one or more reference systems in the **Reference System** section and enter the **ID** this vulnerability has in that system:
   - **Mitre**: vulnerability ID format is CVE-YYYY-XXXX
   - **SecurityFocus**: vulnerability ID format is BID-XXXXX
   - **Microsoft**: vulnerability ID format is MSYY-XXX
   - **Us-Cert**: vulnerability ID format is TAYY-XXXX
6. After you have entered the vulnerability ID, click **Show** next to the ID field to view the information about the vulnerability in the reference system.
7. Type or copy-paste a short description of what the vulnerability is about into the **Description** field.
8. Under **Situations**, browse to the correct Situation elements, select them (one or several at a time) and click the **Add** button to associate them with this Vulnerability. The selected Situations are added to the **Content** field on the right.
9. When you are finished adding Situations, click **OK**. The selected Situations are now associated with this vulnerability, and a link to this Vulnerability is added on the Situations' properties dialog.

Associating Vulnerabilities With Situations

▼ To add/remove a Vulnerability from Situation elements

1. Select **Configuration→Configuration→Security Engine** from the menu. The Security Engine Configuration view opens.
2. Expand the **Other Elements→Situations→By Vulnerability**.
3. Right-click the correct Vulnerability in the tree and select **Properties** from the menu. The Vulnerability Properties dialog opens.
4. Select the Situation elements (one or several at a time) and click the **Add** or **Remove** button to change the selection in the **Content** field on the right.
5. When you are finished adding Situations, click **OK**. The selected Situations are now associated with this vulnerability, and a link to this Vulnerability is added on the Situations' properties dialog.

CHAPTER 44

WORKING WITH APPLICATIONS

Application elements collect combinations of identified characteristics and detected events in traffic to dynamically identify traffic related to the use of a particular application.

The following sections are included:

Getting Started With Applications

Prerequisites: None

Applications are elements that provide a way to dynamically identify traffic patterns related to the use of a particular application.

What Applications Do

Applications allow you to more flexibly identify traffic beyond specifying a network protocol and ports for TCP and UDP traffic with a Service element. Applications are matched against the payload in the packets.

Limitations

Applications are not supported on Master Engines or Virtual Security Engines.

What Do I Need to Know Before I Begin?

- There are several predefined Application elements available that define the criteria for matching commonly-used applications. No configuration is required to be able to use Application elements in Access rules.
- Predefined TLS Match elements are used in the properties of some predefined Application elements to allow the Application to match the use of the TLS (transport layer security) protocol in traffic.
- You cannot edit the predefined Application elements. However, Access rules can override the properties of a predefined Application element.
- Creating new Application elements requires detailed knowledge of the application(s) you want to detect and the traffic patterns related to their use. In most cases, creating new Application elements is not recommended.

Configuration Overview

1. (*Optional*) Create TLS Match elements to override the properties of predefined Applications as explained in Creating TLS Matches.
2. Use the Application in the Access rules. See Creating Access Rules for Application Detection (page 774).

Creating TLS Matches

Prerequisites: None

TLS Matches define matching criteria that is checked against the server certificate in TLS connections. In addition to the predefined TLS Matches, you can optionally define your own TLS Matches to override the properties of an Application in an Access rule.

▼ To create a TLS Match

1. Select **Configuration→Configuration→Security Engine**. The Security Engine Configuration view opens.
2. Expand the **Other Elements** branch.
3. Right-click **TLS Matches** and select **New TLS Match**. The TLS Match Properties dialog opens.
4. Enter a unique **Name**.
5. (*Optional*) Select **Deny Decrypting** to prevent connections from being decrypted for inspection. See Getting Started with TLS inspection (page 794) for more information about inspection of TLS traffic.

> **Note – Selecting Deny Decrypting prevents the Application where the TLS Match is used from being identified if the traffic is encrypted. If you want to specify that decryption is not necessary for identifying the Application, use the Application Identifiable by TLS Match Alone option in the Application Properties instead.**

6. Select the **Match Certificate Validation** to define whether the server certificate validity is checked, and what to match:

Setting	Description
Validation Succeeded	Checks the server certificate validity and matches server certificates that are considered valid.
Validation Failed	Checks the server certificate validity and matches server certificates that are considered invalid.
No Validation	The server certificate validity is not checked. Matches any certificate.

7. Configure the additional settings depending on the Match Certificate Validation you selected:

Match Certificate Validation	Configuration
Validation Succeeded	Click **Add** and enter the fully qualified domain name to match in the server certificate.
Validation Failed	(*Optional*) Select the specific types of invalid certificates to match.
No Validation	There are no additional settings to configure.

8. Click **OK**.

What's Next?

- ▶ Use the TLS Match to override the properties of an Application in a Service Definition as instructed in Creating Access Rules for Application Detection (page 774).

Creating Access Rules for Application Detection

Prerequisites: None

To detect application use, you must create Access rules that define the matching criteria. For detailed instructions on creating rules, see Getting Started with Editing the Rules in Policies (page 660).

The Service cell defines the protocols that are compared to the protocol-related information in each packet's header. You can use Application elements directly in the Service cell, or as part of the Service Definition.

Alternatively, you can use Application Types and Tags directly in the Service cell to match any of the Applications that belong to the Application Type or Tag group. Application Types represent general categories of Applications. Tags represent all Applications that are associated with that Tag.

Some Applications can open several related connections. If a related connection is identified by an Access rule that detects Application use, the related connection is matched against the Access rules again. If the rule that detected the Application use has Deep Inspection enabled and the related connection matches a rule that has Deep Inspection enabled, the related connection is matched against the Inspection Policy. No NAT payload modifications are done for the connection that matches the rule that detected the Application use. NAT payload modifications may be done for the related connections according to the policy.

▼ To create Access rules for application detection

1. Add an Access rule and define the **Source**, **Destination**, and **Action** according to your needs.
2. Drag and drop an Application, Application Type, or Tag element to the **Service** cell.
3. Select the Logging options according to your needs.
   - If you want to include information about Application use in the logs, select **Enforced** for Log Application Information.

What's Next?

- ▶ If you want to override properties of the Application, continue by Overriding Application Properties in Service Definitions.
- ▶ Otherwise, if you are finished editing the policy, save and install the policy to start using the new configuration.

Overriding Application Properties in Service Definitions

Applications may include protocols, port numbers, and TLS Matches. Normally, the Access rule that contains the Application only matches when the properties of the Application match:

- When Applications include a protocol and port number, the rule matches only when the protocol and port number match.
- When a TLS Match is specified in the properties of an Application, the rule matches TLS connections only if the TLS Match also matches.

You can specify additional criteria in the Service Definition to override the properties of the Application. In this case, the Access rule that contains the Application matches when the Service and/or TLS Matches specified in the rule match.

Example **The predefined Google Mail Application matches only TCP ports 80 and 443. Adding the Any TCP Service to the Service Definition allows the Application to match any traffic pattern that resembles the use of Google Mail regardless of the port.**

▼ To override application properties in Service Definitions

1. Right-click the **Service** cell and select **Edit Service**. The Service Definition dialog opens.

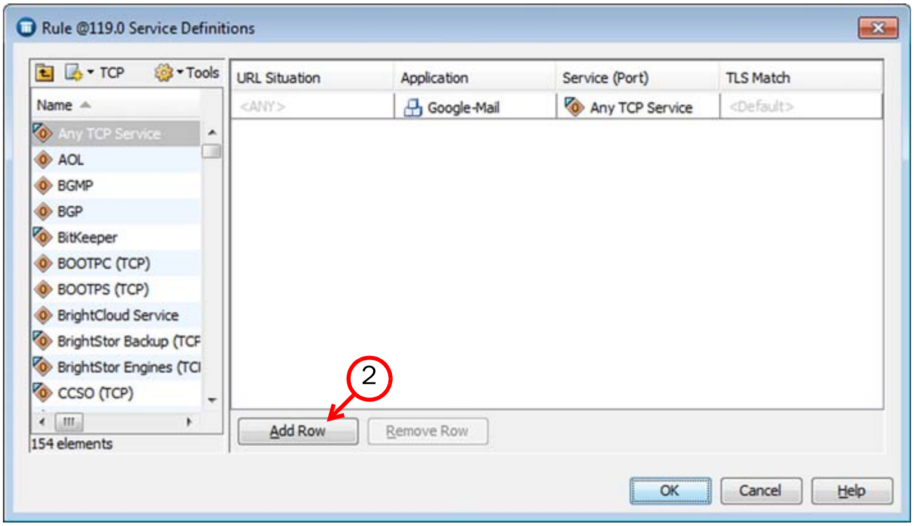

URL Situation	Application	Service (Port)	TLS Match
<ANY>	Google-Mail	Any TCP Service	<Default>

2. Click **Add Row** to add a new row to the Service Definition.
3. Drag and drop Service and/or TLS Match elements to the correct cells.
   - TLS Matches can only be used with Applications that include a TLS Match in the Application properties.
   - URL Situations cannot be used with Applications.

Note – All items on the same row must match the traffic for the row to match. You cannot use an Application element and a Service element on different rows of the Service Definition.

4. Click **OK**.

What's Next?

- If you are finished editing the policy, save and install the policy to start using the new configuration.

Logging Application Use

Prerequisites: None

You can optionally enable the logging of Application use in a Continue rule to log the use of Applications without using Application detection for access control. For detailed instructions on creating rules, see Getting Started with Editing the Rules in Policies (page 660).

To log Application use

1. Add an Access rule and define the **Source**, **Destination**, and **Service** according to your needs.
   - It is not necessary to add an Application element to the Service cell if you only want to log the use of Applications.
2. Select **Continue** as the **Action**.
3. Double-click the **Logging** cell. The Logging options dialog opens.

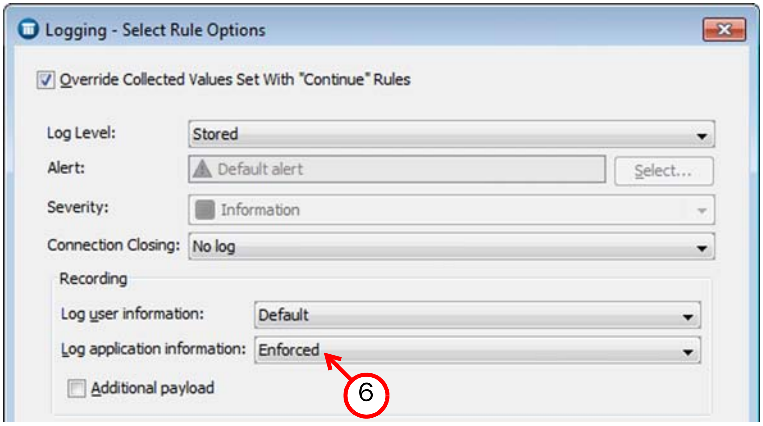

Logging - Select Rule Options

☑ Override Collected Values Set With "Continue" Rules

Log Level: Stored

Alert: Default alert

Select...

Severity: Information

Connection Closing: No log

Recording

Log user information: Default

Log application information: Enforced

□ Additional payload

4. Select **Override Collected Values Set with "Continue" Rules**.
5. Select one of the following options as the Log Level:
   - Transient
   - Stored
   - Essential
   - Alert
6. Select **Enforced** for Log Application Information.
7. Click **OK**.

What's Next?

- If you are finished editing the policy, save and install the policy to start using the new configuration.

CHAPTER 45

DEFINING USER RESPONSES

The User Response element allows you to send a customized reply to the user instead of just ending the connection when an HTTP connection is refused, terminated, or blacklisted. User Responses can be used in Access rules and in Inspection Policies for Firewalls and IPS engines.

The following sections are included:

Getting Started with User Responses

Prerequisites: None

What User Responses Do

User Responses allow you to define custom responses that are sent to the user when an HTTP or HTTPS connection is not allowed to continue. This makes it possible to explain to the user why the connection was stopped, instead of simply closing the connection with no notification. User Responses help administrators differentiate cases where a connection was blocked by the IPS from cases where a technical problem prevented the connection from going through.

Limitations

User Responses are not supported on Layer 2 Firewalls, Master Engines, or Virtual Security Engines.

Configuration Overview

1. Create a User Response element as explained in Creating User Responses.
2. Define the responses that are sent to the users when a connection is not allowed to continue. See Defining User Response Entries (page 779).
3. Use the User Response element in the Access rules and Inspection rules as required. See Editing Access Rules (page 672) and Editing Inspection Policies (page 699).

What's Next?

▶ To begin the configuration, proceed to Creating User Responses.

Creating User Responses

Prerequisites: None

▼ To create a User Response

1. Select **Configuration→Configuration→Security Engine** from the menu. The Security Engine Configuration view opens.
2. Expand the **Other Elements** branch and select **Engine Properties**.
3. Right-click **User Responses** and select **New User Response**. The User Response Properties dialog opens.
4. Enter a unique **Name**.

What's Next?

▶ Customize the responses that are sent to the users. Proceed to Defining User Response Entries (page 779).

Defining User Response Entries

Prerequisites: Creating User Responses

You can define a different User Response entry for each of the following cases in which an HTTP or HTTPS connection is not allowed to continue:

- **Connection Blacklisted**: (*HTTP only*) The connection was discarded according to a rule with the Apply Blacklist action. See Getting Started with Blacklisting (page 816) for more information.
- **Connection Discarded by Access Rule**: (*HTTP only*) The connection was dropped according to an Access rule with the Discard action.
- **Connection Terminated by Inspection Rule**: The connection was terminated according to the Inspection Policy.
- **URL Not Allowed**: The connection was terminated by one of the special URL filtering situations. See Getting Started with Web Filtering (page 790) for more information.
- **Virus Found (Firewall)**: The anti-virus feature found a virus in the web page. See Configuring Anti-Virus Settings (page 540) for more information.

Note – In some cases, such as when inspecting a large file transfer, it is not possible to apply a User Response to HTTPS traffic. In this case, the engine applies the default action for matching traffic.

▼ To define a User Response entry

1. Switch to the **HTTP(S)** tab of the User Response Properties dialog.

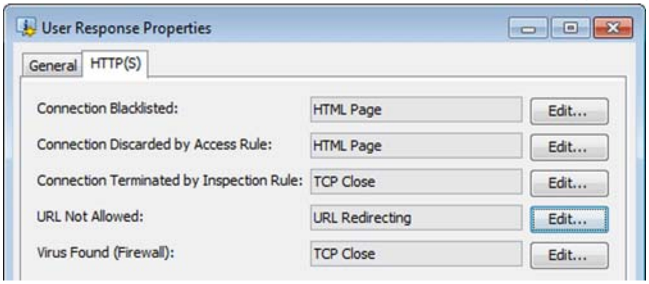

Connection Blacklisted:	HTML Page
Connection Discarded by Access Rule:	HTML Page
Connection Terminated by Inspection Rule:	TCP Close
URL Not Allowed:	URL Redirecting
Virus Found (Firewall):	TCP Close

2. Click **Edit** for the entry you want to change. The User Response Entry Properties dialog opens.

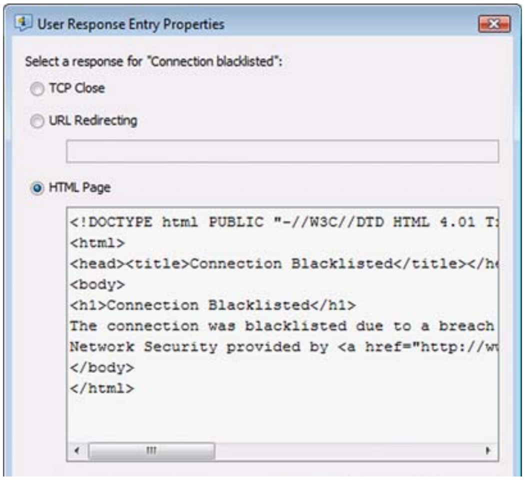

Select a response for "Connection blacklisted":

○ TCP Close

○ URL Redirecting

● HTML Page

```
<!DOCTYPE html PUBLIC "-//W3C//DTD HTML 4.01 T
<html>
<head><title>Connection Blacklisted</title></h
<body>
<h1>Connection Blacklisted</h1>
The connection was blacklisted due to a breach
Network Security provided by <a href="http://w
</body>
</html>
```

3. Select one of the following responses:
   - **TCP Close**: The connection is closed. No response is sent to the user.
   - **URL Redirecting**: Enter the URL to which the connection is redirected. The URL must begin with `http://`
   - **HTML Page**: Enter the HTML source code to display to the user. The HTML source code must be a complete HTML page, including the `<html>` and `<body>` tags.
4. Click **OK**.

What's Next?

▶ Select the User Response element in the Action options of Access rules and Inspection Exceptions as required. See Editing Access Rules (page 672) and Editing Inspection Policies (page 699).

Chapter 46

Quality of Service (QoS)

The Quality of Service (QoS) features allow you to manage bandwidth and prioritize connections on the Security Engine. QoS features are available on Firewalls, IPS engines, Layer 2 Firewalls, Master Engines, and Virtual Firewalls.

The following sections are included:

Getting Started with QoS

Prerequisites: None

QoS (Quality of Service) allows you to manage the available network bandwidth and ensure that important network services are given priority over less important traffic.

What QoS Does

The QoS features help you in the following ways:

- You can set up a *Guarantee* for a type of traffic that must always be given a certain minimum share of the available bandwidth.
- You can set up a *Limit* for maximum bandwidth for a type of traffic that must never use more than a certain share of the available bandwidth.
- You can set a *Priority* value for the traffic. Higher priority traffic is sent forward to its destination before lower priority traffic if the Security Engine needs to queue packets due to congestion.
- *Active Queue Management* (AQM) reduces the volume of dropped or retransmitted packets when there is network congestion. AQM monitors the average queue size and uses a scheduling algorithm to determine the statistical probability for dropping incoming packets.
- You can read or write DiffServ Code Point (DSCP) type of service (ToS) field markers, so that the Security Engine is aware of the priorities set by other network equipment and other equipment is aware of the priorities set in the QoS Policy.
- You can collect statistics about traffic that matches Access rules that apply a QoS Class to the traffic. QoS Class-based statistics items are used in Overviews and Reports.

Limitations

QoS features have the following notable limitations:

- QoS is available only on *Physical interfaces*, *VLAN interfaces*, *ADSL interfaces*, and *SSID interfaces*, as well as in the properties of policy-based *VPNs* on Firewalls. QoS is not available on *Modem interfaces*.
- QoS is available only on *Physical interfaces* that are part of an *Inline interface* pair on IPS engines and Layer 2 Firewalls.
- It is not possible to apply a bandwidth guarantee to incoming Internet traffic on your Internet link. By the time the Security Engine sees the traffic, the bandwidth has already been used. If you want guaranteed bandwidth for a specific portion of your incoming Internet traffic, contact your ISP and ask if they can enforce this guarantee for you.
- If you want to create QoS rules for both incoming and outgoing traffic, you must assign a QoS policy to at least two interfaces. Incoming traffic is processed according to the Firewall, IPS, or Layer 2 Firewall policy, and then the QoS policy is applied to the allowed traffic on the outgoing interface.

What Do I Need to Know Before I Begin?

- By default, the DSCP mark for the encrypted ESP packet in VPN traffic is inherited from the plaintext packet. Selecting a QoS Policy in the properties of the policy-based VPN makes it possible to mark the ESP packet after encryption. See Defining Policy-Based VPNs (page 949).
- Any priorities, limits, and guarantees are applied, and DSCP codes are written to outgoing packets on the interface that the traffic uses to exit the Security Engine according to the QoS policy and interface speed defined for that interface.

- For packets that enter the Security Engine, the QoS policy on that interface is only used for reading DSCP codes and matching them to QoS Classes for further use. This is the only QoS operation that is done on the interface that the traffic uses to enter the Security Engine.

Example **A new packet enters a Firewall through interface A and leaves the Firewall through interface B. Any priorities, guarantees, and limits are configured and applied on interface B. The priorities, guarantees, and limits configured on interface A are ignored for packets in this direction. If the packet contains a DSCP code when entering the Firewall, the DSCP code is read and matched to a QoS Class on interface A, but if a new DSCP codes is (over)written in the packet, the new code is written on interface B.**

Configuration Overview

Illustration 46.1 Elements in the Configuration

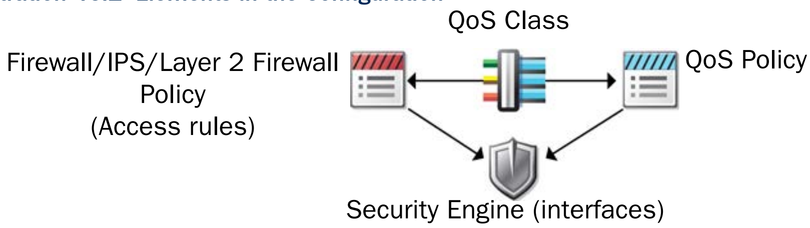

1. Create a QoS Class element for each type of traffic that you want to handle differently on any single network interface. See Creating QoS Classes (page 784).
2. Create one or more QoS Policies to define how each type of traffic is handled on the interfaces. See Defining QoS Policies (page 785).
3. Assign QoS Classes to different types of traffic in your Access rules. See Matching QoS Rules to Network Traffic (page 788).
4. Define the QoS Mode of each interface. See Getting Started with Interface Configuration (page 420).
   - The *QoS Mode* for each interface defines how QoS is applied to the interface.
   - Depending on the QoS Mode, you may also have to define the QoS Policy the interface uses and the speed of the interface.
   - You can select a QoS Mode and define a bandwidth for traffic in the properties of a Physical, VLAN, ADSL, Tunnel, or SSID Interface. Each physical, VLAN, ADSL, and SSID interface has separate QoS settings.

What's Next?

▶ To create your custom QoS scheme, proceed to Creating QoS Classes.

▶ If the default QoS Policy is sufficient for you, proceed to Matching QoS Rules to Network Traffic (page 788).

Creating QoS Classes

Prerequisites: None

QoS Classes are used to collect QoS statistics about traffic, or to create a link between the Access rules in Firewall, IPS, and Layer 2 Firewall policies and the QoS Policy. When traffic matches an Access rule, the QoS Class defined in the rule is applied to the traffic. QoS classes can also be used in Outbound Multi-Link elements to adjust the load balancing of different types of connections.

You must create one QoS Class for each rule you plan to add in any single QoS Policy, as the QoS policy cannot contain any overlapping rules. There is a default QoS Policy and three QoS Classes in the system that can be used to set priorities for high, normal, and low priority traffic without any bandwidth guarantees or limits.

▼ To create a QoS Class

1. Select **Configuration→Configuration→Security Engine** from the menu. The Security Engine Configuration view opens.
2. Expand the **Other Elements** branch of the tree.
3. Right-click **QoS Classes** and select **New QoS Class**. The QoS Class Properties dialog opens.
4. Give the QoS Class a **Name**.
5. Click **OK**. The new QoS Class is added to the resources and can now be used in policies.

Create one QoS class for each rule you plan to add in the QoS Policy.

What’s Next?

- ▶ If you only want to collect QoS statistics about traffic, proceed to Matching QoS Rules to Network Traffic (page 788).
- ▶ Otherwise, continue by Defining QoS Policies.

Defining QoS Policies

Prerequisites: Creating QoS Classes

QoS policies determine the rules that the Security Engine follows when it decides which traffic is given priority and how the available bandwidth is divided. One QoS policy can be assigned for each Physical Interface, VLAN interface, ADSL interface, SSID interface, and policy-based VPN. You can assign the same QoS policy to several interfaces.

Creating New QoS Policies

▼ To create a new QoS Policy

1. Select **Configuration→Configuration→Security Engine** from the menu. The Security Engine Configuration view opens.
2. Expand the **Policies** branch.
3. Right-click **QoS Policies** and select **New QoS Policy**. The QoS Policy Properties dialog opens.
4. Type in a **Name** and click **OK**. The new QoS policy opens.

What's Next?

▶ Continue by Editing QoS Rules (page 785) or Editing DSCP Match/Mark Rules (page 787).

Editing QoS Rules

Follow these general guidelines when editing QoS rules:

- The order of the rules in a QoS policy does not affect how traffic is handled by the Security Engine, as the match is made only based on the QoS Class.
- Enter Guarantees and Limits as percentages instead of kilobits per second values if you want to use the same QoS Policy on interfaces that have different throughputs and use the Full QoS Mode.
- Enter Guarantees and Limits as kilobits per second values if you want to use the QoS Policy on interfaces that use the DSCP Handling and Throttling QoS Mode.
- All operations are always made according to the matching rule in the QoS Policy that is assigned to the interface that the traffic uses to exit the Security Engine.
- The rules do not need to cover all traffic. When Full QoS is used, traffic that is not covered is given a priority of 8 without limits or guarantees.

▼ To edit a QoS rule

1. Right-click an existing rule (or the Not Classified rule) and select **Add Rule Before** or **Add Rule After**. A new blank rule is added.
2. Click the **QoS Class** cell in the new rule and select the QoS class.
3. Define how the Security Engine handles traffic in this QoS Class using any combination of the following options (each is optional):

Table 46.1 QoS Rule Fields

Field	Explanation
Guarantee	Sets the minimum bandwidth given to this type of traffic under any conditions. Enter a number as kilobits per second (for example, 300) or as a percentage of the available bandwidth (for example, 10%). You can also use the abbreviations M or G to enter the value in megabytes or gigabytes respectively (for example, 3M for three megabytes).
Limit	Sets the maximum bandwidth that this type of traffic is allowed to consume at any single moment. Enter a number as kilobits per second (for example, 300) or as a percentage of the available bandwidth (for example, 10%). You can also use the abbreviations M or G to enter the value in megabytes or gigabytes respectively (for example, 1G for one gigabyte).
Priority	Assigns this type of traffic a number that is used to determine the order in which the Firewall sends queued packets onwards or drops them if the queue fills up. Enter a number between 1 (highest priority) and 16 (lowest priority).
Weight	The weight of the QoS class controls the distribution of bandwidth between QoS Classes with the same priority after the Guarantees for the QoS Classes are reached. Enter the weight of the QoS Class as a value from 0 to 100. The relative weight of each QoS Class is displayed in parentheses as a percentage.
Latency	The average time packets are held in the queue for Active Queue Management (AQM). The engine makes a best effort to handle the packets within the specified time, but the Latency value is not a guarantee. Enter the value in microseconds. The default value is 100000 microseconds (0.1 seconds).

4. Save the QoS Policy.

What's Next?

- If you want to read and/or write DSCP markers for traffic, continue by Editing DSCP Match/Mark Rules (page 787).
- If you added new QoS rules, add the corresponding QoS Classes to the Access rules as explained in Matching QoS Rules to Network Traffic (page 788).
- If the necessary QoS Classes are already in the Access rules, but you have not defined how QoS is applied to all physical, VLAN, ADSL, or SSID interfaces, proceed to Getting Started with Interface Configuration (page 420).
- Otherwise, refresh the Security Engine's policy to transfer the changes.

Editing DSCP Match/Mark Rules

It is possible to read and/or write DSCP markers for a particular type of traffic without configuring Access rules to apply a QoS class to the traffic. The matching is done based only on the QoS Policy. When a packet that matches a particular protocol comes in, the DSCP markers are read and assigned a QoS Class according to the DSCP Match/Mark rules of the QoS policy. When the packet is sent out, a DSCP mark is written in packets based on the QoS Class according to the DSCP Match/Mark rules of the QoS Policy on the interface through which the traffic leaves the firewall. This allows you to:

- Communicate the priority of this traffic to other devices that support QoS.
- Convert the packet to use a different classification scheme if the QoS Class was originally assigned to matching traffic by a DSCP match in the source interface's QoS Policy.
- Clear the DSCP classification set by other devices by entering 0 as the value (shown in the policy as 0x00).

▼ To edit a DSCP Match/Mark rule

1. Switch to the **DSCP Match/Mark** tab.
2. Right-click an existing rule (or the Not Classified rule) and select **Add Rule Before** or **Add Rule After**. A new blank rule is added.
3. Click the **QoS Class** cell in the new rule and select the QoS class.
4. Define how the engine reads and/or writes DSCP code for traffic in this QoS Class:

Field	Explanation
DSCP Match	Assigns the rule's QoS Class to traffic when the DSCP code (ToS field) defined in this cell is seen in traffic. This field is the only field that is applied on the interface that the packets use to enter the firewall.
DSCP Mark	Defines the DSCP code (ToS field) that is written to packets that match this QoS rule when the packets exit the firewall. The DSCP Mark allows you to communicate the priority of this traffic to other devices that support QoS. You can also use the field to clear the DSCP classification set by other devices by entering **0** as the value (shown in the policy as 0x00).

5. Save the QoS Policy.

What's Next?

▶ If you also added new rules on the QoS tab, add the corresponding QoS Classes to the Access rules as explained in Matching QoS Rules to Network Traffic (page 788).

▶ If you have not defined how QoS is applied to all physical, VLAN, ADSL, or SSID interfaces, proceed to Getting Started with Interface Configuration (page 420).

▶ Otherwise, refresh the Security Engine's policy to transfer the changes.

Matching QoS Rules to Network Traffic

Prerequisites: Creating QoS Classes, Defining QoS Policies

The rules on the QoS tab of the QoS Policy are linked to different types of traffic using the QoS classes. QoS Classes are matched to traffic in the Access rules with following actions:

- Access rules with the **Allow** action set a QoS class for traffic matching that rule.
- Access rules with the **Continue** action set a QoS class for all subsequent matching rules that have no specific QoS class defined.
- Access rules with the **Use VPN** action (*Firewall only*) set a QoS class for VPN traffic. Incoming VPN traffic may also match a normal Allow rule after decryption. Otherwise, for outgoing traffic, encryption is done after the QoS Policy is checked. For incoming traffic, decryption is done before the QoS Policy is checked.

However, if you only want to read and use DSCP markers set by other devices, the QoS Class is assigned according to the rules on the DSCP Match/Mark tab of the QoS Policy.

Note – If traffic is assigned a QoS Class using a DSCP Match rule, the same traffic must not match any Access rule that assigns a different QoS Class to the same traffic. Such Access rules override the QoS Class set using a DSCP Match.

▼ To match a QoS rule to network traffic

1. Open the Firewall, IPS, or Layer 2 Firewall Policy for editing.
2. Click the **QoS Class** cell of a rule that allows traffic or a Continue rule and drag and drop one QoS Class element into the cell.
   - The QoS Class links connections to a rule on the QoS tab of the QoS Policy. There can be different rules in different QoS Policies for the same QoS Class.
   - Packets in both directions of a connection are assigned the same QoS Class (when the connection is handled statefully, that is, connection tracking is active for the rule).
   - The applied QoS Class is shown in the logs, and reports can also be generated based on this information.

What's Next?

- ▶ If you have just created a new QoS policy, define the interface speeds and select this QoS policy for the relevant interfaces. See Getting Started with Interface Configuration (page 420).
- ▶ Otherwise, refresh the Security Engine's policy to transfer the changes.

Related Tasks

- ▶ Editing Access Rules (page 672)

Chapter 47

Filtering Web Addresses

Web filtering allows you to filter web access based on categories of content and/or lists of individual websites.

The following sections are included:

Getting Started with Web Filtering

Prerequisites: None

What Web Filtering Does

Web filtering prevents users from intentionally or accidentally accessing most websites that provide content that is objectionable (for example, pornographic or violent imagery), potentially harmful (for example, various scam sites), or leisure-oriented (for example, entertainment news). This type of content filtering can increase network security and enforce an organization's policy on acceptable use of resources.

Web filtering compares URLs seen in clients' requests to lists of URLs that you want to match. The lists of URLs can be defined as manual blacklists and/or as dynamically-populated content categories provided by a categorization service. Currently, you can use categorizations provided by the BrightCloud service for category-based filtering. When category-based filtering is active, you can additionally use the manual URL lists as whitelists of individual addresses that must not be blocked by category-based filtering. You can check the category BrightCloud has assigned to a URL at http://www.brightcloud.com/support/lookup.php.

Limitations

- Category-based web filtering (BrightCloud) is a separately licensed feature.
- Security Engines in the Firewall/VPN role must be licensed for deep packet inspection to use web filtering.
- Category-based web filtering is not supported on Virtual Security Engines.
- Web filtering and CIS redirection cannot be done for the same traffic. See Getting Started with External Content Inspection (page 810) for more information about CIS redirection.
- Category-based web filtering is based on the classifications of the external service, so it is not possible to manually add or directly edit the web filtering categories. Use the manual blacklisting/whitelisting feature and the Inspection Policy to make your adjustments.

Configuration Overview

Illustration 47.1 Elements in the Configuration

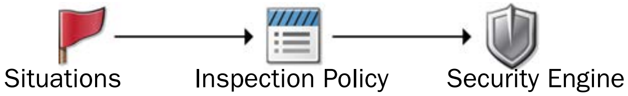

1. (*For category-based filtering*) Make sure that the engine and the network are set up so that the engine can fetch the lists of URLs directly from the BrightCloud servers.
   - Make sure that DNS server(s) are defined on the **Advanced** tab in the engine's properties. See Getting Started with Advanced Engine Settings (page 552).
   - Make sure that the engines have access to the DNS server (UDP port 53, Service element "DNS (UDP)"), and to the BrightCloud servers (TCP port 2316, Service element "BrightCloud update"). The Firewall Template policy allows these connections from the firewall engine on which the policy is installed, but not from other components.
2. (*Optional*) Create User Responses to notify users about matches that are found. See Getting Started with User Responses (page 778).
3. (*Optional when using category-based filtering*) Define lists of individual URLs that you want to filter. See Blacklisting/Whitelisting Web URLs Manually (page 791).
4. Add rules to the Inspection Policy. See Creating Web Filtering Rules (page 792).

Blacklisting/Whitelisting Web URLs Manually

Prerequisites: To use on a firewall, check that the firewall is licensed for deep inspection

HTTP URL Filtering Situations allow you to define lists of URLs that blacklist (block) access to the specified URLs. When used in combination with category-based web filtering, these types of lists can additionally be used to whitelist (allow) individual URLs that are included in an otherwise blocked category. URL whitelisting only affects URL-based filtering. It does not exclude the traffic from other inspection checks.

Note – To be able to filter HTTPS URLs, you must decrypt the traffic. See Getting Started with TLS inspection (page 794).

▼ To configure website access control

1. Create a new Situation element with the appropriate basic properties (see Creating New Situation Elements (page 757) for more information). Leave the **Situation Type** selection empty if you are creating a whitelist.
2. Switch to the **Context** tab and select the Context: **Application Protocols→HTTP→HTTP URL Filter**.
3. Click the **Add** button below the entry panel. You can add up to 20 URLs per Situation.
4. Double-click the newly added row and enter the URL without specifying the protocol, for example, `www.example.com`.
   - Each URL can be a maximum of 64 characters long.
   - To block or allow access to a domain, type: `example.com`
   - To block or allow access to an exact URL, type: `www.example.com/index.html`. Users can still access other pages such as `www.example.com/main.html`.
   - An asterisk (*) can be used *at the beginning of the URL*. For example: `*.example.com` (matches a request for `www.example.com` but not the shorter variant `example.com`).
5. (*Optional*) Add Tags to organize this Situation and click **OK**.

What's Next?

▶ Creating Web Filtering Rules (page 792)

Creating Web Filtering Rules

Prerequisites: Category-based web filtering is a separately licensed feature

Category-based URL filtering and URL filtering based on URL lists can be configured in the Inspection Policy. Category-based filtering Situations are listed in the Rules tree under **Web Filtering**. The categories are updated through dynamic update packages. Firewalls do not deep inspect traffic by default, so make sure there is an Access rule that matches HTTP traffic and has the **Deep Inspection** option selected in the Action options.

Add your whitelist to the Inspection Policy in an Exception rule that is set to use the **Permit** action. Your manual URL lists can be found under **Situations→Custom Situations**. The whitelist must be in the Exceptions to be able to override category-based filtering settings in the Rules tree (which is the purpose of whitelisting).

Blacklists can be included as an Exception or added to the Rules tree (by selecting a Situation Type for the Situation).

You may also want to include a User Response in the rule to display warnings or notes in the users' browsers when web filtering prevents access. See Editing Inspection Policies (page 699) and Getting Started with User Responses (page 778) for more information.

What's Next?

▶ Refresh the policy of the Security Engine to activate the configuration.

CHAPTER 48

SETTING UP TLS INSPECTION

TLS Inspection allows you to decrypt TLS connections so that they can be inspected for malicious traffic.

The following sections are included:

Getting Started with TLS inspection

Prerequisites: None

What TLS inspection Does

The TLS protocol allows applications to communicate across a network in a way designed to ensure the confidentiality and integrity of the communications. For example, HTTPS uses the TLS protocol to encapsulate HTTP connections. However, the encrypted connection can also be used to hide malicious traffic. The TLS inspection feature decrypts TLS traffic so that it can be inspected in the same way as unencrypted traffic, and then re-encrypts the traffic before sending it to its destination.

The TLS inspection feature helps you in the following ways:

- Server protection allows you to inspect incoming TLS connections to servers in the protected network.
- Client protection allows you to inspect outgoing TLS connections initiated by users in the protected network.

You can use client protection alone, server protection alone, or client and server protection together.

Limitations

TLS inspection has the following limitations:

- IPS engines must be deployed in inline mode to use TLS inspection. TLS Inspection cannot be done for traffic picked up through Capture interfaces.
- TLS inspection is not supported on Single IPS engines or on Single Layer 2 Firewalls if the engines are deployed alongside a Firewall Cluster that using dispatch clustering.
- TLS inspection and CIS redirection cannot be done for the same traffic. See Getting Started with External Content Inspection (page 810) for more information about CIS redirection.
- Default Trusted Certificate Authority elements are automatically added to the system from dynamic update packages and cannot be edited or deleted.
- TLS Inspection is not supported on Master Engines.

What Do I Need to Know Before I Begin?

- Traffic that uses TLS may be protected by laws related to the privacy of communications. Decrypting and inspecting this traffic may be illegal in some jurisdictions.
- Because the TLS communications mediated by the engine are decrypted for inspection, and because the private keys of the servers are stored unencrypted in the Server Protection Credentials elements on the Management Server and on the engine, you must carefully consider security precautions when using TLS inspection. The following recommendations are general guidelines for ensuring the security of the engine and the Management Center:
  - Run the Management Server on a hardened operating system.
  - Disable SSH access to the engine's command line.
  - Ensure that the engine's Control interface is in a controlled network.
  - Save Management Server backups as encrypted files.

Configuration Overview

Illustration 48.1 Elements in the Configuration

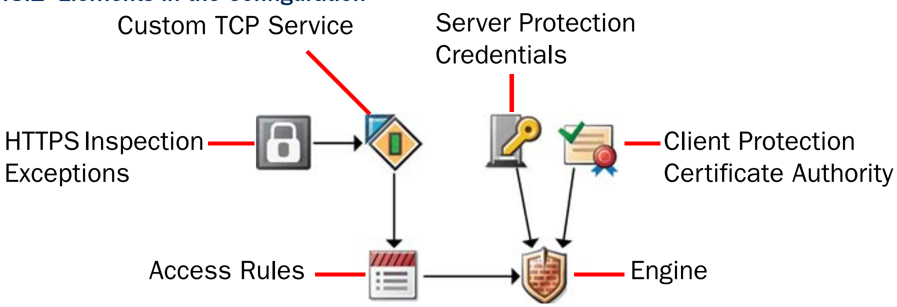
Custom TCP Service
Server Protection Credentials
HTTPS Inspection Exceptions
Client Protection Certificate Authority
Access Rules
Engine

1. If you want to configure server protection, create Server Protection Credentials elements. See Configuring Server Protection (page 796).
2. If you want to configure client protection, create Client Protection Certificate Authority elements. See Configuring Client Protection (page 797).
3. (*Optional*) Define custom Trusted Certificate Authority elements in addition to the default system elements. See Defining Trusted Certificate Authorities for TLS Inspection (page 800).
4. Activate client protection and/or server protection in the properties of the Firewall, IPS engine, or Layer 2 Firewall. See Activating TLS inspection on the Engine (page 803).
5. (*Optional*) Define an HTTPS Inspection Exceptions element to exclude certain domains from decryption and inspection. See Excluding Connections from TLS Inspection (page 804).
6. Create a custom HTTPS Service that enables TLS inspection. See Defining a Custom HTTPS Service (page 806).
7. Create Access rules to select traffic for TLS inspection. See Creating Access Rules for TLS inspection (page 807).

What's Next?

- If you want to configure server protection, proceed to Configuring Server Protection (page 796).
- Otherwise, proceed to Configuring Client Protection (page 797).

Configuring Server Protection

Prerequisites: None

The Server Protection Credentials element stores the private key and certificate of an internal server. The private key and certificate allow the Firewall, IPS engine, Layer 2 Firewall, or Virtual Firewall to decrypt TLS traffic for which the internal server is the destination so that it can be inspected. The certificate and the associated private key must be compatible with OpenSSL and be in PEM format. Make sure that the server's private key and certificate are accessible from the computer where you use the Management Client.

▼ To define a Server Protection Credentials element

1. Select **Configuration→Configuration→Security Engine** from the menu. The Security Engine Configuration view opens.
2. Browse to **Other Elements→Engine Properties→Certificates**.
3. Right-click **Server Protection Credentials** and select **New Server Protection Credentials**. The Server Protection Credentials Properties dialog opens.

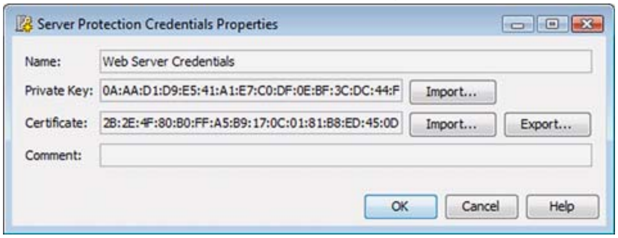

Server Protection Credentials Properties

Name: Web Server Credentials

Private Key: 0A:AA:D1:D9:E5:41:A1:E7:C0:DF:0E:BF:3C:DC:44:F

Certificate: 2B:2E:4F:80:B0:FF:A5:B9:17:0C:01:81:B8:ED:45:0D

Comment:

4. Enter a unique **Name**.
5. Click the **Import** button for the **Private Key** field and browse to the server's private key.
   - If the private key is encrypted, you are prompted to enter the password.
6. Click the **Import** button for the **Certificate** field and browse to the server's certificate.
7. Click **OK**.

What's Next?

- ▶ If you want to inspect connections between internal clients and external HTTPS servers, proceed to Configuring Client Protection (page 797).
- ▶ Otherwise, proceed to Activating TLS inspection on the Engine (page 803).

Configuring Client Protection

Prerequisites: None

When an internal client makes a connection to an external server that uses TLS, the engine generates a substitute certificate that allows it to establish a secure connection with the internal client. The engine adds a Netscape Certificate Comment to the Extensions in the substitute certificate to indicate that the certificate is a dynamically created certificate for Stonesoft SSL/TLS deep inspection. The Client Protection Certificate Authority element contains the credentials the engine uses to sign the substitute certificate it generates. If the engine does not use a signing certificate that is already trusted by users' web browsers when it signs the substitute certificates it generates, users receive warnings about invalid certificates. To avoid these warnings, you must either import a signing certificate that is already trusted, or configure users' web browsers to trust the engine's signing certificate.

Note – Traffic that uses TLS may be protected by laws related to the privacy of communications. Decrypting and inspecting this traffic may be illegal in some jurisdictions.

Creating Client Protection Certificate Authority Elements

▼ To create a Client Protection Certificate Authority element

1. Select **Configuration→Configuration→Security Engine**. The Security Engine Configuration view opens.
2. Browse to **Other Elements→Engine Properties→Certificates**.
3. Right-click **Client Protection Certificate Authorities** and select **New Client Protection Certificate Authority**. The Client Protection Certificate Authority Properties dialog opens.

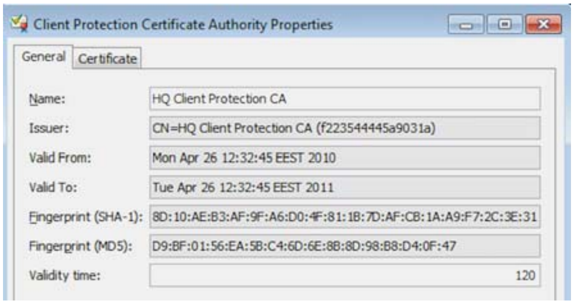

Name: HQ Client Protection CA
Issuer: CN=HQ Client Protection CA (f223544445a9031a)
Valid From: Mon Apr 26 12:32:45 EEST 2010
Valid To: Tue Apr 26 12:32:45 EEST 2011
Fingerprint (SHA-1): 8D:10:AE:B3:AF:9F:A6:D0:4F:81:1B:7D:AF:CB:1A:A9:F7:2C:3E:31
Fingerprint (MD5): D9:BF:01:56:EA:5B:C4:6D:6E:8B:8D:98:B8:D4:0F:47
Validity time: 120

4. Enter a unique **Name**.
5. (*Optional*) Enter the **Validity time** (in minutes) for the substitute certificates the engine creates.
   - Each substitute certificate expires at the end of the validity time, and the engine automatically generates a new certificate. This may produce warnings or error messages in the users' web browsers. To avoid excessive warnings, define a sufficiently long validity time, for example several hours.

Note – All fields except the Name and Validity time on the General tab are grayed out. The grayed out fields are always filled in automatically based on information contained in the certificate you generate or import, and you cannot change them.

What's Next?

- If you want to import an existing certificate, proceed to Importing a Private Key and Signing Certificate for HTTPS Client Protection.
- If you want to generate a new certificate for the Client Protection Certificate Authority, proceed to Generating a Private Key and Signing Certificate for HTTPS Client Protection.

Importing a Private Key and Signing Certificate for HTTPS Client Protection

If you already have a certificate authority that is trusted by users' web browsers, you can import its private key and signing certificate for the engine to use when it signs the substitute certificates it creates. This removes the need to separately configure users' web browsers to trust the engine's signing certificate. You can also import a private key and signing certificate that you generated outside of the Stonesoft Management Center even if you do not already have a certificate authority that is trusted by users' web browsers. The certificate and the associated private key must be compatible with OpenSSL and be in PEM format. Make sure that the private key and certificate are accessible from the computer where you use the Management Client.

To import a private key and signing certificate for HTTPS client protection

1. Switch to the **Certificate** tab.

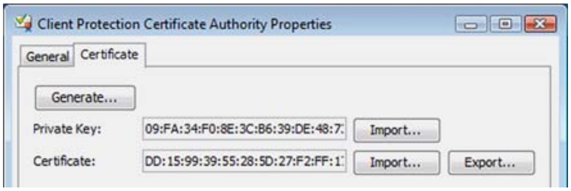

Private Key: 09:FA:34:F0:8E:3C:B6:39:DE:48:7:

Certificate: DD:15:99:39:55:28:5D:27:F2:FF:1:

2. Click the **Import** button for the **Private Key** field and browse to the private key.
3. Click the **Import** button for the **Certificate** field and browse to the certificate.

If users' web browsers are not already configured to trust the certificate authority whose signing certificate you imported here, you must add it to the list of certificate authorities that are trusted by users' web browsers when you are finished configuring TLS inspection in the Management Center.

What's Next?

- Proceed to Activating TLS inspection on the Engine (page 803).

Generating a Private Key and Signing Certificate for HTTPS Client Protection

If you do not already have a private key and signing certificate for the engine, you can generate a private key and signing certificate in the Client Protection Certificate Authority properties. When you generate a private key and signing certificate in the Management Center, you must export the certificate and add it to the list of certificate authorities that are trusted by users' web browsers.

▼ To generate a private key and signing certificate for HTTPS client protection

1. Switch to the **Certificate** tab.

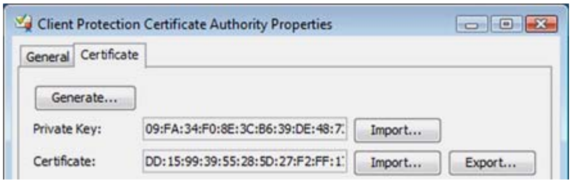

Client Protection Certificate Authority Properties

Private Key: 09:FA:34:F0:8E:3C:B6:39:DE:48:7:

Certificate: DD:15:99:39:55:28:5D:27:F2:FF:1:

2. Click **Generate**. The Signing Certificate Details dialog opens.

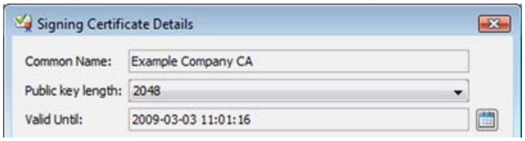

Signing Certificate Details

Common Name: Example Company CA

Public key length: 2048

Valid Until: 2009-03-03 11:01:16

3. Enter the Certificate Authority's **Common Name**.
4. (*Optional*) Select the **Public Key Length**.
5. Specify the date and time the signing certificate is **Valid Until**.
6. Click **OK**. The Signing Certificate Details dialog closes, and a private key and signing certificate are generated.

What's Next?

▶ Export the engine's signing certificate so that it can be added to the list of certificate authorities that are trusted by users' web browsers. Proceed to Exporting an HTTPS Client Protection Certificate.

Exporting an HTTPS Client Protection Certificate

If users' web browsers are not configured to trust the engine's signing certificate, users receive warnings about invalid certificates. If you generated the signing certificate for client protection in the Management Center, you must export the certificate and add it to the list of certificate authorities that are trusted by users' web browsers.

These instructions assume that you already have the Client Protection Certificate Authority Properties dialog open.

▼ To export an HTTPS client protection certificate

1. Switch to the **Certificate** tab.

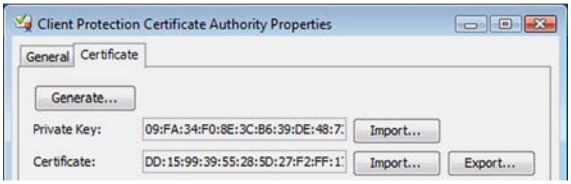

Private Key: 09:FA:34:F0:8E:3C:B6:39:DE:48:7

Certificate: DD:15:99:39:55:28:5D:27:F2:FF:1

2. Click the **Export** for the **Certificate** field and browse to the location where you want to save the file.

When you are finished configuring TLS inspection in the Management Center, add the exported certificate to the list of certificate authorities that are trusted by users' web browsers.

What's Next?

- ▶ Proceed to Activating TLS inspection on the Engine (page 803).

Defining Trusted Certificate Authorities for TLS Inspection

Prerequisites: Configuring Client Protection

Trusted Certificate Authority elements represent the certificates that identify certificate authorities. There are default Trusted Certificate Authority elements for many major certificate authorities. When a client in the protected network connects to an HTTPS server, the engine checks whether the certificate authority that signed the server's certificate is one of the Trusted Certificate Authorities. If the certificate was signed by one of the Trusted Certificate Authorities, the engine makes a substitute certificate that matches the server's certificate and signs it with the Client Protection Certificate Authority signing certificate. If the server's certificate is not signed by a Trusted Certificate Authority, the engine makes a new self-signed certificate, and users receive a warning that the issuer of the certificate is not trusted. In both cases, client protection continues to function normally.

If you are using client protection and users need to connect to domains whose certificates are not signed by one of the default Trusted Certificate Authorities, you can define a Trusted Certificate Authority element to represent it. When you define a CA as trusted, all certificates signed by that CA are considered valid until their expiration date.

Creating Trusted Certificate Authority Elements

▼ To create a Trusted Certificate Authority element

1. Select **Configuration→Configuration→Security Engine**. The Security Engine Configuration view opens.
2. Browse to **Other Elements→Engine Properties→Certificates**.
3. Right-click **Trusted Certificate Authorities** and select **New Trusted Certificate Authority**. The Trusted Certificate Authority Properties dialog opens.

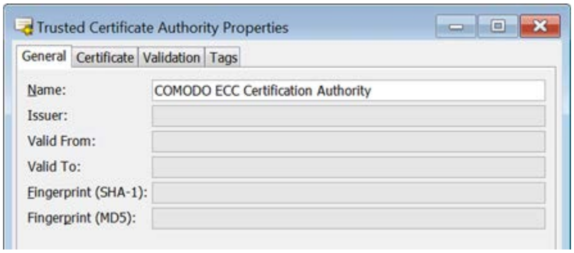

Name: COMODO ECC Certification Authority
Issuer:
Valid From:
Valid To:
Fingerprint (SHA-1):
Fingerprint (MD5):

4. Enter a unique **Name**.

Note – All fields except the Name on the General tab are grayed out. The grayed out fields are always filled in automatically based on information contained in the certificate you import and you cannot change them.

Importing a Trusted Certificate Authority Certificate for TLS inspection

Make sure that the certificate authority's certificate is accessible from the computer where you use the Management Client.

▼ To import a Trusted Certificate Authority certificate

1. Switch to the **Certificate** tab.

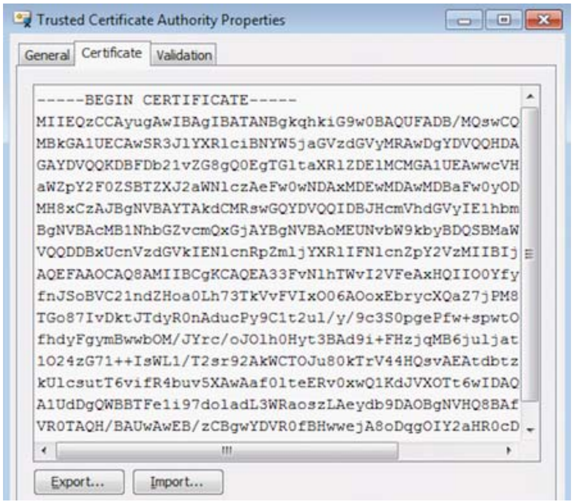

```
-----BEGIN CERTIFICATE-----
MIIEQzCCAyugAwIBAgIBATANBgkqhkiG9w0BAQUFADB/MQswCQ
MBkGA1UECAwSR3J1YXR1ciBNYW5jaGVzdGVyMRAwDgYDVQQHDA
GAYDVQQKDBFDb21vZG8gQ0EgTG1taXR1ZDE1MCMGA1UEAwwcVH
aWZpY2F0ZSBTZXJ2aWN1czAeFw0wNDAxMDEwMDAwMDBaFw0yOD
MH8xCzAJBgNVBAYTAkdCMRswGQYDVQQIDBJHcmVhdGVyIE1hbm
BgNVBAcMB1NhbGZvcmQxGjAYBgNVBAoMEUNvbW9kbyBDQSBMaW
VQQDDBxUcnVzdGVkIEN1cnRpZm1jYXR1IFN1cnZpY2VzMIIBIj
AQEFAAOCAQ8AMIIBCgKCAQEA33FvN1hTWvI2VFeAxHQII00Yfy
fnJSoBVC21ndZHoa0Lh73TkVvFVIxO06AOoxEbrycXQaZ7jPM8
TGo87IvDktJTdyR0nAducPy9C1t2ul/y/9c3S0pgePfw+spwtO
fhdyFgymBwwbOM/JYrc/oJO1h0Hyt3BAd9i+FHzjqMB6juljat
1O24zG71++IsWL1/T2sr92AkWCTOJu80kTrV44HQsvAEAtdbtz
kUlcsutT6vifR4buv5XAwAaf0lteERv0xwQ1KdJVXOTt6wIDAQ
A1UdDgQWBBTFe1i97doladL3WRaoszLAeydb9DAOBgNVHQ8BAf
VR0TAQH/BAUwAwEB/zCBgwYDVR0fBHwwejA8oDqgOIY2aHR0cD
```

2. Click **Import**. A file browser dialog opens.

3. Browse to the certificate and click **Open**.

What's Next?

- If you want to configure certificate revocation list (CRL) checks for this certificate authority, proceed to Configuring Certificate Revocation List Checks for TLS Inspection.
- Otherwise, click **OK** to close the Trusted Certificate Authority Properties dialog and proceed to Activating TLS inspection on the Engine (page 803).

Configuring Certificate Revocation List Checks for TLS Inspection

You can optionally set up the Management Center to query certificate revocation lists (CRLs) on which the certificate authority lists certificates that are cancelled, for example, because they have been compromised. By default, the system Refresh Certificate Revocation Lists Task runs daily. The Management Server connects to the specified CRL servers, checks the CRL information, and transfers the updated CRLs to the engines. The Management Server must have an Internet connection to connect to external CRL servers.

To configure certificate revocation list checks

1. Switch to the **Validation** tab.

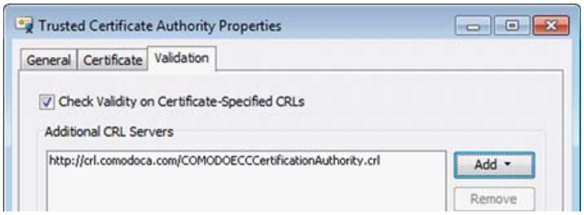

☑ Check Validity on Certificate-Specified CRLs

Additional CRL Servers

http://crl.comodoca.com/COMODOECCCertificationAuthority.crl

2. Select **Check Validity on Certificate-Specified CRLs**.
3. Click **Add** and select **Manual Server Address**. The Add CRL Server dialog opens.
4. Enter the URL of the CRL server and click **OK**.

Example `http://example.com`

5. Add any other CRL servers in the same way, then click **OK** to close the Trusted Certificate Authority Properties dialog.

What's Next?

- Proceed to Activating TLS inspection on the Engine.

Activating TLS inspection on the Engine

Prerequisites: Configuring Server Protection, Configuring Client Protection

Depending on the elements you select in the engine properties, you can activate the capability for client protection alone, server protection alone, or client and server protection together.

▼ To specify TLS inspection options in engine properties

1. Select **Configuration→Configuration→Security Engine**. The Security Engine Configuration view opens.
2. Click **Security Engines**. A list of engines opens on the right.
3. Right-click the engine and select **Properties**. The Engine Properties dialog opens.
4. Switch to the **Add-Ons** tab.

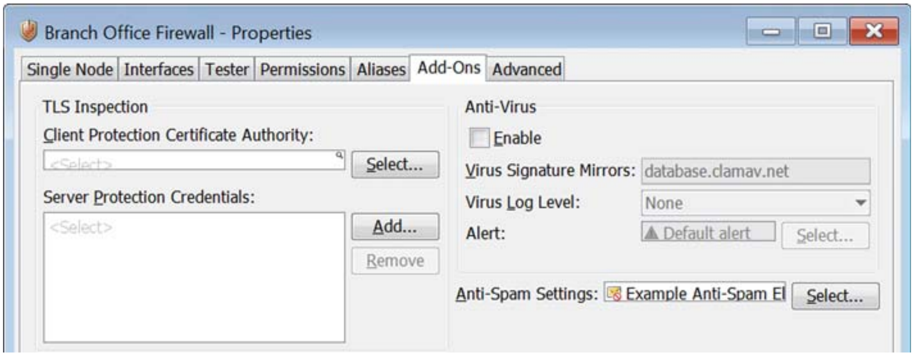

Branch Office Firewall - Properties

TLS Inspection

Client Protection Certificate Authority: <Select>

Server Protection Credentials: <Select>

Anti-Virus

□ Enable

Virus Signature Mirrors: database.clamav.net

Virus Log Level: None

Alert: Default alert

Anti-Spam Settings: Example Anti-Spam El

5. (*For client protection*) Click **Select** and select the **Client Protection Certificate Authority** element.
6. (*For server protection*) Click **Add** and select the **Server Protection Credentials** element(s). The selected elements are added to the list.
7. Click **OK** to close the engine properties dialog.

What's Next?

- ▶ If you want to exclude certain domains from HTTPS decryption and inspection, proceed to Excluding Connections from TLS Inspection.
- ▶ Otherwise, proceed to Defining a Custom HTTPS Service (page 806).

Excluding Connections from TLS Inspection

Prerequisites: None

Traffic to and from some servers that use TLS may contain users' personal information that is protected by laws related to the privacy of communications. Decrypting and inspecting this traffic may be illegal in some jurisdictions. You can optionally exclude connections from decryption and inspection in two ways: globally with a TLS Match element, or for specific matching traffic with an HTTPS Inspection Exception element. In both cases, connections to the specified domains are allowed to pass through the Firewall, IPS engine, or Layer 2 Firewall without being decrypted.

Globally Excluding Domains From Decryption

TLS Match elements allow you to prevent the specified traffic from being decrypted. TLS Matches that deny decrypting are applied globally, so even if the TLS Match elements are not used in the policy, matching connections are never decrypted. In most cases, TLS Matches are the recommended way to prevent traffic from being decrypted and inspected.

Note – Globally excluding domains from decryption also prevents Applications from being detected in encrypted connections to the specified domain. If you only want to exclude a domain from TLS inspection, see Excluding Domains from Inspection of HTTPS Traffic (page 805)

▼ To globally exclude domains from decryption

1. Select **Configuration→Configuration→Security Engine**. The Security Engine Configuration view opens.
2. Expand the **Other Elements** branch.
3. Right-click **TLS Matches** and select **New TLS Match**. The TLS Match Properties dialog opens.

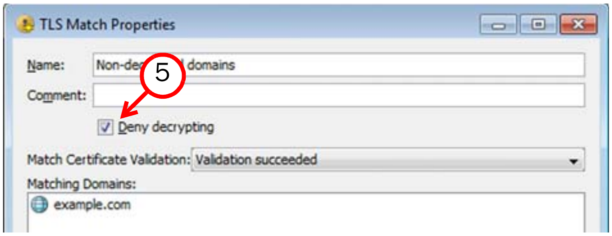
TLS Match Properties

Name: Non-de domains

Comment:

☑ Deny decrypting

Match Certificate Validation: Validation succeeded

Matching Domains:
example.com

4. Enter a unique **Name**.
5. Select **Deny Decrypting**.
6. Select **Validation succeeded** for Match Certificate Validation.
7. Click **Add** and specify the **Matching Domains** to exclude from decryption. If no domains are specified, any connection for which validation succeeded is excluded from decryption.
8. Click **OK**. Connections are excluded from decryption as specified in the TLS Matches.

Related Tasks

- ▶ Getting Started With Applications (page 772)

Excluding Domains from Inspection of HTTPS Traffic

The HTTPS Inspection Exceptions element is a list of domains that are excluded from decryption and inspection. HTTPS Inspection Exceptions are used in a custom HTTPS service to define a list of domains for which HTTPS traffic is not decrypted. The custom HTTPS service must be used in a rule, and only traffic that matches the rule is excluded from decryption and inspection. HTTPS Inspection Exceptions are primarily intended for backwards compatibility.

▼ To exclude domains from inspection of HTTPS Traffic

1. Select **Configuration→Configuration→Security Engine**. The Security Engine Configuration view opens.
2. Expand the **Other Elements** branch.
3. Right-click **HTTPS Inspection Exceptions** and select **New HTTPS Inspection Exceptions**. The HTTPS Inspection Exceptions Properties dialog opens.

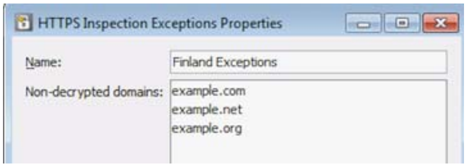

Name: Finland Exceptions

Non-decrypted domains: example.com
example.net
example.org

4. Enter a unique **Name**.
5. Click **Add**. The Add Non-decrypted Domain dialog opens.
6. Enter the domain name or NetBIOS name and click **OK**. The domain is added to the list of non-decrypted domains.

Caution – The domain name or NetBIOS name must exactly match the DNSName or IPAddress fields in the Subject Alternative Name or Common Name in the certificate of the server you want to exclude from decryption. Otherwise, the domain is not excluded from decryption and inspection.

7. Repeat these steps for any additional domains, then click **OK**.

What's Next?

▶ Create a custom HTTPS Service that enables TLS inspection and select the HTTPS Inspection Exceptions in the Protocol Parameters. Proceed to Defining a Custom HTTPS Service.

Defining a Custom HTTPS Service

Prerequisites: (*Optional*) Excluding Connections from TLS Inspection

By default, TLS inspection is not enabled in the HTTPS Service. To enable TLS inspection, you must create a custom HTTPS Service. You can optionally also specify the HTTPS Inspection Exceptions to exclude domains from decryption and inspection.

▼ To define a custom HTTPS Service

1. Select **Configuration→Configuration→Security Engine**. The Security Engine Configuration view opens.
2. Browse to **Other Elements→Services→TCP**. A list of TCP Services opens on the right.
3. Right-click the default HTTPS Service and select **New→Duplicate**. The TCP Service Properties dialog opens with the properties of the HTTPS Service.
4. Enter a unique **Name** for the custom Service.
5. Switch to the **Protocol Parameters** tab.

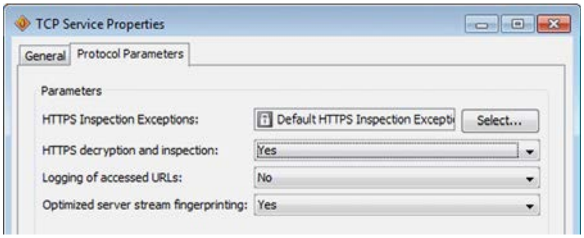

HTTPS Inspection Exceptions: Default HTTPS Inspection Excepti
HTTPS decryption and inspection: Yes
Logging of accessed URLs: No
Optimized server stream fingerprinting: Yes

6. (*Optional*) Click the **Select** button for the **HTTPS Inspection Exceptions** field and select the HTTPS Inspection Exceptions element in the dialog that opens.
7. Select **Yes** for **HTTPS Decryption and Inspection**.
8. Click **OK**.

What's Next?

▶ Use the custom HTTPS Service in an Access rule to select HTTPS traffic for inspection. Proceed to Creating Access Rules for TLS inspection (page 807).

Creating Access Rules for TLS inspection

Prerequisites: Activating TLS inspection on the Engine, Defining a Custom HTTPS Service

To select HTTPS traffic for inspection, you must create Access rules that use the custom HTTPS Service that you have created. You must select Deep Inspection in the Action options of the Firewall Access rules to enable inspection of matching traffic. Deep Inspection is enabled by default in the IPS and Layer 2 Firewall Access rules. Traffic that matches the Access rules for TLS inspection is decrypted and matched against HTTP Situations in the Inspection rules in the same way as unencrypted HTTP traffic. Any traffic that is allowed to continue by an Inspection rule is re-encrypted and sent to its destination.

▼ To create Access rules for TLS inspection

1. (*Client Protection*) Add a rule with the following properties to select traffic from clients in the internal network for inspection:

Table 48.1 Access Rules for Client Protection

Source	Destination	Service	Action
The elements that represent clients in your internal network or ANY.	The elements that represent the HTTPS server(s) to which internal clients connect, or ANY.	Your custom HTTPS Service	Allow **Deep Inspection** selected in the Action options.

2. (*Server Protection*) Add a rule with the following properties to select traffic to internal servers for inspection:

Table 48.2 Access Rules for Server Protection

Source	Destination	Service	Action
The elements that represent the clients that connect to your HTTPS server, or ANY.	The elements that represent your internal HTTPS server(s).	Your custom HTTPS Service	Allow **Deep Inspection** selected in the Action options.

What's Next?

- Define Inspection rules to match the decrypted traffic against HTTP Situations. Proceed to Editing Inspection Policies (page 699).
- Otherwise, if you are finished editing the policy, save and install the policy to start using the new configuration.

Related Tasks

- Editing Access Rules (page 672)

Chapter 49

External Content Inspection

Content inspection, such as virus scanning or URL filtering, can be done by integrating an external content inspection server (CIS) in the traffic inspection process. Any Single Firewall or Firewall Cluster can be configured to forward traffic to an external CIS.

The following sections are included:

Getting Started with External Content Inspection

Prerequisites: Install and configure the CIS server before you configure CIS redirection

You can redirect traffic to an external content inspection server (CIS), such as a virus scanning server, for content screening before the traffic continues to its final destination.

What CIS Redirection does

Connections arriving at the Firewall can be redirected (using NAT) to a content inspection server that works as a proxy. The content inspection server filters out unwanted traffic before inspected traffic is forwarded through the Firewall toward its original destination (possibly with a second NAT operation).

Limitations

- CIS redirection is only supported on Single Firewalls and Firewall Clusters.
- Only the FTP, HTTP, and SMTP Protocol Agents can be used for redirecting traffic to a content inspection server.
- You can apply either external content inspection (CIS redirection) or internal content inspection (TLS Inspection, web filtering, internal virus scanning, deep inspection) on traffic. It is not possible to use both forms of content inspection simultaneously on the same connection.
- Only IPv4 addresses are supported in CIS redirection.

Configuration Overview

Illustration 49.1 Elements in the Configuration

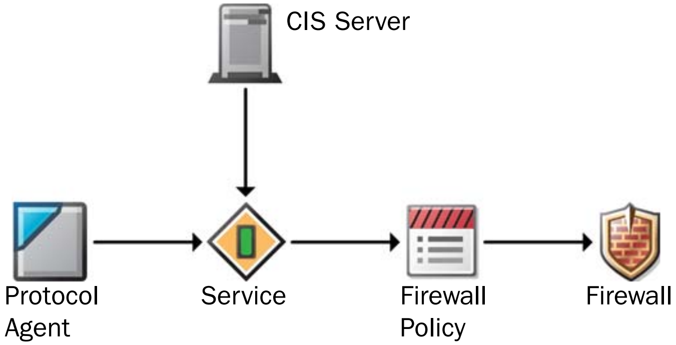

1. Define the CIS Server element to define the IP address and services the server inspects. See Defining a Content Inspection Server Element (page 811).
2. Create a custom Service element that includes the CIS Server. See Defining a Service for CIS Redirection (page 812).
3. Define the Access rules that select traffic for CIS redirection. See Defining Access Rules for CIS Redirection (page 813).
4. Define the NAT rules for CIS redirection. See Defining NAT Rules for CIS Redirection (page 814).

Defining a Content Inspection Server Element

Prerequisites: Install and configure the CIS server before you configure CIS redirection

The CIS Server element defines the IP address and the ports for the service(s) on the server.

▼ To define a CIS Server element

1. Select **Configuration→Configuration→Security Engine**. The Security Engine Configuration view opens.
2. Expand the **Network Elements** branch.
3. Right-click **Servers** and select **New→CIS Server.** The CIS Server Properties dialog opens.

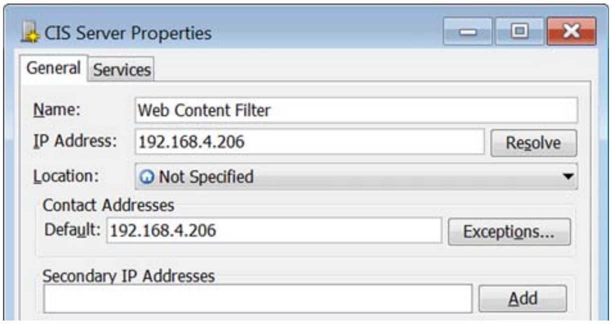

Name: Web Content Filter
IP Address: 192.168.4.206
Location: Not Specified
Contact Addresses
Default: 192.168.4.206
Secondary IP Addresses

4. Specify a unique **Name** and the **IP Address** of the server.
5. (*Optional*) Specify the **Location** and the NAT addresses for the server.
   - The Location and NAT addresses need to be modified if there is a NAT device between a Firewall and the server, so that the Firewall cannot connect directly to the IP address defined for the server. For more, see Defining Locations (page 66) and Defining Contact IP Addresses (page 66).
6. (*Optional*) Enter a **Secondary IP Address** for the server.
   - The Firewall associates the secondary IP address to the correct element when the IP address is used as the source or destination address in pass-through communications, but the secondary IP address is never used as the destination address when the Firewall initiates the communications.
7. Switch to the **Services** tab.

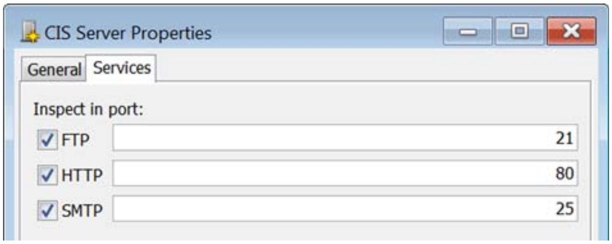

Inspect in port:
☑ FTP 21
☑ HTTP 80
☑ SMTP 25

8. Select the protocol to be inspected and enter the **Port Number** that the server uses for the service.
9. Click **OK**.

What's Next?

▶ Continue the configuration by Defining a Service for CIS Redirection (page 812).

Defining a Service for CIS Redirection

Prerequisites: Defining a Content Inspection Server Element

Creating a Service for CIS Redirection

▼ To define a service for CIS redirection

1. Select **Configuration→Configuration→Security Engine**. The Security Engine Configuration view opens.
2. Expand the **Other Elements→Services** branch.
3. Right-click the **TCP** branch and select **New→TCP Service**. The TCP Service Properties dialog opens.

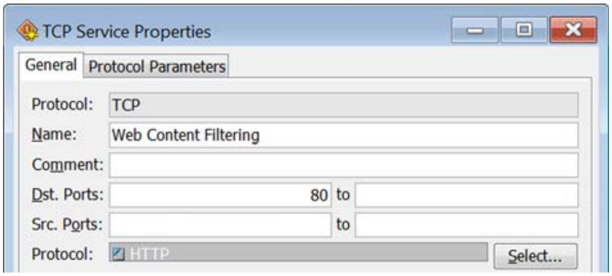

Protocol: TCP
Name: Web Content Filtering
Comment:
Dst. Ports: 80 to
Src. Ports: to
Protocol: HTTP

4. Enter a unique **Name**.
5. Enter the destination port or port range in **Dst. Ports** and/or enter the source port or port range in **Src. Ports**. Leave fields empty as necessary.
6. Click the **Select** button for the **Protocol** field and select the Protocol Agent.
   - The **FTP**, **HTTP**, or **SMTP** Protocol Agents can be used for CIS redirection.

What's Next?

▶ Continue the configuration by Defining Protocol Parameters for CIS Redirection.

Defining Protocol Parameters for CIS Redirection

▼ To define CIS redirection protocol parameters

1. Switch to the **Protocol Parameters** tab.
2. Click the **Edit** button for **Redirect Connections to CIS**. The CIS Server Entry dialog opens.

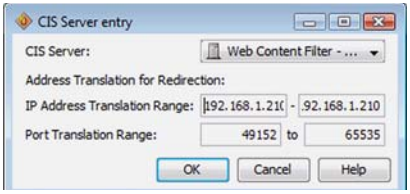

CIS Server: Web Content Filter - ...
Address Translation for Redirection:
IP Address Translation Range: 192.168.1.21(- .92.168.1.210
Port Translation Range: 49152 to 65535

3. Select the **CIS Server**.
4. Specify the **IP Address Translation Range** and the **Port Translation Range** used for source NAT of the CIS-redirected connections. When the traffic is sent from the Firewall to the CIS server, the addresses are translated as defined here.
   - To specify a single IP address, enter the same IP address in both fields. Only IPv4 addresses are supported in CIS redirection.

Note – NAT is mandatory for CIS redirection to forward the traffic to the CIS server. Your NAT rules must not contain overlapping definitions.

5. Click **OK**. The CIS Server Entry dialog closes.
6. Click **OK**. The TCP Service Properties dialog closes.

What's Next?

▶ Continue the configuration in Defining Access Rules for CIS Redirection.

Defining Access Rules for CIS Redirection

Prerequisites: Defining a Content Inspection Server Element, Defining a Service for CIS Redirection

Typically, two IPv4 Access rules are needed when the traffic to and from the CIS server goes through the Firewall. One rule redirects the matching traffic to the CIS server and another rule allows the forward connection from the CIS server to the actual destination.

▼ To create IPv4 Access rules for CIS redirection

1. Select **Configuration→Configuration→Security Engine**. The Security Engine Configuration view opens.
2. Expand the **Policies** branch and select **Firewall Policies**.
3. Right-click a policy and select **Edit Firewall Policy**. The policy opens for editing.
4. Add a new IPv4 Access rule that redirects the traffic to the CIS server:

Table 49.1 Access Rule for Redirecting Traffic to CIS Server

Source	Destination	Service	Action
Original source address of the traffic to be inspected (such as clients in the internal network).	Original destination address of the traffic (such as a web server).	Your custom Service element with the CIS redirection.	Allow

5. Add a new rule that allows forward connections from the CIS server to the actual destination. This rule is not needed if the forward connection does not go through the Firewall.

Table 49.2 Access Rule for Forwarding Traffic From CIS Server

Source	Destination	Service	Action
CIS server's address.	Original destination address (such as a web server).	Normal service element without CIS redirection. (The Agent with redirection is used only for the connection that is redirected, not the forward connection).	Allow

What's Next?

- Continue the configuration in Defining NAT Rules for CIS Redirection.

Defining NAT Rules for CIS Redirection

Prerequisites: See Getting Started with External Content Inspection

The address translation defined in your custom Service for CIS redirection must not overlap with NAT rules. The NAT rules outlined here ensure that no duplicate NAT is applied and define the NAT applied when connections return from the external server.

To create NAT rules for CIS redirection

1. Add a "no-NAT" rule for the CIS redirection to ensure no conflicts with NAT rules. Place the following rule as the first rule in the IPv4 NAT rules table:

Table 49.3 No-NAT Rule

Source	Destination	Service	NAT
The original source address of the traffic to be redirected (such as clients in the internal network, source network, etc.).	The original destination address of the traffic (such as a web server).	Any Service that matches the traffic	Leave empty to ensure that no NAT is applied.

2. If NAT is needed for the forward connection from the CIS server to the actual destination, add a second NAT rule:

Table 49.4 NAT Rule for Inspected Traffic

Source	Destination	Service	NAT
The CIS server's address.	The original destination address of the traffic (such as a web server).	Any Service that matches the traffic	Define the NAT that would be used even without CIS redirection (such as dynamic source NAT with NetLinks for outbound connections.)

What's Next?

- If your CIS server is ready to inspect the traffic, save and install the policy to start using the new configuration.

Related Tasks

- Editing Firewall NAT Rules (page 687)

Chapter 50

Blacklisting IP Addresses

Blacklists contain entries for blocking traffic temporarily based on traffic patterns that the Security Engines detect or on administrator commands. Firewalls, IPS engines, Layer 2 Firewalls, and Virtual Firewalls can use a blacklist for blocking traffic.

The following sections are included:

Getting Started with Blacklisting

Prerequisites: None

Blacklisting is a way to temporarily block unwanted network traffic either manually or automatically with blacklist requests from a Security Engine or Log Server.

What Blacklisting Does

Blacklisting allows you to temporarily stop traffic:

- without editing and installing policies (*manual blacklisting only*)
- based on events detected by Security Engines
- based on correlation of detected events
- on a different engine than the one that detects an event
- on multiple engines with a single administrator command or a single detected event.

Security Engines can add entries to their own blacklists based on detected events in the traffic they inspect. Security Engines and Log Servers can also send blacklist requests to other Security Engines. You can also blacklist IP addresses manually in the Management Client.

Example **If a rule in the Inspection Policy detects a serious attack against a single host in your internal network, you may want the rule to trigger automatic blacklisting of connections from that host to any other host in your internal networks.**

Limitations

- Blacklisting is only supported for IPv4 traffic on Layer 2 Firewalls.
- Virtual Firewalls cannot send blacklist requests to other Virtual Firewalls or Security Engines.

What Should I Know Before I Begin?

By default, the blacklist is not enforced at all. To enforce the blacklist, you must define the point(s) at which the blacklist is checked in the Access rules. If a connection is allowed by a rule placed above the blacklist rule, the connection is allowed regardless of the blacklist entries.

Automatic blacklisting can have unintended consequences that could disrupt business-critical traffic. Legitimate traffic can be incorrectly identified as malicious if the pattern for detecting malicious traffic is inaccurate. If an attacker uses spoofed IP addresses, a legitimate IP address may be blacklisted instead of the attacker’s actual IP address, causing a self-inflicted denial of service (DoS). Potential benefits must be weighed against the potential risk of legitimate communications being blocked in the process. Use automatic blacklisting with careful consideration.

Configuration Overview

1. Define which traffic is matched against the blacklist in the Access rules. See Enabling Blacklist Enforcement (page 817).
2. (*Automatic blacklisting only*) Define the traffic that you want to blacklist automatically in the Exceptions in the Inspection Policy. See Defining Which Traffic is Blacklisted Automatically (page 818).

After blacklisting is configured, you can monitor the currently active blacklist as explained in Monitoring Connections, Blacklists, VPN SAs, Users, and Routing (page 111).

Enabling Blacklist Enforcement

Prerequisites: None

Access rules define which connections are checked against the blacklist. By default, the blacklist is not enforced at all. To enforce the blacklist, you define the point(s) at which the blacklist is checked in the Access rules. Blacklisting is applied with Access rules that contain the Apply Blacklist action. Only traffic that matches such a rule is checked against the blacklist entries. If the traffic checked against the blacklist does not match any of the blacklist entries, the next rule in the policy is checked as usual.

You can have several Apply Blacklist rules with different matching criteria at different points in the policy.

▼ To enable blacklist enforcement

1. Open the Firewall, Layer 2 Firewall, or IPS Policy for editing.
   - Blacklist enforcement for Virtual Firewalls is configured in the Firewall Policy that is used on the Virtual Firewall.
2. Switch to the **IPv4 Access** or **IPv6 Access** tab, and define which **Sources**, **Destinations**, and **Services** are compared with the blacklist.

Note – Blacklisting is only supported for IPv4 traffic on Layer 2 Firewalls.

3. Right-click the **Action** cell and select **Apply Blacklist**.
4. (*Optional*) If you want to restrict which engines and servers are allowed to send blacklist requests, right-click the **Action** cell and select **Edit Options**. The Select Rule Action Options dialog opens. Configure the settings as described below and click **OK**:

Setting	Configuration
Allowed Blacklisters for This Rule	Select Restricted.
Allowed Blacklisters	Select the elements from **Available Blacklisters** that you want to add as Allowed Blacklisters and click **Add**. - Add the Management Server to allow manual blacklisting through the Management Clients. - Add the Log Server to allow it to relay blacklisting requests from other Security Engines.

Note – By default, Security Engines are allowed to add entries directly to their own blacklists for traffic they inspect.

5. Install the policy on the engine to activate the changes.

What's Next?

▶ If you want to configure automatic blacklisting based on events detected by the Security Engines, see Configuring Automatic Blacklisting (page 818).

Related Tasks

- ▶ Editing Access Rules (page 672)
- ▶ Blacklisting Traffic Manually (page 821)

Configuring Automatic Blacklisting

Prerequisites: Enabling Blacklist Enforcement

Security Engines trigger automatic blacklisting based on the Blacklist Scope options in the Exceptions in the Inspection Policy. Security Engines add entries directly to their own blacklists for traffic they inspect. Security Engines can also send blacklisting requests to other Security Engines. In this case, the Security Engine sends the blacklisting request to the Log Server. The Log Server relays the blacklisting request to the Management Server, and the Management Server relays the blacklisting request to the other Security Engine(s) that enforce the blacklisting.

What's Next?

- ▶ Defining Which Traffic is Blacklisted Automatically (page 818).

Defining Which Traffic is Blacklisted Automatically

You can define blacklisting as an option for Exceptions in the Inspection Policy. Security Engines generate blacklist entries based on the patterns they detect in the traffic flow. The blacklist entry that is sent identifies traffic based on IP addresses and optionally the Protocol and port.

The blacklist entries can be configured to include whole networks, even if the detected events that trigger them are related to a single source or destination IP address.

Adding a Rule for Automatic Blacklisting

▼ To add a rule for automatic blacklisting

1. Right-click the Inspection Policy and select **Edit Inspection Policy**. The Inspection Policy opens for editing.
2. Switch to the **Exceptions** tab.
3. Add an Exception Rule with the **Situation**, **Source**, **Destination**, and **Protocol** that match traffic you want to blacklist.
4. Right-click the **Action** cell and select **Terminate**.
5. Right-click the **Action** cell and select **Edit Options**. The Select Rule Action Options dialog opens.

What's Next?

- ▶ Proceed to Defining Blacklisting Rule Action Options (page 819).

Defining Blacklisting Rule Action Options

▼ To define blacklisting rule Action options

1. Switch to the **Blacklist Scope** tab.

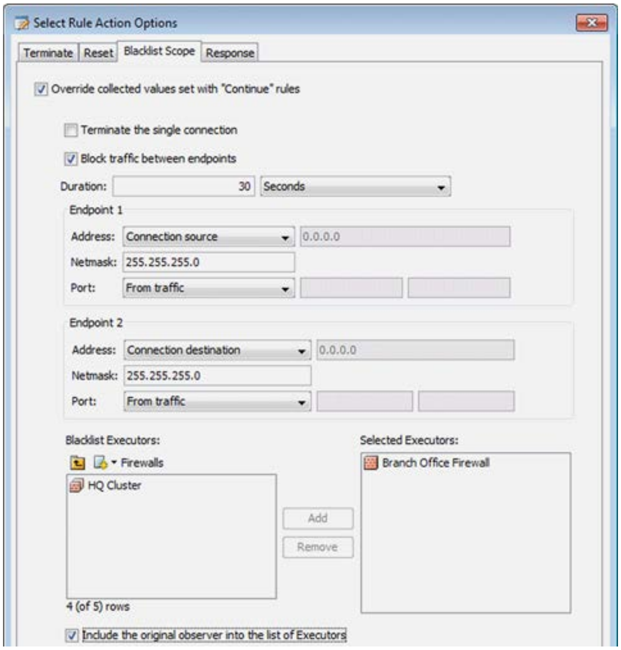

☑ Override collected values set with "Continue" rules

□ Terminate the single connection

☑ Block traffic between endpoints

Duration: 30 Seconds

Endpoint 1

Address: Connection source 0.0.0.0

Netmask: 255.255.255.0

Port: From traffic

Endpoint 2

Address: Connection destination 0.0.0.0

Netmask: 255.255.255.0

Port: From traffic

Blacklist Executors: Firewalls

HQ Cluster

4 (of 5) rows

Selected Executors: Branch Office Firewall

☑ Include the original observer into the list of Executors

2. Select **Override collected values set with “Continue” rules**.
3. Select the type of Blacklist entry to create:
   - Select **Terminate the Single Connection** to create a Blacklist entry that terminates only the current connection using the default options. Proceed to Step 4.
   - Select **Block Traffic Between Endpoints** to block the traffic for defined duration and configure the settings as explained below.

Table 50.1 Block Traffic Between Endpoints Options

Setting	Configuration
Duration	Specify how long the entry is kept on the Security Engine’s blacklist.
Address	Select the type of address to match for Endpoint 1 and Endpoint 2

Table 50.1 Block Traffic Between Endpoints Options (Continued)

Setting		Configuration
	Any	Matches any IP address.
	Attacker/Victim	Matches the IP address identified as the originator/target of an attack by the Situation element that is triggered.
	IP Source/IP Destination	Matches the IP address that is the source/destination of the packet(s) that trigger the detected situation.
	Connection Source/ Connection Destination	Matches the IP address that is the source/destination of the TCP connection that triggers the detected situation.
	Predefined	Matches only the fixed IP address you enter in the field to the right of the Address type list.
Netmask		Define the range of IP addresses in the same network as the source IP address that is blacklisted for Endpoint 1 and Endpoint 2. For example, the netmask 255.255.255.0 blacklists all the addresses in the same C-class network. The default netmask 255.255.255.255 blacklists only one specific IP address.
Port		Select the Port type to blacklist for Endpoint 1 and Endpoint 2
	Ignored	Matches any IP traffic regardless of the protocol or port.
	From traffic	Matches the IP protocol and the port number in the traffic that triggered the blacklist entry.
	Predefined TCP	Matches only TCP traffic through the TCP port or the range of TCP ports that you enter in the fields to the right of the Port type list.
	Predefined UDP	Matches only UDP traffic through the UDP port or the range of UDP ports that you enter in the fields to the right of the Port type list.

4. Select **Blacklist Executors** (the engines where the blacklist entry is sent) and click **Add**.
5. (*Optional*) Select **Include the Original Observer in the List of Executors** to include the engine that detects the situation in the list of blacklist executors.
6. Click **OK**.

Related Tasks

- Blacklisting Connections Manually (page 228)
- Monitoring Connections, Blacklists, VPN SAs, Users, and Routing (page 111)

Blacklisting Traffic Manually

Prerequisites: Enabling Blacklist Enforcement

You can blacklist traffic manually by creating blacklist entries in the Blacklist view, Connections view, Monitoring view, and Logs view. See Blacklisting Connections Manually (page 228).

Related Tasks

▶ Monitoring Connections, Blacklists, VPN SAs, Users, and Routing (page 111)

USERS AND AUTHENTICATION

In this section:

CHAPTER 51

SETTING UP DIRECTORY SERVERS

A directory server provides access to information about user accounts in a user database. Both internal and external directory servers can be used. Directory servers can be used for user authentication with Firewalls (not with IPS engines or Layer 2 Firewalls).

The following sections are included:

Getting Started with Directory Servers

Prerequisites: None

A directory server is a server that stores user information in a user database. The directory server is queried during the user authentication process. Authentication is based on the user information, but is a separate process and is not necessarily done on the same server that stores the user information.

Optionally, an integrated external Active Directory Server can be used with the User Agent for Access Control by User to allow the use of Active Directory users as the source and destination of rules. See Enabling Access Control by User (page 838) for more information.

Limitations

- The internal LDAP user database does not allow external authentication servers or the Authentication Server component to query user information. An external user database is always required if you want to use external authentication servers with the SMC.
- The internal LDAP database limits the length of the User and User Group DN (distinguished name) to a maximum of 254 characters. Check the restrictions of external LDAP servers from the external server's documentation.
- If there are administrative Domains configured, the internal user database is always in the Shared Domain and the user accounts stored in the internal database are also always in the Shared Domain. If you want to limit the visibility of end-user accounts, you must configure external LDAP database(s) separately for each Domain.
- User authentication is only supported on Firewalls (not with IPS engines or Layer 2 Firewalls).

What Do I Need to Know Before I Begin?

- The Management Server has an internal LDAP user database.
- Alternatively, you can use external LDAP user databases (including Active Directory).
- The optional Authentication Server component includes its own proprietary user database, which is not based on LDAP. If you want to authenticate users with the authentication methods offered by the Authentication Server, you must link users in an external LDAP user database to user accounts in the Authentication Server's user database.
- Different users can be stored in different databases.

See the *Firewall/VPN Reference Guide* for more information on the different configuration options.

Depending where user information is stored, different authentication options are available. The following table explains the possible combinations of internal and external directory servers and authentication servers:

Table 51.1 Combinations of Internal and External Directory Servers and Authentication Servers

		Authentication Server	
		Internal	**External**
Directory Server	**Internal**	User and User group information is maintained using the Management Client and can be used for creating rules. Authentication can be done with password, IPsec certificate, or pre-shared key.	To make the configuration work, a second, external user database is required because the external authentication server has no access to the internal database. The same user information is maintained separately in the Management Server's internal user database and in the external user database. User and User Group information can be used for creating rules. Any authentication method supported by the external authentication server can be used.
	External	The Management Server is defined as an LDAP client for the external server. User and User Group information is shown in the Management Client, and can be used for creating rules. Authentication can be done with password, IPsec certificate, or pre-shared key.	You can optionally introduce the user information to the Management Server. To do this, define the Management Server as an LDAP client (see the table cell on the left), or manually maintain the information in two databases (see the table cell above). If you do not make the user information available to the Management Server, each authentication rule includes all users. There can be several rules, but any user that can authenticate in one rule can also authenticate when any of the other rules is triggered. Any authentication method supported by the external authentication server can be used.

Configuration Overview

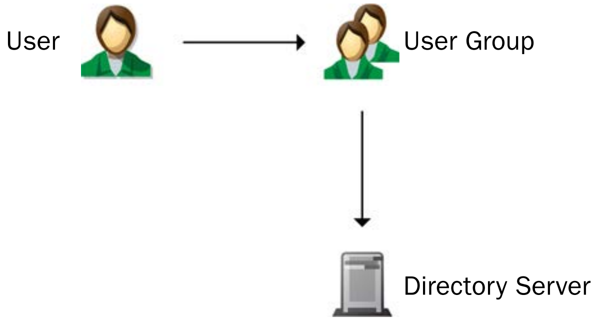

Illustration 51.1 Elements in the Configuration

1. (*Optional*) Integrate an external LDAP directory server:
   - **1a.** Define an LDAP Server or Active Directory Server element. See Integrating External Directory Servers. The Active Directory Server element also contains the settings for IAS authentication; generic LDAP Server elements define a directory server only.
   - **1b.** (*Optional*) Define an LDAP Domain. See Defining LDAP Domains (page 837).
2. Define the User Group and User information:
   - **2a.** To import existing user information from some other Management Server, see Importing and Exporting User Information (page 852).
   - **2b.** To create new accounts or modify accounts stored in an external LDAP database, see Defining User Accounts (page 843).

Integrating External Directory Servers

Prerequisites: None

You can use an external directory server to store user group and user information instead of or in addition to the internal user database.

The Management Server and the firewall engines each use their own integrated LDAP client to query the external LDAP directory directly. The External LDAP directory is not replicated into the internal directory on the Management Server or into the local directory of the firewalls. Instead, the external LDAP directory is queried separately each time by the Management Server when you view the User elements in the Management Client and by the firewalls when a user attempts to authenticate.

You can also use an external directory server without integrating it with the Stonesoft components. The Firewall relays the information it receives from the user to the external authentication server and it is up to the external authentication server to check the supplied information and report back to the Firewall whether authentication succeeds or fails.

In this configuration, a single element (User element named *external*) is used to represent all externally stored users in the Firewall's policy. It is not possible to create different rules for different externally stored users.

The Management Server and the Firewall engines both have an integrated LDAP client that can query the external directory server. You can configure access to the directory server for both the Management Server and the Firewall, or for the Firewall only. To take full advantage of user authentication features, we recommend configuring access to the directory server for both the Management Server and the Firewall.

Configuring access to the external directory server for both the Firewall and the Management Server allows the following:

- There is no need to manually duplicate user account information. User and User Group elements are automatically added to the SMC from the external directory.
- Externally stored user accounts are shown in the Management Client and can be used to create different rules for different users.
- In most cases, users can be also added, removed, and modified through the Management Client.
- Internal authentication methods can be used to authenticate externally stored users.

If only the Firewall engines can access the external directory server, the following restrictions apply:

- You can authenticate externally stored users only against authentication methods provided by an external authentication server or the Authentication Server component. Internal authentication methods are not available for externally stored users.
- A single element (User element named *external*) is used to represent all externally stored users in the Firewall's policy. It is not possible to create different rules for different externally stored users.

You can view user information and use it for authentication against an external authentication service simply by allowing the Stonesoft components to connect to the LDAP database. To be able to authenticate users with an LDAP password or to be able to modify information in the LDAP directory through the Management Client, you must add parameters for the Stonesoft components on the LDAP directory server by configuring schema files.

What's Next?

▶ To authenticate users with an LDAP password or to modify information in the LDAP directory through the Management Client, proceed to Configuring Schema Files on External Directory Servers.

▶ To view user information and/or use it for authentication against an external authentication service, proceed to Defining LDAP Server Elements (page 831) or Defining Active Directory Server Elements (page 830).

Configuring Schema Files on External Directory Servers

A schema file defines the attributes (individual pieces of data) that an account can contain. Extending the external server’s schema with Stonesoft-specific attributes is optional, but extending the schema also allows you to add Stonesoft-specific information to Users and User Groups through the Management Client. You must update the schema file if you want to configure authentication requirements for specific Users or User Groups. Otherwise, you can configure authentication only at the LDAP domain level.

The method of configuring Schema files varies depending on which LDAP server you are using. The Schema update is done outside the Management Client. In general, the schema update means that you add the Stonesoft-specific attributes to the existing user information on the external LDAP server. These include attributes for the Stonesoft-specific user name, password, and allowed authentication methods for the user. These are documented in more detail in Schema Updates for External LDAP Servers (page 1209).

What’s Next?

▶ If the external directory server is an Active Directory server, continue by Defining Active Directory Server Elements (page 830).

▶ For other types of LDAP servers, see Defining LDAP Server Elements (page 831).

Defining Active Directory Server Elements

The Active Directory Server element contains both the user database and authentication service options needed to use a Microsoft Active Directory server to store and authenticate users.

▼ To define an Active Directory server element

1. Select **Configuration→Configuration→User Authentication**. The User Authentication Configuration view opens.
2. Right-click **Servers** and select **New→Active Directory Server** from the menu. The Active Directory Server Properties dialog opens.
3. Configure the General settings as described below:

Setting	Description
Name	Enter a unique name for the element.
IP Address	The IP address of the server. Only IPv4 addresses are supported.
Location/ Contact Addresses	A Location and Contact Address are needed if NAT is applied between a Firewall or Management Server and the Active Directory server. For more information, see Defining Locations (page 66) and Defining Contact IP Addresses (page 66).
Timeout	Enter the time (in seconds) that Stonesoft components wait for the server to reply.

What’s Next?

▶ Continue by Configuring LDAP Connection Settings (page 832).

Defining LDAP Server Elements

The LDAP Server element can be used to configure access to any LDAP server as a user database for the Firewalls and/or the Management Server.

▼ To define an LDAP Server element

1. Select **Configuration→Configuration→User Authentication**. The User Authentication Configuration view opens.
2. Right-click **Servers** and select **New→LDAP Server**. The LDAP Server Properties dialog opens.

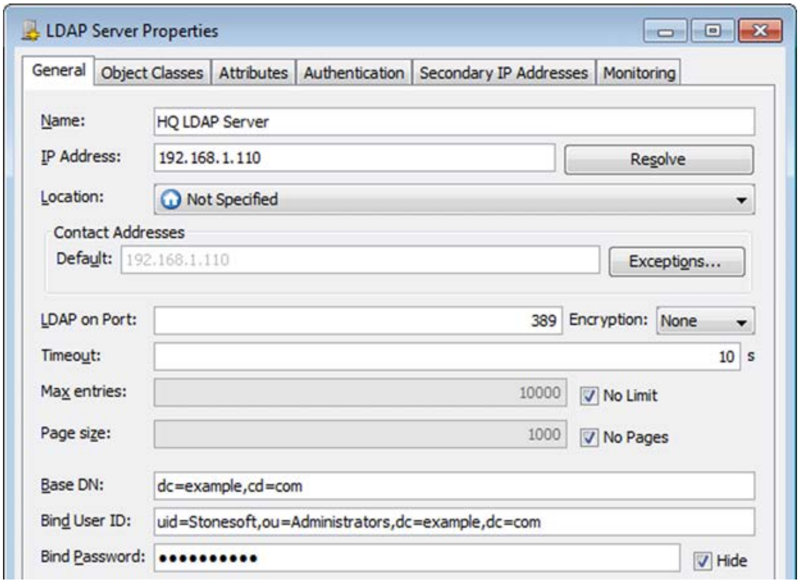

Name: HQ LDAP Server
IP Address: 192.168.1.110
Location: Not Specified
Default: 192.168.1.110
LDAP on Port: 389 Encryption: None
Timeout: 10 s
Max entries: 10000 ☑ No Limit
Page size: 1000 ☑ No Pages
Base DN: dc=example,cd=com
Bind User ID: uid=Stonesoft,ou=Administrators,dc=example,dc=com
Bind Password: •••••••••• ☑ Hide

3. Configure the General settings as described below:

Setting	Description
Name	Enter a unique name for the element.
IP Address	The IP address of the server. Only IPv4 addresses are supported.
Location/ Contact Addresses	A Location and Contact Address are needed if NAT is applied between a Firewall or Management Server and the LDAP server. For more information, see Defining Locations (page 66) and Defining Contact IP Addresses (page 66).
Timeout	Enter the time (in seconds) that Stonesoft Components wait for the server to reply.

4. (*Optional*) Switch to the **Secondary IP Addresses** tab and add one or more secondary IP addresses for the server. These IP addresses are included when the LDAP Server element is used in rules, but Stonesoft components never use the addresses when initiating contact with the server.

What's Next?

▶ Continue by Configuring LDAP Connection Settings (page 832).

Configuring LDAP Connection Settings

The Active Directory or LDAP Server element's LDAP connection settings include an account that the Management Server and Firewall engines use to connect to query the directory. The user account must exist in the Active Directory or LDAP Server's user database. Make sure the account you use has the privileges to manage other user accounts.

There are also values for some attributes that the Management Server and Firewall engines look for in the directory. We recommend that you do not use special characters or accented letters in the Distinguished Names or user ID attributes. Currently, Active Directory has a limit of 160 characters for the Base DN, 24 characters for a UID (user ID) and 64 characters for the OU (organizational unit). For other LDAP servers, check the documentation of your LDAP server.

▼ To configure LDAP connection settings

➥ Configure the Connection settings as described below:

Table 51.2 Connection Settings

Setting	Description
Base DN	Enter the LDAP tree under which the authenticating users' accounts are stored. **Example** (DNS-based tree) dc=example,dc=com **Example** (*"O-based" tree used, for example, in Novell eDirectory*) ou=astronauts,o=government,st=Florida,c=US

Bind User ID	Define the Distinguished Name of the User ID that the Firewalls and Management Servers use to connect to the LDAP server. This user account must exist in the user database. Make sure the account you use has the privileges to manage other user accounts.
	Example (*DNS-based tree*)
	uid=Stonesoft,ou=Administrators,dc=example,dc=com
	Example (*"O-based" tree used, for example, in Novell eDirectory*)
	uid=Stonesoft,ou=Administrators,ou=astronauts,
	o=government,st=Florida,c=US

Bind Password	Enter the password for the user account that the Firewalls and Management Servers use to connect to the LDAP server.
LDAP on Port	(*Optional*) Enter the port number if the server communicates on a port other than the default port (TCP port 389). The predefined Firewall Template allows the engines to connect to the default port. If you change to a custom port, you must add a new Access Rule to allow the traffic.
LDAP StartTLS	(*Recommended*) Select this option to enable Transport Layer Security (TLS) for connections to the LDAP server.

What's Next?

- If your Active Directory or LDAP server has LDAP object classes that are not defined in the SMC by default, continue by Adding LDAP Object Classes (page 833).
- If you do not want to add object classes, and you want to define how attributes in the LDAP directory are mapped, proceed to Configuring LDAP Attribute Mapping (page 834).
- Otherwise, proceed to Adding Authentication Methods (page 836)

Adding LDAP Object Classes

If your Active Directory or LDAP server has LDAP object classes that are not defined in the SMC by default, you must add those object classes to the LDAP Object classes in the server properties. This way, the existing classes on the Active Directory or LDAP server can also be used for authentication.

▼ To add an object class

1. Switch to the **Object Classes** tab.

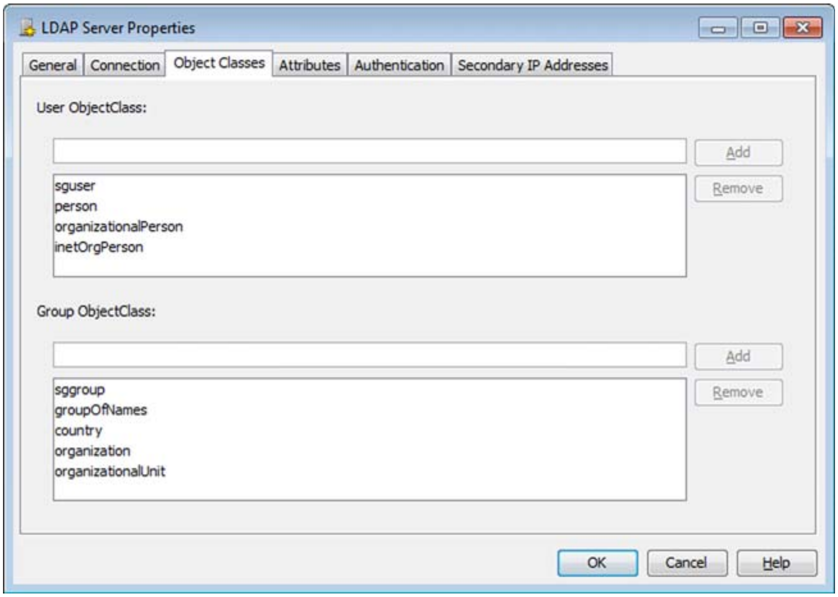

User ObjectClass:

sguser
person
organizationalPerson
inetOrgPerson

Group ObjectClass:

sggroup
groupOfNames
country
organization
organizationalUnit

2. Enter the name of the **User Object Class** or **Group Object Class** and click **Add**. The object class appears in the list. Repeat to add more object classes.

What's Next?

- If you want to define how attributes in the LDAP directory are mapped, continue by Configuring LDAP Attribute Mapping.
- Otherwise, proceed to Adding Authentication Methods (page 836).

Configuring LDAP Attribute Mapping

The settings on the Attributes tab allow you to define how attributes in the LDAP directory are mapped to user properties in the SMC. You may need to change or fill in these values according to the server's configuration. Enter the same values for the attributes that are defined in the LDAP Server's schema file.

▼ To configure LDAP attributes

1. Switch to the **Attributes** tab.

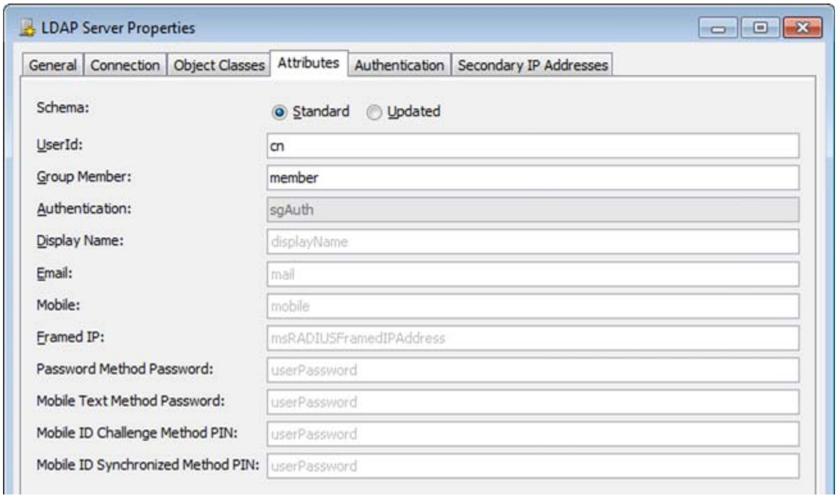

Schema: ● Standard ○ Updated
UserId: cn
Group Member: member
Authentication: sgAuth
Display Name: displayName
Email: mail
Mobile: mobile
Framed IP: msRADIUSFramedIPAddress
Password Method Password: userPassword
Mobile Text Method Password: userPassword
Mobile ID Challenge Method PIN: userPassword
Mobile ID Synchronized Method PIN: userPassword

2. Configure the Attribute settings as described below:

Table 51.3 Attribute Settings

Setting	Description
Schema	Select **Standard** if you are using the external LDAP server's standard schema files. Select **Updated** if you have extended the schema file to include Stonesoft-specific attributes. See Configuring Schema Files on External Directory Servers (page 830).
UserId	The name that the server uses for the UserID Attribute.
Group Member	The name that the server uses for the Group Member Attribute. By default, the attribute is set to *member* for standard schema, and *sgMember* for updated schema,
Authentication (*Updated Schema only*)	The Authentication Attribute for storing the authentication method information. By default, the attribute is set to *sgauth*.
Display Name (*Updated Schema only*)	The name that the server uses for the Display Name attribute.
Email (*Updated Schema only*)	(*Updated Schema only*) The name of the attribute for storing the users' e-mail address. This attribute is primarily used for linked Authentication Server Users. It can also be used to identify users by their e-mail address in certificate authentication.

Table 51.3 Attribute Settings (Continued)

Setting	Description
Mobile (*Updated Schema only*)	The name of the attribute for storing the users' mobile phone numbers. This attribute is primarily used for linked Authentication Server users.
Framed IP (*Updated Schema only*)	When the framed IP address option has been enabled for the Authentication Server component, this IP address is sent to an access point for services when a user successfully authenticates. This information can be used in authorization decisions made by the access point.
Password Method Password (*Updated Schema only*)	The name of the password attribute for the Password Authentication Method.
Mobile Text Method Password (*Updated Schema only*)	The name of the password attribute for the Mobile Text Authentication Method.
Mobile ID Challenge Method PIN (*Updated Schema only*)	The name of the PIN attribute for the Mobile ID Challenge Authentication Method.
Mobile ID Synchronized Method PIN (*Updated Schema only*)	The name of the PIN attribute for the Mobile ID Synchronized Authentication Method.

What's Next?

▶ Continue by Adding Authentication Methods.

Adding Authentication Methods

Authentication Methods specify the allowed authentication methods for the users stored on the Active Directory or LDAP server.

You can optionally use the Internet Authentication Service (IAS) in previous Windows Server versions or the Network Policy Server (NPS) in Windows Server 2008 to authenticate end-users. You must configure the IAS/NPS as a RADIUS server, and define each Firewall engine that users contact for authentication as a separate RADIUS client for IAS/NPS (use the NDI addresses on a cluster). The IAS/NPS must have access to user information in the Active Directory. The user accounts must have remote access permissions. Set up the IAS/NPS as explained in the Microsoft Server documentation. The SMC does not support the Message-Authenticator attribute option available in the IAS/NPS, and is not NAP-capable. Only PAP authentication is supported.

▼ To add authentication methods

1. Switch to the **Authentication** tab.
2. Configure the settings according to the type of server:

Type of Server	Setting	Configuration
Active Directory	Use Network Policy Server Method (NPS)	Select this option to use the Windows server's IAS/NPS.
	Port Number	Enter the Port Number for your Windows server's IAS/NPS.
	Shared Secret	Enter the Shared Secret that you have defined for RADIUS clients on the Active Directory server.
	Number of Retries	Specify the Number of Retries. If a Stonesoft component's attempt to connect to the Active Directory server fails, it tries to connect again the specified number of times before giving up on the authentication.
	IP Address (*Optional*)	If the authentication service on the Active Directory server uses a different IP address than the server itself, enter the IP address for authentication.
	External Authentication Methods	Click **Add** and select existing Authentication Method(s), or create new Authentication Method(s) as instructed in Defining Authentication Methods for External Servers (page 860).
LDAP	Authentication Methods	Click **Add** and select existing Authentication Method(s), or create new Authentication Method(s) as instructed in Defining Authentication Methods for External Servers (page 860).

What’s Next?

- ▶ To integrate the Management Server with the Active Directory or LDAP server, click **OK** and continue by Defining LDAP Domains (page 837). If you want to use Active Directory users as the source and destination of rules, you must define an LDAP domain for the Active Directory server.
- ▶ Otherwise, click **OK** and continue by Defining User Accounts (page 843).

Defining LDAP Domains

Each LDAP server has its own LDAP domain in the SMC. One LDAP domain can be selected as the default LDAP domain, so that users can leave out this information when they authenticate (users can type “username” instead of “username@domain”). Users that are stored under non-default LDAP domains must always include the domain in the username.

If administrative Domains have been configured, you can create separate LDAP Domains for each administrative Domain and select one LDAP Domain in each administrative Domain as the Default. Alternatively, you can select one of the LDAP Domains in the Shared Domain as the Default LDAP Domain across all the administrative Domains.

▼ To define a new LDAP Domain

1. Select **Configuration→Configuration→User Authentication**. The User Authentication Configuration view opens.
2. Right-click **Users** and select **New LDAP Domain**. The LDAP Domain Properties dialog opens.

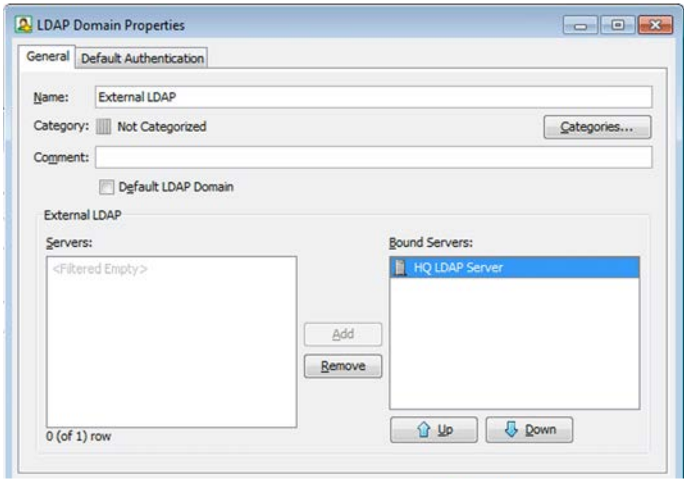

Name: External LDAP

Category: Not Categorized

Comment:

☐ Default LDAP Domain

Bound Servers: HQ LDAP Server

3. Enter the **Name** for the new LDAP Domain.
   - If the LDAP Domain you are creating is not the default LDAP Domain, users must type in the domain name when they authenticate.
4. Select **Default LDAP Domain** if this LDAP Domain will be used for all authentication unless otherwise specified in the IPv4 Access rules.
   - Only one LDAP Domain can be the default LDAP Domain. The selection is automatically cleared from the previous default LDAP Domain.
5. Select a Server and click **Add** to bind the LDAP Server to the LDAP Domain.

6. (*Optional*) Switch to the **Default Authentication** tab and click **Select** to define the allowed Authentication Methods for all accounts in this LDAP Domain.
   - You can override the default setting by selecting different Authentication Method(s) in the User Group or User properties.
   - We recommend that you set a default authentication method. If the Authentication Method you are planning to use is not defined yet, you can return to this dialog to complete the information after your custom Authentication Method is created.
7. Click **OK**.

You have now completed all of the steps required for setting up your external directory server in the SMC.

What's Next?

▶ If you want to use Active Directory users as the source and destination of rules, continue by Enabling Access Control by User.

▶ Otherwise, if you plan to configure user authentication using the SMC's internal authentication methods, continue by Defining User Accounts (page 843).

Enabling Access Control by User

Prerequisites: Defining Active Directory Server Elements, Defining LDAP Domains

Access control by user allows you to use User and User Group elements as the source or destination of a rule to create user-specific rules without requiring user authentication. You can use user-specific rules together with user authentication rules to allow some user groups to access a service, while otherwise requiring authentication for the same service. See Getting Started with Editing the Rules in Policies (page 660) for more information about creating rules.

The User Agent is an optional software component that can be installed on a Windows system in an Active Directory domain that communicates with the Active Directory domain controller to associate users with IP addresses. The User Agent monitors the domain controller's Security Event Log to keep track of when a user logs on or logs off, a user's IP address changes, or a user acquires the same IP address that was previously associated with another user. When the User Agent receives information about a particular user, the User Agent sends an LDAP query to the Active Directory server to look up the user's group membership. The Firewall periodically queries the User Agent for information about which IP address is associated with the user, and which User Group the user belongs to.

For detailed background information about the User Agent, see the *Firewall Reference Guide*.

Limitations

- User-specific rules do not replace user authentication; they are a tool to simplify the configuration of access control, and improve the end-user experience by allowing transparent access to services. They are intended to be used for trusted users in a trusted environment where strong authentication is not required.
- Information about users' IP addresses is cleared from the Firewall's cache if the Firewall has been unable to contact the User Agent for five minutes. This can prevent rules that block connections from matching. For this reason, it is recommended to use Users and User Groups only in rules that allow a connection.

- Users are associated with IP addresses based on logs collected by the Active Directory Domain Controller. For this reason, it is only possible to associate one user with each IP address from a client computer. It is not recommended to use this feature for terminal servers or other computers that have many users logged on at the same time.
- The User Agent periodically sends ICMP echo (ping) requests to users' workstations to monitor which users are active. If a user's workstation does not respond, the user is removed from the list of IP addresses. In cases where ping requests to workstations are not allowed, users' connections may be incorrectly closed. Workstation monitoring can optionally be disabled in the User Agent interface in Windows to prevent this.

Defining the Active Directory Domain Controllers for Access Control by User

The Active Directory Domain Controller settings define which domain controllers the User Agent queries for information from the Windows Security Event Log, and the credentials for logging in to the domain controllers. If you do not already have an Active Directory user account that has permission to execute Windows Management Instrumentation (WMI) queries from a remote computer, you must create a user account before defining the domain controllers.

▼ To define the Active Directory domain controllers

1. Open the Active Directory Server Properties and switch to the **Domain Controllers** tab.

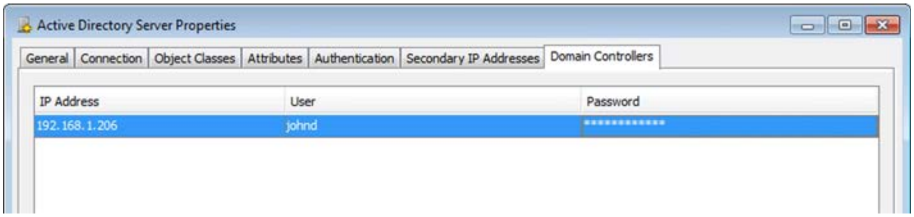

IP Address	User	Password
192.168.1.206	johnd	************

2. Click **Add**. A row is added to the table.
3. Define the settings as described below:

Setting	Configuration
IP Address	Double-click the IP Address cell and enter the IP address of the domain controller. Only IPv4 addresses are supported.
User	Double-click the User cell and enter the user name of a user in the domain that has permission to execute WMI queries from a remote computer. **Note!** Enter only the user name without any domain information. The domain information is automatically added to the user name.
Password	Double-click the Password cell and enter the password for the user account with Domain Admin credentials.

4. Repeat Step 2-Step 3 for any additional domain controllers and click **OK**.

What's Next?

▶ Continue by Creating User Agent Elements (page 840).

Creating User Agent Elements

▼ To create a User Agent element

1. Select **Configuration→Configuration→Security Engine**. The Security Engine Configuration view opens.
2. Browse to **Other Elements→Engine Properties**.
3. Right-click **User Agents** and select **New User Agent**.

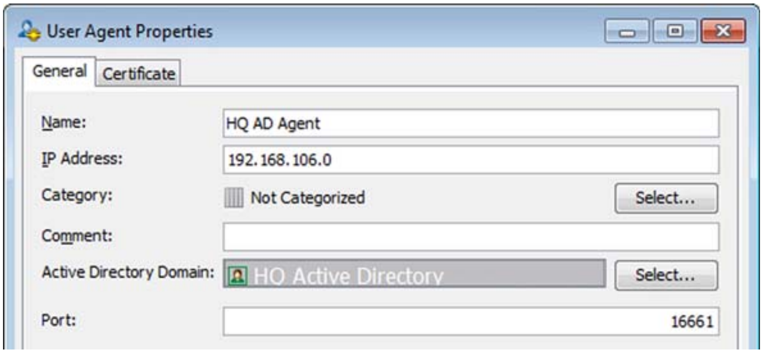

Name: HQ AD Agent

IP Address: 192.168.106.0

Category: Not Categorized

Comment:

Active Directory Domain: HQ Active Directory

Port: 16661

4. Configure the User Agent properties as described below:

Setting	Description
Name	A unique name for the element.
IP Address	The IP address of the Windows Server where the User Agent is installed.
Active Directory Domain	The LDAP domain for the Active Directory users.
Port (*Optional*)	The port on which the User Agent communicates with the Firewall. If you change the port from the default, you must configure the same port in the User Agent Properties on the Windows system. You must also change the rule that allows communication between the Firewall and the User Agent.

5. Click **OK**.

What's Next?

▶ Continue by Selecting User Agents for Security Engines (page 841).

Selecting User Agents for Security Engines

Each User Agent can be associated with one or more Security Engines, but only one User Agent can be selected for each Security Engine. You must select the User Agent in the properties of each Security Engine that communicates with the User Agent before you save and apply the User Agent configuration.

▼ To select a User Agent for a Security Engine

1. Select **Configuration→Configuration→Security Engine**. The Security Engine Configuration view opens.
2. Select **Security Engines**.
3. Right-click the Security Engine for which you want to select a User Agent and select **Properties**.
4. Switch to the **Add-Ons** tab.
5. Select the **User Agent** that this Security Engine communicates with.
6. Click **OK** to close the Security Engine properties.

What's Next?

▶ Continue by Generating a Certificate for a User Agent.

Generating a Certificate for a User Agent

The User Agent configuration file contains the certificates used to secure communication between engines and the User Agent.

Caution – The configuration .zip file contains the private key and certificate for the User Agent, and the SMC Certificate Authority certificate. Handle configuration files securely.

▼ To generate a certificate and save the configuration

1. Select **Configuration→Configuration→Security Engine**. The Security Engine Configuration view opens.
2. Browse to **Other Elements→Engine Properties→User Agents**.
3. Open the properties of the User Agent and switch to the **Certificate** tab.

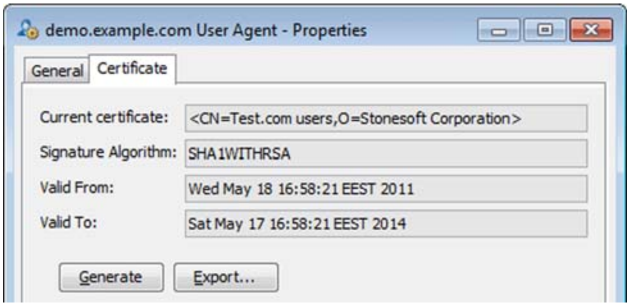

Current certificate: <CN=Test.com users,O=Stonesoft Corporation>
Signature Algorithm: SHA1WITHRSA
Valid From: Wed May 18 16:58:21 EEST 2011
Valid To: Sat May 17 16:58:21 EEST 2014

4. Click **Generate**. A certificate and private key are generated.

5. Click **Export** and save the configuration `.zip` file.

What's Next?

▶ Continue by Allowing Communication With the User Agent (page 842).

Allowing Communication With the User Agent

The connection between the engine and User Agent is opened by the engine. By default, the connection uses port 16661.

▼ To allow communication with the User Agent

➥ Create the following IPv4 Access rule:

Table 51.4 IPv4 Access Rule to Allow Communication With the User Agent

Source	Destination	Service	Action
Firewall or Virtual Firewall element(s)	User Agent element	SG Engine to User Agent Service	Allow

What's Next?

▶ Continue by Installing User Agents.

Installing User Agents

The User Agent software can be installed either locally on Domain Controller or on another Windows system in the domain. If there is only one Domain Controller, and the User Agent is installed on it, no special configuration is needed before installation. If the User Agent is installed on a separate system from the Domain Controller, or there is more than one Domain Controller, you must configure the Windows system and the Domain Controller before installing the User Agent. See Technical Note #7023: User Agent: Preparing the Windows Environment and Installing the User Agent for instructions on configuring the server and the Domain Controller.

User Agent installation requires administrator rights on the Windows system. After installation, the User Agent service must run with sufficient permissions to allow it to monitor the domain controller's Security Event Log.

▼ To install a User Agent

1. Log in to the system where you are installing the User Agent with the correct administrative rights.
2. Transfer the User Agent installation files and the configuration`.zip` file to the computer.
3. Run `UIA_Installer.exe`. The Installation Wizard starts.
4. Click **Next**. The License Agreement opens.
5. Click **I Agree** to accept the license.
6. (*Optional*) Click **Browse** and select the installation folder.
7. Click **Install**. You may be prompted to install additional components that the User Agent requires.

8. Install any additional components as instructed in the Installation Wizards for those components. When the installation returns to the User Agent Installation Wizard, click **Next**.

9. Click **Finish**. The User Agent properties open.

10. Click **Import Configuration** and select the configuration `.zip` file.

11. Click **OK**.

What's Next?

▶ If you plan to configure user authentication using the SMC's internal authentication methods, continue by Defining User Accounts.

▶ Otherwise, the configuration of the User Agent is complete. You can now use User and User Group elements in rules. See Getting Started with Editing the Rules in Policies (page 660).

Defining User Accounts

Prerequisites: See Getting Started with User Authentication for more information

User Group and User elements define the user account information for end-users. User Group and User elements can be inserted in Firewall IPv4 Access rules to add a requirement for authentication as a condition of matching the end-users' connections. If a User Agent is configured for an integrated Active Directory Server, User Group and User elements can also be used as the source and destination of Access, Inspection, and NAT rules without user authentication. See Enabling Access Control by User (page 838).

Options for Adding User Accounts

If you are using the Management Server's internal user database:

- If you have existing user accounts stored in an internal user database on another Management Server, you can export/import the information between the databases (see Importing and Exporting User Information (page 852)).
- Otherwise, create the User Groups and Users individually.

If you are using an external directory server:

- If the LDAP database is integrated with the Management Server, you can view the user information in the Management Client. However, for the accounts to be valid in Access rules, you must configure at least one Authentication Method as allowed for the users. This can be done as a default setting for the LDAP Domain (see Defining LDAP Domains (page 837)) and/or for the User Groups and Users.
- If the LDAP database is not integrated with the Management Server, the user accounts are not shown in the Management Client and are not available for configuration.

If you are using the Authentication Server component:

- You must create user accounts manually in the Authentication Server's internal user database and link them to user accounts on an integrated external directory server. See Linking Authentication Server Users to External Directories (page 847).

What's Next?

- ▶ If you want to use the Authentication Server for user authentication, see Linking Authentication Server Users to External Directories (page 847).
- ▶ If you want to create User Groups for organizing the individual users, see Defining User Groups (page 844).
- ▶ If you do not need new User Groups or your directory server is not integrated with the SMC, you can proceed directly to Defining Users (page 845).

Defining User Groups

If you have many users, you may want to organize your users into several different User Groups. You can organize the groups, for example, according to different services. A single user can belong to several groups at the same time. You must have at least one User Group.

▼ To create a User Group

1. Select **Configuration→Configuration→User Authentication**. The User Authentication Configuration view opens.
2. Expand the branch of the LDAP Domain that represents the correct user database. The internal database is represented by the default **InternalDomain** LDAP Domain.
3. Right-click the parent group below the LDAP Domain (this is called *stonegate* for the internal database) and select **New→User Group**. The User Group Properties dialog opens.

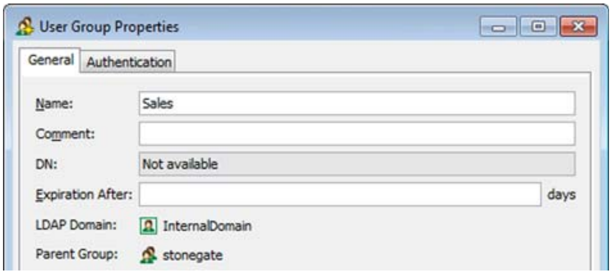

Name: Sales

Comment:

DN: Not available

Expiration After: days

LDAP Domain: InternalDomain

Parent Group: stonegate

4. Enter a **Name** for the User Group.
   - The name is used as the common name (CN) for the group.
   - The distinguished name (**DN**) is inherited from the LDAP Domain to which this User Group belongs.
5. (*Optional*) Enter a number of days after which the User Group expires in the **Expiration After** field. When the User Group expires, it stays in the system but is invalid and does not allow authentication until it is re-enabled.
6. Switch to the **Authentication** tab.

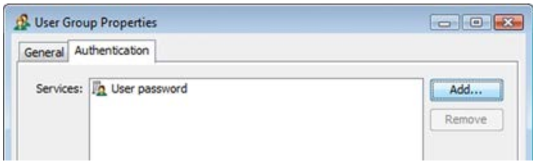

Services: User password

7. Click **Add** and select one or more **Authentication Methods**.
   - If you select several Authentication Methods for the User Group, you can restrict the Authentication Methods allowed in the User element properties and in each Access rule that requires authentication.
8. Click **OK**.

What's Next?

▶ Continue by Defining Users (page 845).

Defining Users

The User element defines who your users are and how they can identify themselves to get access to certain networks and services as defined in your Firewall Access rules. When using the internal user database, you can create the users one by one or import multiple users from an .ldif file to transfer the user information from another Management Server's internal user database (see Importing Users from an LDIF File (page 852)).

If you use an external directory server that is not integrated with the SMC, create a User called *external* in some User Group and use the User Group in rules to represent all users. You must also define the correct authentication method for the *external* User.

Note – Although you cannot edit User Group memberships in the User element properties, each user can belong to several User Groups. After creating the User element, drag and drop it to other User Groups to add more group memberships.

▼ To define a new User

1. Select **Configuration→Configuration→User Authentication**. The User Authentication Configuration view opens.
2. Browse to **Users**.
3. Create a User element in one of the following ways:
   - Right-click a User Group (to add the user to that group) and select **New→User**.
   - Right-click an LDAP Domain (to create a User that is not in any User Group) and select **New User**.

   The User Properties dialog opens.
4. Enter a unique **Name** to identify the User in the directory. The name is used as the common name (CN) for the User. The distinguished name (DN) is inherited from the LDAP Domain to which the User belongs.
5. (*External directory only*) Enter additional user information as described below:

Table 51.5 Additional User Information

Setting	Description
Display Name	The user's full name. User accounts in external directories must always contain a display name.
E-mail Address	An e-mail address where the user can receive notification of changes to the user account. The user account must contain an e-mail address if you use SMTP notification.

Table 51.5 Additional User Information (Continued)

Setting	Description
Mobile Phone Number	A mobile phone number where the user can receive SMS messages for Mobile Text authentication and notification of changes to the user account. The user account must contain a phone number if the Mobile Text Authentication Method is enabled for the user, or if you use SMS notification.
Framed IP (*Authentication Server Users Only*)	When the framed IP address option has been enabled for the Authentication Server component, this IP address is sent to an access point for services when a user successfully authenticates. This information can be used in authorization decisions made by the access point.

6. (*Optional*) Change the Activation settings for the user account as described below:

Setting	Description
Always Active	The user account is considered active immediately and is never automatically disabled.
Active	The user account is considered active starting on the **from** date and optionally ending on the **to** date. If an end date is specified, the user account is automatically disabled on the end date.
Disabled	The user account is not active.

7. Switch to the **Authentication** tab.
8. Click **Add** to select the **Authentication Methods** for the user.
   - You can add more than one authentication method for each user. This way, you can put the User in more than one User Group when the User Groups have different authentication methods.
   - If you have not configured any Authentication Methods yet, you can create them in this dialog.
9. Define the properties for the selected Authentication Method:

Authentication Method	Configuration
External authentication method (including Network Policy Server)	There are no options to configure.
IPsec Certificate	Enter the **Alternative Subject Name or CN** (for example, “alice.smith@example.com” or 192.168.42.103). The value you enter here must match the value entered in the corresponding field of the certificate request.
User Password	Enter and confirm the user’s password.
Pre-Shared Key	Enter and confirm the pre-shared key.

Caution – Use strong passwords that are at least eight characters long and that contain numbers, letters, and special characters. Do not base passwords on personal information such as names, birthdays, ID numbers, phone numbers, street names, registration plate numbers, relatives' names, not even if spelled backwards or partially.

10. Click **OK**.

The user account is created. If the user is stored in the internal LDAP database, the information is automatically synchronized to the local databases on the Firewalls without further action, unless user database replication has been disabled in your system.

What's Next?

- To configure user authentication, proceed to Getting Started with User Authentication (page 856).
- Otherwise, the configuration is complete.

Related Tasks

- Managing User Information (page 852)

Linking Authentication Server Users to External Directories

Authentication Server Users cannot be directly referenced in rules for authentication. Instead, you can create user accounts manually in the Authentication Server's user database and link them to users stored in an integrated external directory. This allows you to reference the User elements in the LDAP Domains of the external directories in rules for authentication. Additionally, linking users allows the Authentication Server to automatically retrieve the values of the LDAP attributes you define in the LDAP Server or Active Directory Server properties.

The number of named users allowed is limited by the Authentication Server's license. User groups cannot be used with the Authentication Server's internal user database.

What's Next?

- Begin by Selecting Domain Nodes for User Linking.

Selecting Domain Nodes for User Linking

To be able to link users, you must specify the domain nodes from which users can be linked, and where to search for the users in the directory.

▼ To specify directory servers for user linking

1. Select **Configuration→Configuration→User Authentication**. The User Authentication Configuration view opens.
2. Open the Authentication Server Properties and switch to the **User Linking** tab.

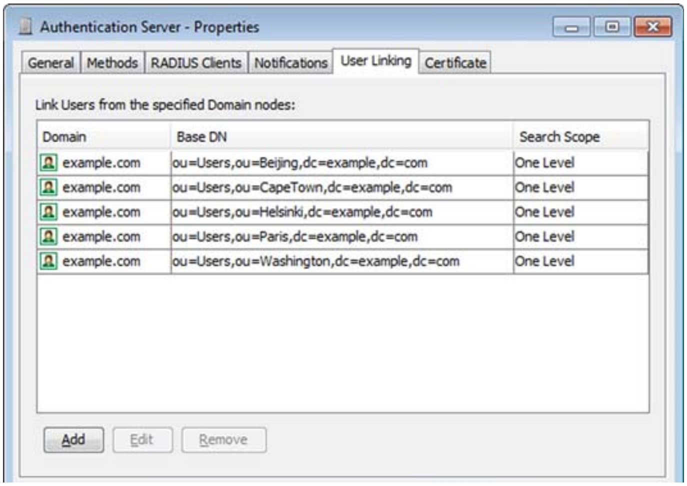

Link Users from the specified Domain nodes:

Domain	Base DN	Search Scope
example.com	ou=Users,ou=Beijing,dc=example,dc=com	One Level
example.com	ou=Users,ou=CapeTown,dc=example,dc=com	One Level
example.com	ou=Users,ou=Helsinki,dc=example,dc=com	One Level
example.com	ou=Users,ou=Paris,dc=example,dc=com	One Level
example.com	ou=Users,ou=Washington,dc=example,dc=com	One Level

3. Click **Add**. The Select Object dialog opens.

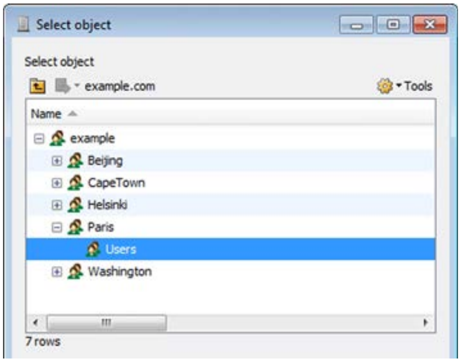

4. Browse to the domain node from which you want to allow users to be linked and click **Select**.
   - The Base DN string is automatically generated based on your selection.
   - The Domain and Base DN are automatically added to the table.
5. Select the Search Scope to define how many levels below the base DN to search for users:
   - **One level**: the search is limited to the specified based DN.
   - **Object level**: The search is limited to entries immediately below the base DN. The base DN is not included.
   - *Sub-tree level*: The search includes the entire subtree below the base DN. The base DN is not included.

6. Click **OK** to close the Authentication Server properties. To activate the changes, right-click the Authentication Server and select **Apply Configuration**.

What's Next?

- Continue by Creating and Linking Authentication Server User Accounts (page 849).

Creating and Linking Authentication Server User Accounts

▼ To create and link Authentication Server user accounts

1. Select **Configuration→Configuration→User Authentication** from the menu. The User Authentication Configuration view opens.
2. Browse to **Users→Domain for Authentication Server**. A list of users opens on the right.
3. Right-click **Authentication Server Users** and select **New User**. The User Properties dialog opens.

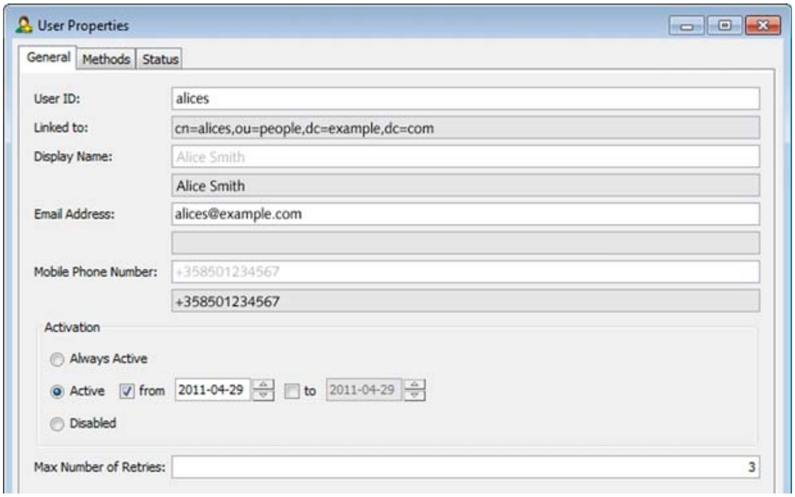

User ID: alices

Linked to: cn=alices,ou=people,dc=example,dc=com

Display Name: Alice Smith

Alice Smith

Email Address: alices@example.com

Mobile Phone Number: +358501234567

+358501234567

Activation

○ Always Active

● Active ☑ from 2011-04-29 ☐ to 2011-04-29

○ Disabled

Max Number of Retries: 3

4. Enter the same **User ID** that is used in the integrated external user database.
   - The external user database is searched, and the user account is automatically linked when a match is found.
   - The Display Name, E-mail Address, and Mobile Phone Number are automatically filled in with the user information retrieved during linking.
   - If no match is found, you can manually create an account.
5. (*Optional for linked users*) Edit the user information to update existing information, or enter missing information as described below. These changes only apply to the user accounts in the Authentication Server's database. The linked user accounts in the external user database server are not changed.

Table 51.6 User Properties for Linked Users

Setting	Description
Display Name	The user's full name. The user account must always contain a display name.
E-mail Address	An e-mail address where the user can receive notification of changes to the user account. The user account must contain an e-mail address if you use SMTP notification.

Table 51.6 User Properties for Linked Users (Continued)

Setting	Description
Mobile Phone Number	A mobile phone number where the user can receive SMS messages for Mobile Text authentication and notification of changes to the user account. The user account must contain a phone number if the Mobile Text Authentication Method is enabled for the user, or if you use SMS notification.

6. (*Optional*) Change the Activation settings for the user account as described below:

Setting	Description
Always Active	The user account is considered active immediately and is never automatically disabled.
Active	The user account is considered active starting on the **from** date and optionally ending on the **to** date. If an end date is specified, the user account is automatically disabled on the end date.
Disabled	The user account is not active.

7. (*Optional*) Specify the **Max Number of Retries** before the user account is locked when the user fails to enter the correct credentials.
   - The maximum number of retries should be low enough to discourage attempts to guess passwords by brute force, but high enough to prevent legitimate users from having their accounts locked if they enter their credentials incorrectly.
8. Switch to the **Methods** tab.
9. Click **Add** to select the **Authentication Methods** for the user.
   - You can add more than one authentication method for each user. This way, you can put the User in more than one User Group when the User Groups have different authentication methods.
   - If you have not configured any Authentication Methods yet, you can create them in this dialog.

10. Define the properties for the selected Authentication Method:

Table 51.7 Authentication Methods

Authentication Method	Setting	Configuration
Password, Mobile Text	Use Password From Directory Server (*Linked users only*)	There are no options to configure. The linked user must have a password defined in the external directory server.
	Use Internal Password	Enter and confirm the user's password or click **Generate** to create a random alphanumeric password.
	Password never expires (*Optional*)	The user is never required to change the password.
	User cannot change Password (*Optional*)	The user cannot change the password. Password changes must be made by the administrator.
	User must change Password on next logon (*Optional*)	The user must enter a new password at the next authentication. Select this option if you want to create a temporary password for the user.
Mobile ID - Synchronized, Mobile ID - Challenge	Use PIN From Directory Server (*Linked users only*)	There are no options to configure. The linked user must have a PIN defined in the external directory server.
	Use Internal PIN	Enter and confirm the user's PIN or click **Generate** to create a random PIN.
	PIN never expires (*Optional*)	The user is never required to change the PIN.
	User cannot change PIN (*Optional*)	The user cannot change the PIN. PIN changes must be made by the administrator.
	User must change PIN on next logon (*Optional*)	The user must enter a new PIN at the next authentication. Select this option if you want to create a temporary PIN for the user.
	Seed (*Optional*)	Click **Generate** to generate a new seed.

11. Click **OK**. The user account is created.

What's Next?

- To configure user authentication, proceed to Getting Started with User Authentication (page 856).
- Otherwise, the configuration is complete.

Adding/Removing Users From User Groups

▼ To add Users to new User Groups

➥ Drag and drop the User element on a User Group in the element tree. The User remains a member of the current User Group and becomes a member of the new group.

▼ To remove Users from User Groups

1. Browse into the User Group from which you want to remove the User.
2. Right-click the User element and select **Tools→Remove**. A confirmation message is shown.
3. Click **Yes** to remove the user from the User Group.

Importing and Exporting User Information

Import and export operations allow you to transfer user and user group information from one Management Server's internal LDAP user database to another Management Server's internal LDAP user database. The internal LDAP user database is represented by the default LDAP Domain **InternalDomain**.

Importing Users from an LDIF File

To be able to import user information successfully, the information must meet two conditions:

- The distinguished name (DN) must always be of type `dn: cn=cn-of-the-user-or-group,dc=stonegate`.
- All user groups must be directly attached below the `dc=stonegate` top-level group.

Note – The import feature is meant for importing users exported from another Management Server's internal LDAP database. The import is not meant to work with other LDIF files.

▼ To import users from an LDIF file

1. Select **File→Import→Import Users**. The Import LDIF dialog opens.
2. Select the correct file and click **Import**. A new tab opens showing the progress of the import.
3. When the import is finished, check the messages for any warnings before you close the tab.
   - If warnings were displayed, select **Configuration→Configuration→User Authentication** to open the User Authentication Configuration view and browse to **Users→InternalDomain** to check the properties for those Users/User Groups.

Exporting Users to an LDIF File

You can export the user information stored in the internal user database on one Management Server into an .ldif file to transfer it to another Management Server.

▼ To export users to an LDIF file

1. Select **Configuration→Configuration→User Authentication**. The User Authentication Configuration view opens.
2. Browse to **Users→InternalDomain**.
3. Right-click the correct user group under InternalDomain and select **Tools→Export StoneGate Users**.The Export to LDIF dialog opens.
4. Select the correct location, type in the filename, and click **Export**. A new tab opens showing the progress of the export.
5. When the export is finished, check the messages for any warnings before you close the tab.

Changing User Passwords

If you use LDAP password authentication, you can reset and change a user's password as necessary through the Management Client.

▼ To change a user's password

1. Select **Configuration→Configuration→User Authentication**. The User Authentication Configuration view opens.
2. Browse to **Users** and then to the correct LDAP Domain and User Group.
3. Right-click the User element and select **Tools→Change Password**. The Change Password dialog opens.
4. Type in the new password in both fields and click **OK**. The new password becomes valid immediately.

Clearing the Authentication Settings of a User

You can quickly remove all the authentication settings of a user without deleting the user from the system. This is useful, for example, when a user's account is known to have been compromised and access rights need to be revoked very quickly.

▼ To clear a user's authentication attributes

1. Select **Configuration→Configuration→User Authentication**. The User Authentication Configuration view opens.
2. Browse to **Users** and then to the correct LDAP Domain and User Group.
3. Right-click the User element and select **Tools→Clear Authentication Attributes**.
4. Click **Yes** in the Confirmation message.

Related Tasks

- ▶ Changing User Passwords (page 853)

Resetting Local User Database on Firewalls

If you use the SMC's internal user database, the user information is stored centrally on the Management Server. When changes are made, they are incrementally replicated to each Firewall node to ensure fault-tolerant authentication. If necessary, you can perform a full synchronization manually as instructed below.

▼ To reset a Firewall engine's local user database

1. Right-click a Single Firewall or individual node(s) of a Firewall Cluster and select **Commands→Reset User Database**.
2. Click **Yes** in the Confirmation message. The engine's local copy of the user database is replaced by a new copy of the internal user database on the Management Server.

Setting User Database Replication to Firewalls and Master Engines On or Off

By default, the SMC's internal user database (for end-user authentication) is replicated automatically to Firewall engines and Master Engines, so that the engines can perform the authentication without a connection to the Management Server. If you have some special need to do so, you may turn off the replication as explained below.

> **Caution – If you want to prevent users from authenticating, follow the instructions in Clearing the Authentication Settings of a User (page 853) instead of turning off the replication. Turning off the replication prevents any new users you add after the operation from authenticating, but may not prevent existing users from doing so.**

▼ To set user database replication on or off for a Firewall

↳ Right-click a Single Firewall, Firewall Cluster, or Master Engine and select/deselect **Options→User DB Replication**.

Chapter 52

Setting up User Authentication

You can implement user authentication to control which resources different end-users can access. Authentication can be used as an additional access requirement in IPv4 Access rules in Firewall policies. User authentication is not supported on IPS engines and Layer 2 Firewalls. Both internal and external user authentication servers can be used.

The following sections are included:

Getting Started with User Authentication

Prerequisites: None

Authentication means requiring users proof of their identity before giving access to a network resource. Authentication is mandatory with client-to-gateway VPNs. You can require authentication for any non-VPN connections as well.

Alternatively, if strong authentication is not required, you can allow specific users to access services in a trusted environment without requiring user authentication. See Enabling Access Control by User (page 838).

What User Authentication Does

With user authentication, you can:

- Maintain separation of internal networks that have different security levels in cases where the confidentiality of the information that is accessed does not need to be strictly enforced. For example, user authentication can provide an additional access control measure for applications that already exchange information securely.
- Allow secure and confidential access from any location to any internal resources for Stonesoft IPsec VPN Client users.
- Authenticate Administrator and Web Portal User logins. See Authenticating Administrators Using RADIUS Methods (page 262).

Limitations

- User authentication is only supported on Firewalls (not on IPS engines or Layer 2 Firewalls).
- Only IPv4 addresses are supported in user authentication. All elements used in user authentication must have an IPv4 address.

What Do I Need to Know Before I Begin?

Authentication requires a user database that stores the user information and an authentication method that inspects credentials and grants or denies access.

Firewalls have internal authentication methods. Alternatively, you can use external RADIUS or TACACS+ compatible authentication servers such as IAS, RSA Authentication Manager (SecurID), Stonesoft SSL VPN appliances, or the separately licensed Authentication Server component. See the *Firewall/VPN Reference Guide* for more information on the different configuration options for user authentication.

Configuration Overview

Illustration 52.1 Elements in the Configuration

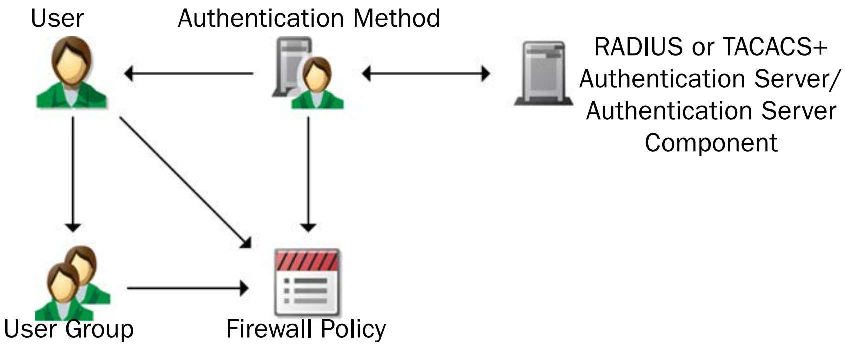
User
Authentication Method
RADIUS or TACACS+ Authentication Server/ Authentication Server Component
User Group
Firewall Policy

1. (*Optional*) Define server elements and authentication methods for external authentication services. See Integrating External Authentication Services (page 858).
2. (*Optional*) Configure the Authentication Server component for user authentication. See Integrating Authentication Server Services (page 861).
3. Add an authentication requirement to the relevant IPv4 Access rules. See Defining IPv4 Access Rules for Authentication (page 873).
4. (*Stonesoft IPsec VPN Client Authentication*) Install the Stonesoft IPsec VPN Client software on the users' computers. See the *Stonesoft IPsec VPN Client Administrator's Guide*.
5. (*Browser-based and Telnet Authentication*) Configure the authentication prompt:
   - Enable users to authenticate using a browser-based authentication prompt. See Enabling Browser-Based User Authentication (page 874).
   - Customize the authentication prompt. See Customizing the HTML Pages Profile for Browser-Based User Authentication (page 880) and Customizing the Telnet Authentication Prompt (page 882).

What's Next?

- To use the Firewall to perform user authentication, or if you have already configured an Active Directory Server element, proceed to Defining IPv4 Access Rules for Authentication (page 873).
- To use an external RADIUS or TACACS+ server for user authentication, proceed to Integrating External Authentication Services (page 858).
- To use the optional Authentication Server component for user authentication, proceed to Integrating Authentication Server Services (page 861).

Related Tasks

- Getting Started with Directory Servers (page 826)

Integrating External Authentication Services

Prerequisites: None

An external authentication server can be any server that supports either the RADIUS or the TACACS+ protocol, including Microsoft Active Directory servers. External authentication servers are integrated with the help of Active Directory Server, RADIUS Authentication Server, TACACS+ Authentication Server, and Authentication Method elements. The RADIUS Authentication Server and TACACS+ Authentication Server elements define the settings necessary for connecting to an external authentication server. The Authentication Method elements define an authentication method, and can include several RADIUS Authentication Servers or TACACS+ Authentication Servers that support the method and can be used as backups to each other.

What's Next?

- ▶ To configure a new Active Directory Server, see Defining Active Directory Server Elements (page 830).
- ▶ To configure a new RADIUS or TACACS+ Authentication Server, start by Defining RADIUS or TACACS+ Authentication Servers.
- ▶ To configure Authentication Methods for an existing RADIUS or TACACS+ Authentication Server, or Active Directory Server proceed to Defining Authentication Methods for External Servers (page 860).

Defining RADIUS or TACACS+ Authentication Servers

You can authenticate end-user access through Firewalls and administrator's logins to the SMC against external authentication servers that support either the RADIUS or the TACACS+ protocol.

▼ To define a RADIUS or TACACS+ Authentication Server element

1. Select **Configuration→Configuration→User Authentication**. The User Authentication Configuration view opens.
2. Right-click **Servers** and select **New→RADIUS Authentication Server** or **New→RADIUS Authentication Server**. The RADIUS Authentication Server Properties or TACACS+ Authentication Server Properties dialog opens.

Illustration 52.2 RADIUS and TACACS+ Authentication Server Properties

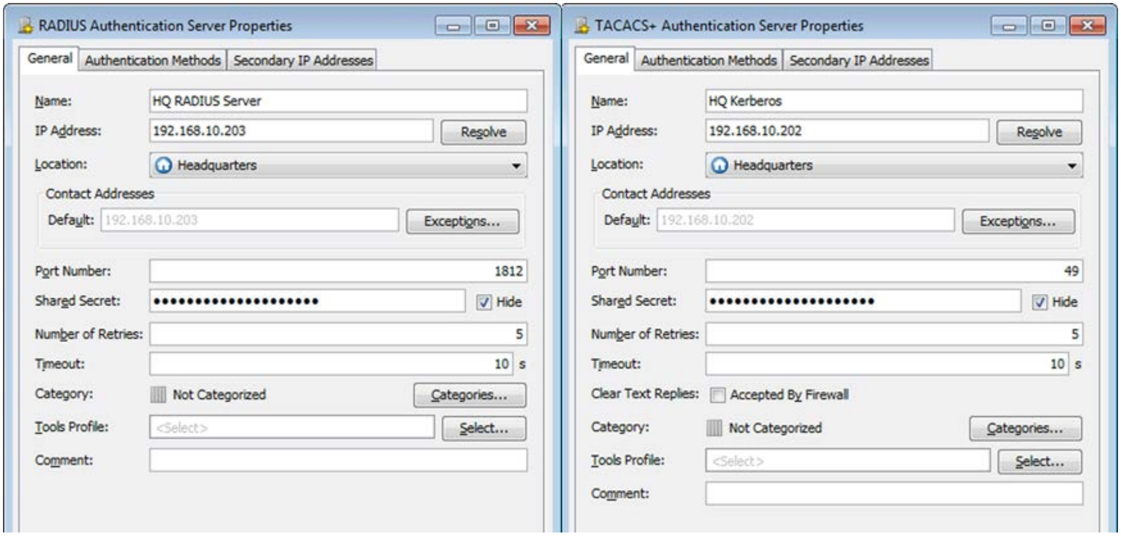

RADIUS Authentication Server Properties

Name: HQ RADIUS Server
IP Address: 192.168.10.203
Location: Headquarters
Contact Addresses
Default: 192.168.10.203
Port Number: 1812
Shared Secret: •••••••••••••••••••• ☑ Hide
Number of Retries: 5
Timeout: 10 s
Category: Not Categorized
Tools Profile: <Select>
Comment:

TACACS+ Authentication Server Properties

Name: HQ Kerberos
IP Address: 192.168.10.202
Location: Headquarters
Contact Addresses
Default: 192.168.10.202
Port Number: 49
Shared Secret: •••••••••••••••••••• ☑ Hide
Number of Retries: 5
Timeout: 10 s
Clear Text Replies: □ Accepted By Firewall
Category: Not Categorized
Tools Profile: <Select>
Comment:

3. Configure the settings as described below:

Setting	Description
Name	Enter a unique name for the element.
IP Address	Enter the IP address of the server. Only IPv4 addresses are supported.
Location/ Contact Addresses	A Location and Contact Address are needed if NAT is applied between a Firewall or Management Server and the external authentication server, so that the Firewall and Management Server cannot connect directly to the IP address defined for the server. For more information, see Defining Locations (page 66) and Defining Contact IP Addresses (page 66).
Port Number (*Optional*)	Enter the port number if the server communicates on a port other than the default port. The predefined Firewall Template allows the engines to connect to the default port. If you change to a custom port, you must add a new IPv4 Access Rule to allow the traffic.
Shared Secret	Enter the secret key for communication with the RADIUS or TACACS+ authentication server. Deselect **Hide** to see the key as you type.
Number of Retries	Enter the number of times Firewalls try to connect to the RADIUS or TACACS+ authentication server if the connection fails.
Timeout	Enter the time (in seconds) that Firewalls wait for the RADIUS or TACACS+ authentication server to reply.
Clear Text Replies (*Optional, TACACS+ only*)	Select **Accepted by Firewall** if you want the Firewall to accept unencrypted replies from the TACACS+ authentication server.

4. (*Optional*) Switch to the **Secondary IP Addresses** tab and define additional IP addresses. Only IPv4 addresses are supported. These are only used in IPv4 Access rules and routing.

Firewalls always use the RADIUS or TACACS+ authentication server's main IP address when they contact the authentication server.

5. Click **OK**.
6. Configure the RADIUS or TACACS+ authentication server to accept connections from your Firewall engines:
   - Take special care to ensure that the shared secret is entered identically in the Management Client and on the RADIUS or TACACS+ authentication server.
   - The identity that the Firewall announces to the server is the IP address of the interface that has been selected as the value for the **Identity for Authentication Requests** option in the properties of the Firewall element's Interface Options.

Note – The IP address used as the identity is a name only. The interface and IP address used for the authentication-related connections is selected based on the Firewall's routing information just like for any other connection.

The connections to RADIUS or TACACS+ authentication servers are allowed in the predefined Firewall Template. Make sure your IPv4 Access and IPv4 NAT rules are configured correctly for these connections.

What's Next?

- ▶ To use the RADIUS or TACACS+ Authentication Server for User Authentication, continue by Defining Authentication Methods for External Servers.
- ▶ If you want to use the server for authenticating administrators' Management Client logins, see Authenticating Administrators Using RADIUS Methods (page 262).

Defining Authentication Methods for External Servers

Authentication Methods are used in directory server properties and in User and User group properties to specify the allowed authentication methods for the users, in the properties of a RADIUS or TACACS+ Authentication Server or Active Directory Server to specify the authentication method offered by the server, and in the IPv4 Access rules to specify which authentication method users are required to use.

The RADIUS and TACACS+ protocols are generic communications protocols for user authentication, so you can use many different types of authentication methods that use the RADIUS and TACACS+ authentication protocols, for example, simple password authentication, one-time passwords, or any other username/passcode-type authentication schemes.

Firewalls can also use the Internet Authentication Service (IAS) in previous Windows Server versions or the Network Policy Server (NPS) in Windows Server 2008 to authenticate end-users. If you use a Windows Server's IAS/NPS service for authentication, see Defining Active Directory Server Elements (page 830) instead of the steps below.

▼ To define a new Authentication Method

1. Select **Configuration→Configuration→User Authentication**. The User Authentication Configuration view opens.
2. Right-click **Authentication Methods** and select **New Authentication Method**.
3. Enter the **Name** and select the authentication method **Type**.
4. Click **Add** and select one or more RADIUS or TACACS+ Authentication Servers.
   - When multiple servers are associated with the same Authentication Method, the servers are used as alternative servers if the first contacted server does not respond.
   - All servers associated with the same Authentication Method must contain identical information on each authenticating user, since it is not possible for the user to determine which of the alternative servers is being contacted.
5. Click **OK**.

What's Next?

▶ Use the Authentication Method in an IPv4 Access rule for user authentication. Proceed to Defining IPv4 Access Rules for Authentication (page 873).

Integrating Authentication Server Services

Prerequisites: The Authentication Server component has been installed

The Authentication Server is a separately licensed component that provides authentication services for other Stonesoft components, and for third-party devices. The Authentication Server component can be installed as a single node or as a cluster. You can install the Authentication Server in the following ways:

- If you want to install a single node Authentication Server on the same computer as the Management Server, you can create and install the Authentication Server using the Installation Wizard as instructed in the *Management Center Installation Guide*.
- If you want to install a single node Authentication Server on a different computer than the Management Server, or if you want to install an Authentication Server cluster, you must first define the Authentication Server element in the Management Client before installing the Authentication Server node(s).

After installation, you must configure the Authentication Methods provided by the Authentication Server, and the notification channels for notifying end-users of changes to their accounts. If you want to allow the SSL VPN or third-party devices to use the Authentication Server's authentication services, you must define those devices as RADIUS clients of the Authentication Server. If you want to use the Authentication Server in CHAPv2 RADIUS authentication, you must create and sign the Authentication Server's certificate.

Note – Whenever you make changes to the Authentication Server properties, or to any server that interacts with the Authentication Server, you must apply the Authentication Server's configuration for the changes to take effect.

What's Next?

- If you are installing a single node Authentication Server on a different computer than the Management Server, or installing an Authentication Server cluster, begin by Defining Authentication Server Elements.
- If you have already installed a new Authentication Server and are configuring it for the first time, begin by Defining Authentication Server Authentication Methods (page 863).
- Otherwise, proceed to the next relevant section:
  - To specify which devices can use the Authentication Server's authentication services, proceed to Defining Authentication Server RADIUS Clients (page 867).
  - To specify how end-users are notified of changes to their accounts, proceed to Defining Authentication Server Notification Channels (page 868).
  - To generate a certificate for CHAPv2 RADIUS authentication, proceed to Creating and Signing Authentication Server Certificates (page 870).

Defining Authentication Server Elements

If you plan to install a single node Authentication Server on a different computer than the Management Server, or plan to install an Authentication Server cluster, you must define the Authentication Server element and install the correct licenses in the Management Client before installing the Authentication Server node(s).

▼ To define an Authentication Server element

1. Select **Configuration→Configuration→User Authentication**. The User Authentication Configuration view opens.
2. Right-click **Servers** and select **New→Authentication Server**. The Authentication Server Properties dialog opens.

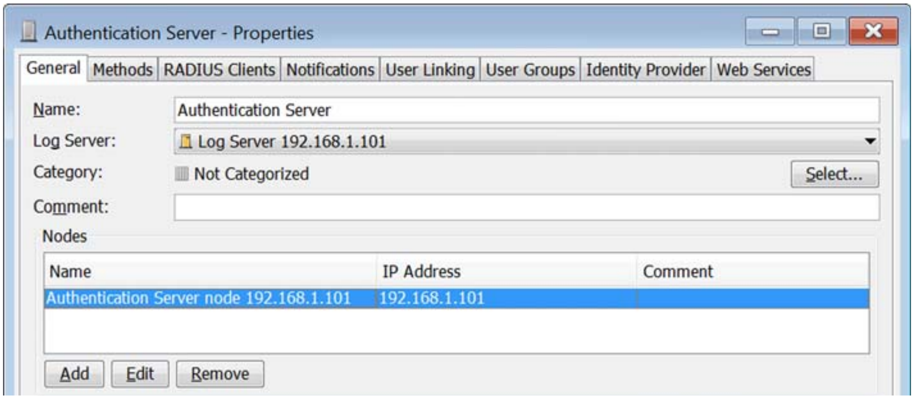

Name: Authentication Server

Log Server: Log Server 192.168.1.101

Category: Not Categorized

Comment:

Nodes

Name	IP Address	Comment
Authentication Server node 192.168.1.101	192.168.1.101	

3. Configure the settings as described below:

Setting	Description
Name	Enter a unique name for the element.
Log Server	Select the Log Server to which the Authentication Server sends status and event information.

4. Configure the Node in one of the following ways:
   - Select the automatically-created node and click **Edit**.
   - (*Cluster only*) Click **Add**.

   The Node Properties dialog opens.
5. Configure the settings as described below:

Setting	Description
Name	You can optionally change the automatically-generated name for the node.
IP Address	Enter the IP address of the node.
Location/ Contact Addresses	A Location and Contact Address are needed if NAT is applied between a Firewall or Management Server and the Authentication Server node, so that the Firewall and Management Server cannot connect directly to the IP address defined for the node. For more information, see Defining Locations (page 66) and Defining Contact IP Addresses (page 66).

6. Click **OK**. Add any additional nodes in the same way.
7. Install the Authentication Server license as instructed in Installing Licenses (page 1050).
8. Install the Authentication Server node(s) as instructed in the *Management Center Installation Guide*.

What's Next?

▶ After all nodes are installed, continue by Defining Authentication Server Authentication Methods.

Defining Authentication Server Authentication Methods

There are four predefined types of Authentication Methods for use with the Authentication Server. You can add one Authentication Method of each type to the Authentication Server. The Authentication Methods for the Authentication Server cannot be used by RADIUS or TACACS+ Authentication Servers.

▼ To define Authentication Server authentication methods

1. Select **Configuration→Configuration→User Authentication**. The User Authentication Configuration view opens.
2. Select **Servers**.
3. Right-click **Authentication Server** and select **Properties**. The Authentication Server Properties dialog opens.
4. Switch to the **Methods** tab.

5. Click **Add** and select the type of Authentication Method to add:

Authentication Method	Description
Password	A static password is created and maintained for authenticating remote access.
Mobile ID - Synchronized	For use with the Stonesoft Mobile ID client. The Mobile ID client must be seeded with a seed generated by the Authentication Server before the first use. During authentication, users enter their user ID and are prompted to enter a one-time password (OTP). Users enter their PIN in the Mobile ID client and the Mobile ID client software generates the OTP.
Mobile ID - Challenge	For use with the Stonesoft Mobile ID client. The Mobile ID client must be seeded with a seed generated by the Authentication Server before the first use. During authentication, users enter their user ID, and are prompted with a challenge to provide the correct response. Users enter their PIN in the Mobile ID client, and the Mobile ID client software generates the response.
Mobile Text	Based on a combination of a PIN and one-time password (OTP) distributed by SMS to the end-user.

6. Enter a unique **Name** and the **Port** on which the Authentication Server listens for incoming RADIUS packets for the Authentication Method. You must allow RADIUS traffic on this port in the Firewall Access rules.

Caution – The security of the system to which authentication provides access requires that these communications remain confidential. We recommend transferring the RADIUS communications to and from the Authentication Server over secure networks only.

7. (*Optional*) To specify and enable default password or PIN options for Authentication Server users, switch to the **Default Settings** tab and select **Enable When a New User is Created**. Configure the settings according to the selected method as described below:

Table 52.1 Password and PIN Settings

Authentication Method	Setting	Description
Password, Mobile Text	Use Password From Directory Server (*Linked users only*)	The user's password is automatically retrieved from the directory server when the user account is linked.
	Use Internal Password	A password must be manually created for each user.
	Password Never Expires (*Optional*)	The user is never required to change the password.
	User Cannot Change Password (*Optional*)	The user cannot change the password. Password changes must be made by the administrator.
	User Must Change Password on Next Logon (*Optional*)	The user must enter a new password at the next authentication. Select this option if you want to create a temporary password for the user.

Table 52.1 Password and PIN Settings (Continued)

Authentication Method	Setting	Description
Mobile ID - Synchronized, Mobile ID - Challenge	Use PIN From Directory Server (*Linked users only*)	The user's PIN is automatically retrieved from the directory server when the user account is linked.
	Use Internal PIN	A PIN must be manually created for each user.
	PIN Never Expires (*Optional*)	The user is never required to change the PIN.
	User Cannot Change PIN (*Optional*)	The user cannot change the PIN. PIN changes must be made by the administrator.
	User Must Change PIN on Next Logon (*Optional*)	The user must enter a new PIN at the next authentication. Select this option if you want to create a temporary PIN for the user.

8. (*Optional*) Switch to the **Password Properties** or **PIN Properties** tab and configure the settings for the authentication method according to your environment:

Authentication Method	Setting	Description
Password, Mobile Text	Minimum (*characters*)	The minimum number of characters that passwords must contain.
	Maximum (*characters*)	The maximum number of characters that passwords may contain.
	Minimum (*letters*)	The minimum number of letters that passwords must contain.
	Minimum (*numbers*)	The minimum number of numerical character that passwords must contain.
	Password Expires in (*days*)	The number of days that passwords are valid.
	Password History Size (*passwords*)	The number of unique new passwords that each user must use before reusing an expired password.
Mobile ID - Synchronized, Mobile ID - Challenge	PIN Expires in	The number of days that PINs are valid.
	PIN History Size	The number of unique new PINs that each user must use before reusing an expired PIN.

9. (*Optional*) Switch to the **Extended Properties** tab and configure the settings according to your environment:

Setting	Configuration
RADIUS Character Encoding	The character set for encoding communication with RADIUS clients. When the system receives a RADIUS packet, it normally transforms the data to strings according to the UTF-8 standard. Some RADIUS clients do not support the UTF-8 standard.
User Name May Not Change During Session	Allow only the user ID associated with linked user accounts in the Authentication Server’s user database to be used during the authenticated session.
Show Users Reason for RADIUS Rejection	Display the reason a request has been rejected to the RADIUS client.
Stonesoft Account Required Before Authentication	Allow only user IDs associated with linked user accounts in the Authentication Server's user database to authenticate.
Allow Users Not Listed in Any Directory Server	Allow users that are not listed in any directory server to authenticate. Usually the sharing of user information between directories is not allowed.
Always Create Users	Automatically create user accounts in the Authentication Server user database whenever a user successfully authenticates.
Create User on Failed Login	Automatically create user accounts in the Authentication Server user database when a user fails to authenticate.

10.(*Optional*) Click **Global Settings** and configure the authentication settings for all users.

Setting	Description
Max. Retries	The maximum number of retries before the user account is locked when the user fails to enter the correct credentials.
Time Lock Timeout	The number of minutes the user account is locked when the user fails to provide the correct credentials and the maximum number of retries has been reached.
Time Lock Interval	The maximum number of unsuccessful login attempts before the user account is locked.
Session Timeout	The number of minutes that the user authentication is valid. The user must re-authenticate after the session timeout has been reached.
Password Expiration Notification	The user is notified about an expiring password the specified number of days before the password expires.

11. Click **OK**. To activate the changes, right-click the Authentication Server and select **Apply Configuration**.

What's Next?

- If you are configuring a new Authentication Server and want to allow other devices to use the Authentication Server's authentication services, continue by Defining Authentication Server RADIUS Clients.
- If you are configuring a new Authentiation Server and do not need to define RADIUS clients, proceed to Defining Authentication Server Notification Channels (page 868).
- Otherwise, you are ready to use the Authentication Server to authenticate users. Continue by Defining IPv4 Access Rules for Authentication.

Defining Authentication Server RADIUS Clients

A RADIUS client connects to the Authentication Server for RADIUS authentication. A RADIUS client can be the Management Server, a Firewall, the SSL VPN, or any external component that uses the Authentication Server for RADIUS authentication. Management Servers and Firewalls with static IP addresses are automatically defined as RADIUS clients. SSL VPN gateways must be manually added as RADIUS clients. The number of RADIUS clients you can add is defined by the Authentication Server license. Other system components do not count against the license limit.

▼ To define Authentication Server RADIUS clients

1. Select **Configuration→Configuration→User Authentication**. The User Authentication Configuration view opens.
2. Open the Authentication Server Properties and switch to the **RADIUS Clients** tab.

Illustration 52.3 Authentication Server Properties - RADIUS Clients Tab

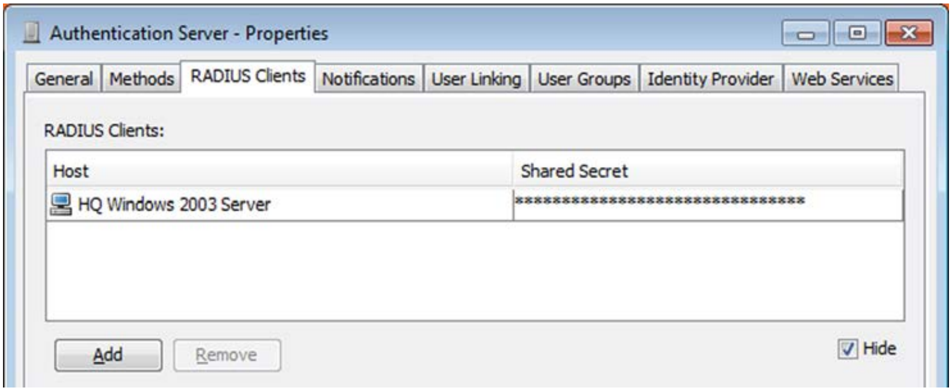

RADIUS Clients:

Host	Shared Secret
HQ Windows 2003 Server	********************************

3. Click **Add**. A row is added to the table.
4. Double-click the **Host** cell and select the Host element that represents the RADIUS client.
5. Double-click the **Shared Secret** cell and enter the shared secret between the RADIUS client and the Authentication Server.

Tip – Deselect **Hide** to see the Shared Secret as plaintext.

6. Repeat Step 4-Step 5 for any additional RADIUS clients.

7. Click **OK**. To activate the changes, right-click the Authentication Server and select **Apply Configuration**.

What's Next?

- If you are configuring a new Authentication Server, continue by Defining Authentication Server Notification Channels.
- Otherwise, you are ready to use the Authentication Server to authenticate users. Continue by Defining IPv4 Access Rules for Authentication (page 873).

Defining Authentication Server Notification Channels

The Authentication Server's notification settings define how end-users can be notified about changes to their user accounts. SMTP notification requires an SMTP server. SMS notification requires an SMS gateway or SMS service provided by a service provider.

▼ To define Authentication Server notification channels

1. Select **Configuration→Configuration→User Authentication**. The User Authentication Configuration view opens.
2. Open the Authentication Server Properties and switch to the **Notifications** tab.

Illustration 52.4 Authentication Server Properties - Notifications Tab

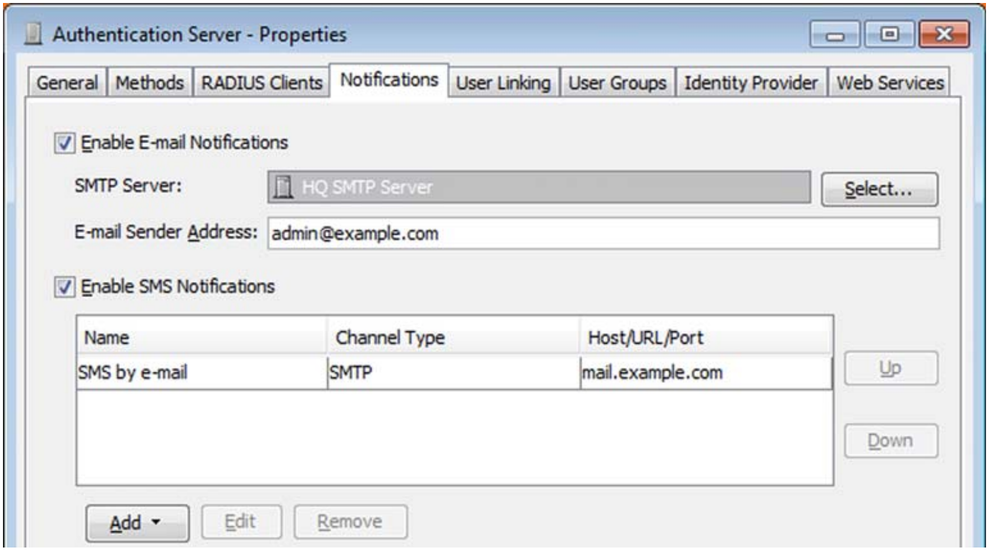

☑ Enable E-mail Notifications

SMTP Server: HQ SMTP Server

E-mail Sender Address: admin@example.com

☑ Enable SMS Notifications

Name	Channel Type	Host/URL/Port
SMS by e-mail	SMTP	mail.example.com

3. Select **Enable E-mail Notifications** to enable the sending of notifications by e-mail, and configure the following settings:

Table 52.2 E-mail Notifications

Setting	Configuration
SMTP Server	Select or create the SMTP Server that is used to send notifications as e-mail.

Table 52.2 E-mail Notifications (Continued)

Setting	Configuration
E-mail Sender Address (*Optional*)	Enter the e-mail address to be used in the From field of the e-mail. If this setting is left blank, the Default Sender Address defined in the SMTP Server Properties is used.

4. Select **Enable SMS Notifications** to enable the sending of notifications by SMS message to a mobile phone. Authentication Server user accounts must contain a mobile phone number if you enable SMS notification.
5. Click **Add** and select the channel type that represents the SMS gateway or SMS service you use:
   - **SMTP**: The SMS gateway converts e-mail into an SMS message to the user.
   - **HTTP Plugin**: The SMS service converts an HTTP request into an SMS message to the user.

Example With the SMTP channel type, sending e-mail to +1234567890@example.com sends the text of the e-mail as an SMS message to +1234567890 (the user's mobile phone number).

6. Configure the properties according to the channel type:

Table 52.3 SMS Notification Channels

Channel Type	Setting	Configuration
SMTP	Name	Enter a unique name.
	SMTP Server	Select the SMTP server that communicates with the SMS gateway or service.
	Account (*If required by SMTP server*)	Enter the user name for connecting to the SMTP server.
	Password (*If required by SMTP server*)	Enter the password for connecting to the SMTP server.
	Local Host Name	Enter the DNS host name of the Authentication Server.
	E-mail Sender Address (*Optional*)	Enter the e-mail address to be used in the From field of the e-mail. If this setting is left blank, the Default Sender Address defined in the SMTP Server Properties is used.
	Recipient	Enter the domain of the SMTP server to which the mail is sent for processing.
	Subject	Enter a subject for the message.

Table 52.3 SMS Notification Channels (Continued)

Channel Type	Setting	Configuration
HTTP Plugin	Name	Enter a unique name.
	URL	Enter the URL of the SMS gateway or service where the request is sent for processing.
	Proxy Host	If the SMS gateway or service must be accessed through a proxy, enter the IP address of the proxy host.
	Proxy Port	If the SMS gateway or service must be accessed through a proxy, enter the port number for contacting the proxy host.
	Account (*If required by proxy host*)	Enter the user name for connecting to the proxy host.
	Password (*If required by proxy host*)	Enter the password for connecting to the proxy host.
	Use HTTP 1.1	Select this option if the proxy host requires the use of the HTTP 1.1 protocol.

7. Click **OK**. To activate the changes, right-click the Authentication Server and select **Apply Configuration**.

What's Next?

▶ If you are configuring a new Authentication Server and want to use the Authentication Server in CHAPv2 RADIUS authentication, continue by Creating and Signing Authentication Server Certificates.

▶ Otherwise, you are ready to use the Authentication Server to authenticate users. Continue by Defining IPv4 Access Rules for Authentication (page 873).

Creating and Signing Authentication Server Certificates

The Authentication Server's certificate is used to authenticate the server to the client in CHAPv2 RADIUS authentication. You can either self-sign the certificate directly in the Authentication Server properties dialog or use an externally signed certificate.

▼ To generate an Authentication Server certificate

1. Select **Configuration→Configuration→User Authentication**. The User Authentication Configuration view opens.
2. Open the Authentication Server Properties and switch to the **Certificate** tab.

Illustration 52.5 Authentication Server Properties - Certificate Tab

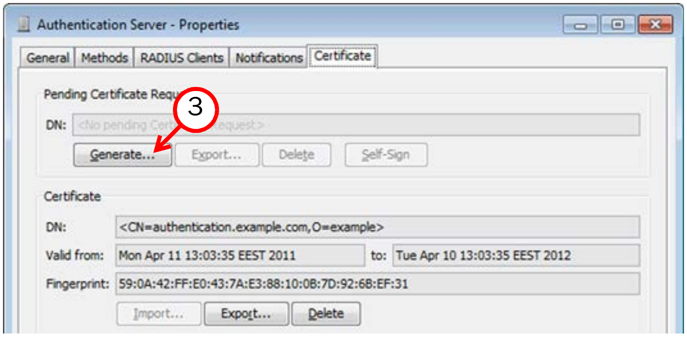

DN: <No pending Cert equest>

DN: <CN=authentication.example.com,O=example>

Valid from: Mon Apr 11 13:03:35 EEST 2011 to: Tue Apr 10 13:03:35 EEST 2012

Fingerprint: 59:0A:42:FF:E0:43:7A:E3:88:10:0B:7D:92:6B:EF:31

3. Click **Generate**. The Server Certificate Request dialog opens.
4. Enter the certificate information:

Certificate Field	Description
Common Name(CN)	The fully qualified domain name (FQDN) of the authentication page as it should appear in the certificate.
Organization(O) *(Optional)*	The name of your organization as it should appear in the certificate.
Organizational Unit(OU) *(Optional)*	The name of your department or division as it should appear in the certificate.
Country/Region(C) *(Optional)*	Standard two-character country code for the country of your organization.
State/Province(ST) *(Optional)*	The name of state or province as it should appear in the certificate.
City/Locality(L) *(Optional)*	The name of the city as it should appear in the certificate.
Key Length	Length of the key for the generated public-private key pair. The default is 1024 bits.

5. Click **OK** to close the Server Certificate Request dialog.
6. Sign the certificate request:
   - To sign the certificate in the Authentication Server properties dialog, click **Self-Sign**.
   - To sign using an external certificate authority, click **Export** to save the certificate request. After signing, **Import** the signed certificate file through this dialog.
7. When the certificate is signed, you are prompted to remove the pending certificate request. Click **OK**
8. Click **OK** to close the Authentication Server properties. To activate the changes, right-click the Authentication Server and select **Apply Configuration**.

What’s Next?

▶ You are ready to use the Authentication Server to authenticate users. Continue by Defining IPv4 Access Rules for Authentication (page 873).

Enabling Federated Authentication With the Authentication Server

In Federated Authentication, user identities and authentication are managed separately from services. This allows entities that provide services to delegate the authentication process and maintenance of user accounts to another entity. Federated Authentication also allows the user to use the same credentials for authentication in multiple security domains, optionally as part of a single-sign-on (SSO) configuration.

Entities in a Federated Authentication scenario have the following roles:

- **Subject**: the user who requests access to an application or service.
- **Identity Provider**: Verifies the credentials of the user and creates a unique, signed SAML assertion that contains the information about the user and the user’s privileges. This assertion is then used to authenticate the user to the Service Provider.
- **Service Provider**: Provides applications or services. When a user requests an application or service, the Service Provider sends an authentication request to the Identity Provider. The Identity Provider authenticates the user and replies with an authentication response to the Service Provider.

The Authentication Server can act as the Identity Provider. All Authentication Methods provided by the Authentication Server can be used to create SAML assertions. For detailed configuration instructions, see Technical Note #6731: Federated Authentication in StoneGate Authentication Server at the Stonesoft website.

Enabling RADIUS Accounting With the Authentication Server

RADIUS Accounting allows you to keep track of a user’s session for billing, statistical, or network monitoring purposes. RADIUS Accounting packets contain information about when a user logs in and out of a system, and what resources the user accesses. The Authentication Server can collect RADIUS accounting information and provide a summary of the data.

For detailed configuration instructions, see Technical Note #6732: Radius Accounting in StoneGate Authentication Server at the Stonesoft website.

Enabling Web Services With the Authentication Server

You can use the Authentication Server component to authenticate users of web services. Any applications that can make requests over HTTPS can be used as Authentication Server web services. For detailed configuration instructions, see Technical Note #7732: How to Configure Web Services in Stonesoft Authentication Server at the Stonesoft website.

Defining IPv4 Access Rules for Authentication

Prerequisites: Getting Started with Directory Servers

In Firewall IPv4 Access rules, the Authentication cell specifies matching criteria for accessing a particular service and for setting options for the authentication. Authentication rules can be used to require authentication to access services, and for authenticating VPN client users. The authentication settings in a rule are configured in the same way regardless of whether a VPN is used.

Caution – Only a VPN ensures confidential information exchange. A rule that only requires authentication does not significantly increase the security of access from external networks to internal networks.

Connections from users who have not successfully authenticated, or whose authentication has expired do not match rules that require authentication. The connection matching continues to rules further down in the policy.

▼ To define an IPv4 Access rule for authentication

1. Open the Firewall Policy for editing.
2. Add an IPv4 Access rule and define the **Source**, **Destination**, and **Service**.
3. Right-click the **Action** cell and select the action:
   - Select **Allow** to create a rule that allows access to a particular service.
   - Select **Use IPsec VPN** if you want to direct connections into a VPN.
4. Double-click the **Authentication** cell. The Authentication Parameters dialog opens.

Illustration 52.6 Authentication Parameters

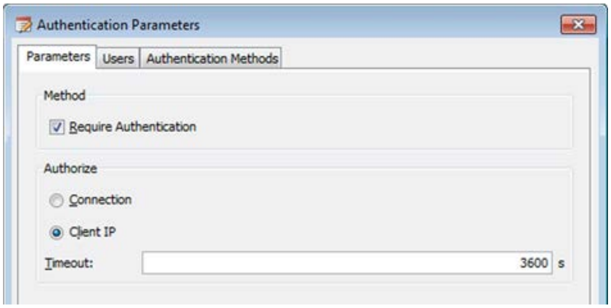

5. Select the Method for the authentication process:
   - **Require Authentication**: The users are required to authenticate themselves. The request is sent to the appropriate external authentication server or Authentication Server component and evaluated.
6. Select one of the authorization types:
   - **Connection** authorization is granted for a single connection. Each connection requires that the users type in their credentials again. This option is useful for granting restricted access to a particular service that does not open many separate connections.
   - **Client IP** authorization permits connections from the client's IP address for as long as indicated in the **Timeout** setting in seconds. The default Timeout is 3600 seconds. This option is better if you want to grant access to several services or a service that opens many connections during a single session.

7. Switch to the **Users** tab and select the Users or User Groups that this rule applies to.
   - Authentication Server users cannot be used directly in rules. Instead, you must use the User element from the external directory server to which the Authentication Server user is linked. See Linking Authentication Server Users to External Directories (page 847).
   - User Groups cannot be used with the Authentication Server.
8. Switch to the **Authentication Methods** tab and select the Authentication Method(s) or click **Set to ANY** to allow any authentication method.
9. Click **OK** to close the Authentication Parameters dialog. The rule is now finished.

What's Next?

▶ This concludes the User Authentication configuration. Install the policy to transfer the changes to the Firewall.

Related Tasks

▶ Enabling Browser-Based User Authentication (page 874)

Enabling Browser-Based User Authentication

Prerequisites: None

Browser-based user authentication allows the end-users to authenticate to a Firewall or Virtual Firewall using any standard web browser. Alternatively, you can use browser-based user authentication with external RADIUS or TACACS+ compatible authentication servers.

▼ To enable browser-based user authentication

1. Open the Firewall or Virtual Firewall properties and switch to the **Add-Ons** tab.
2. Click **User Authentication**. The Browser-Based User Authentication dialog opens.

Illustration 52.7 Browser-Based User Authentication - Properties

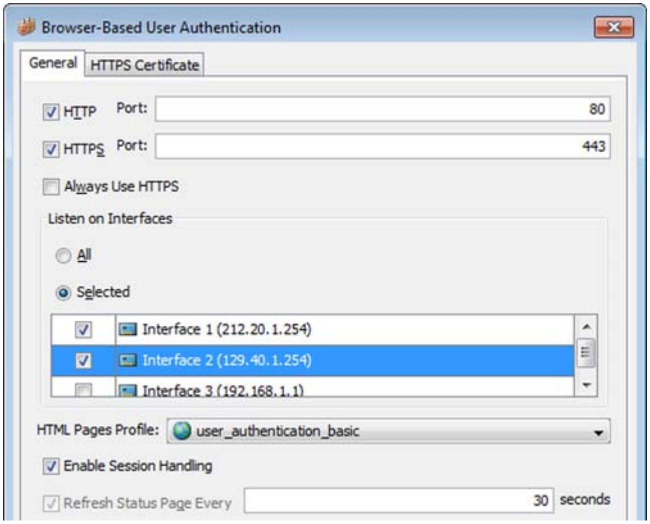

☑ HTTP Port: 80

☑ HTTPS Port: 443

□ Always Use HTTPS

Listen on Interfaces

○ All

● Selected

☑ Interface 1 (212.20.1.254)

☑ Interface 2 (129.40.1.254)

□ Interface 3 (192.168.1.1)

HTML Pages Profile: user_authentication_basic

☑ Enable Session Handling

☑ Refresh Status Page Every 30 seconds

3. Select **HTTPS** to allow authentication using encrypted HTTPS connections and/or **HTTP** to allow authentication using plain HTTP connections.

> ❗ **Caution – Plain HTTP connections are unsecured and transfer the username and password in cleartext. Use encrypted HTTPS connections to avoid loss of sensitive information.**

4. *(Optional)* Modify the port settings if you want to use a different port for the authentication interface. You must use the same port settings when you define an IPv4 Access rule that allows the authentication connections. See Defining IPv4 Access Rules for Browser-Based User Authentication *(page 877)*.

5. *(Recommended)* To prevent unauthorized access to resources, select the interfaces through which users can authenticate to the Firewall or Virtual Firewall in the **Listen on Interfaces** section. Only interfaces that have IPv4 addresses are supported.

6. Select the **HTML Pages Profile** that defines the look of the login page, challenge page, and status page shown to end-users when they authenticate.
   - If the HTML Pages Profile you want to select is not listed, select Other, or select New to create a new HTML Pages Profile as instructed in Customizing the HTML Pages Profile for Browser-Based User Authentication (page 880).

7. (*Optional*) Select **Enable Session Handling** to enable cookie-based strict session handling.
   - If the option is selected, the end-user must keep the status page open. If the status page is closed or cannot be refreshed, the connection is terminated.

8. *(Optional)* Select **Refresh Status Page Every** and define how often the status page is automatically refreshed.
   - The option is automatically selected when you select **Enable Session Handling**.
   - If the option is selected, the end-user must keep the status page open. If the status page is closed or cannot be refreshed, the connection is terminated.

What's Next?

▶ If HTTPS is enabled, proceed to Creating and Signing HTTPS Certificates for Browser-Based User Authentication.

▶ Otherwise, continue by Defining IPv4 Access Rules for Browser-Based User Authentication (page 877).

Creating and Signing HTTPS Certificates for Browser-Based User Authentication

If HTTPS is enabled for browser-based user authentication, you must have a signed HTTPS certificate.

▼ To create and sign a certificate

1. Switch to the **HTTPS Certificate** tab in the Browser-Based User Authentication dialog.
2. Enter the certificate information:

Certificate Field	Description
Common Name(CN)	The fully qualified domain name (FQDN) of the authentication page as it should appear in the certificate.
Organization(O)	*(Optional)* The name of your organization as it should appear in the certificate.
Organizational Unit(OU) *(Optional)*	The name of your department or division as it should appear in the certificate.
Country/Region(C) *(Optional)*	Standard two-character country code for the country of your organization.
State/Province(ST) *(Optional)*	The name of state or province as it should appear in the certificate.
City/Locality(L) *(Optional)*	The name of the city as it should appear in the certificate.
Key Length	Length of the key for the generated public-private key pair. The default is 1024 bits.

3. Select how you want to **Sign** the certificate:
   - Select **Internally with** to sign the certificate using the SMC's Internal CA for Gateways.
   - Select **With External Certificate Authority** if you want to create a certificate request for an external Certificate Authority to sign.
4. Click **Generate Request**. The signed certificate or the certificate request is displayed.
5. *(External certificate authorities only)* Click **Export** when the certificate request is displayed and sign the certificate with an external certificate authority.
6. Click **Import Certificate** to import the signed certificate.
7. Click **OK** to close the Certificate Request dialog.
8. Click **OK** to close the Browser-Based User Authentication dialog.

What's Next?

▶ Continue by Defining IPv4 Access Rules for Browser-Based User Authentication.

Defining IPv4 Access Rules for Browser-Based User Authentication

Browser-based user authentication is not allowed by default in the Firewall Template policy. You must add a rule that allows this traffic in the Firewall Policy. To reduce the risk of resource consumption or DoS (denial of service) attacks, we recommend limiting the number of connections from each source IP address. Under normal conditions, there should only be one connection at a time from each source IP address. However, incomplete connections or other network errors could temporarily result in more than one simultaneous connection attempt from the same IP address. For this reason, we recommend taking your network environment into consideration when setting the limit for simultaneous connections so that the limit does not interfere with legitimate connection attempts.

▼ To add a rule that allows browser-based user authentication

1. Right-click the Firewall Policy and select **Edit**. The policy opens for editing.
2. Add the following IPv4 Access rule:

Table 52.4 IPv4 Access Rule for Browser-Based User Authentication

Source	Destination	Service	Action
ANY	$$Local Cluster (CVI addresses only) or $$Interface ID X. ip (if specific listening interfaces are selected on the General tab in the Browser-Based User Authentication Properties.)	HTTP and/or HTTPS (Port settings must be the same as defined on the General tab in the Browser-Based User Authentication Properties.)	Allow Connection tracking: Default Connection limit by Source: the number of simultaneous connection attempts you want to allow

What's Next?

- ▶ If you want to automatically redirect unauthenticated HTTP connections to the login page, continue by Enabling Redirection of Unauthenticated HTTP Connections (page 878).
- ▶ Otherwise, the configuration is complete. Install the policy to transfer the changes to the Firewall.

Enabling Redirection of Unauthenticated HTTP Connections

To allow users to access the login page, to require authentication for HTTP connections, and to automatically redirect unauthenticated HTTP connections to the login page, you must define the following IPv4 Access rules:

- An Access rule that allows all clients to access the login page.
- An Access rule that allows authenticated users to establish HTTP connections.
- An Access rule that refuses all HTTP traffic.
- An Access rule that redirects unauthenticated HTTP traffic to an Inspection rule.

▼ To enable redirection of unauthenticated HTTP connections

1. Open the Firewall Policy for editing and add the following IPv4 Access rules:

Table 52.5 Example Access Rules for Unauthenticated HTTP Connections

Source	Destination	Service	Action	Authentication
ANY	IP addresses of interfaces through which users can authenticate.	HTTP (Port settings must be the same as defined on the General tab in the Browser-Based User Authentication Properties.)	Allow	
ANY	IP addresses of network services that require authentication.	HTTP	Allow	Users/User Groups who are allowed to access services, and appropriate Authentication Method(s).
ANY	IP addresses of network services that require authentication.	HTTP	Allow Connection tracking: Default Deep Inspection: on	
ANY	ANY	HTTP	Refuse	

2. Define an IPv4 Inspection Exception and specify a User Response that redirects HTTP traffic terminated by the Inspection rules to the URL of the login page. For more information on User Responses, see Getting Started with User Responses (page 778).

Table 52.6 Example Inspection Exception for Unauthenticated HTTP Connections

Situation	Severity	Source	Destination	Protocol	Action
HTTP_Request-GET	ANY	ANY	IP addresses of services that require authentication	HTTP	Terminate Response: redirect to the login page

What's Next?

- To customize the default HTML pages, see Customizing the HTML Pages Profile for Browser-Based User Authentication (page 880).
- Otherwise, the configuration is complete. Install the policy to transfer the changes to the engine.

Authenticating to a Stonesoft Firewall or Virtual Firewall

Prerequisites: Getting Started with Directory Servers, Enabling Browser-Based User Authentication

Users can authenticate using the Stonesoft IPsec VPN Client, a web browser, or a standard Telnet client. If the users are authenticating for VPN access, they must authenticate using the Stonesoft IPsec VPN Client.

Caution – If users authenticate over an unsecured connection, use a one-time password scheme to reduce the risk of unauthorized access.

▼ To authenticate to a Stonesoft Firewall or Virtual Firewall

1. Access the authentication prompt in one of the following ways:
   - Follow the instructions in *IPsec VPN Client User's Guide* to authenticate using the Stonesoft IPsec VPN Client.
   - Enter the IP address and port of the Firewall to open an authentication page in a web browser.
   - Enter the following command:
     `telnet <firewall's IP address or DNS name> 2543`
2. Enter the user access credentials.
   - If you only enter your username without specifying the LDAP domain, the Firewall assumes the default LDAP Domain. If your user account does not belong to the default LDAP Domain, you must add the LDAP Domain to the username with an @-character as a separator.

Example **Type "`fred@mobileusers`" for user "fred" in the LDAP Domain "mobileusers".**

Customizing the HTML Pages Profile for Browser-Based User Authentication

Prerequisites: None

You can change the look of the HTML pages shown to end-users who authenticate using browser-based user authentication. You can customize the login page, challenge page, and status page that is shown to end-users when they authenticate using a web browser. The pages you define are the same for all user groups. The default HTML pages are stored as a `.zip` file and they are included in dynamic update packages.

Exporting the Default HTML Pages Profile

You can export a default HTML Pages Profile and modify the HTML pages using any HTML editor.

▼ To export the default HTML pages profile

1. Select **Configuration→Configuration→Security Engine** from the menu. The Security Engine Configuration view opens.
2. Browse to **Other Elements→Engine Properties→HTML Pages Profiles**. The default HTML Pages Profiles are displayed in the right panel.
3. Right-click the appropriate default HTML Pages Profile and select **Properties**. The Properties dialog opens.

Illustration 52.8 Browser-Based User Authentication - Default HTML Pages Profile

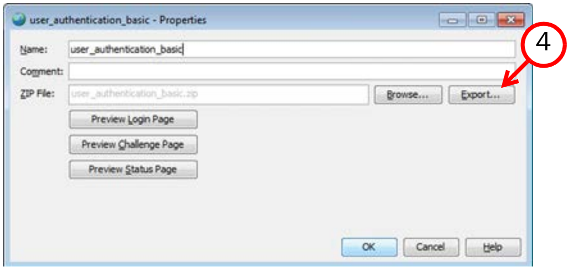

user_authentication_basic - Properties

Name: user_authentication_basic

Comment:

ZIP File: user_authentication_basic.zip

4. Click **Export** and save the `.zip` file under a different name.
5. Click **OK** to close the dialog.

Customizing the Default HTML Pages

You can use an HTML editor to modify the HTML pages. For example, you can add images and CSS files to the default pages.

▼ To customize the default HTML pages

1. Decompress the exported.`zip` file to your computer.
2. Open the HTML file you want to modify in an HTML editor.

Caution – The following fields and buttons must not be removed from the default HTML files.

HTML File	Required Fields
Login Page	The **Username** and **Password** fields, and the **Login** button.
Challenge Page	The **Challenge Response** field and the **Login** button.
Status Page	The **Logged On**, **Connection Status**, **Authentication Message**, and **Message** fields.

3. Compress the modified files into a `.zip` file.

Importing the Custom HTML Pages

To use the customized HTML pages, you must create a new HTML Pages Profile and import the custom files. The HTML files are validated during the import to check that the HTML pages conform to basic HTML standards.

▼ To import the custom HTML pages

1. Select **Configuration→Configuration→Security Engine**. The Security Engine Configuration view opens.
2. Browse to **Other Elements→Engine Properties**.
3. Right-click **HTML Pages Profiles** and select **New HTML Pages Profile**. The HTML Pages Profile Properties dialog opens.

Illustration 52.9 HTML Pages Profile Properties

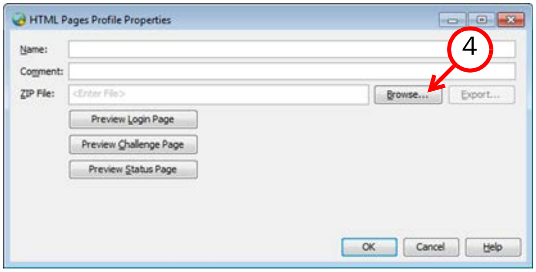

4. Click **Browse** in the HTML Pages Profile Properties dialog and select the `.zip` file that you created. The files are validated during import. The **Name** is detected automatically based on the name of the `.zip` file.

Tip – You can preview the HTML pages by clicking **View Login Page, View Challenge Page,** or **View Status Page.**

5. Click **OK**.

Customizing the Telnet Authentication Prompt

Prerequisites: None

Anyone who is allowed to connect to the Firewall or Virtual Firewall with a Telnet client on port 2543 (TCP) is greeted with an authentication message. If your Firewall Policy is not based on the predefined Firewall Template, make sure that access to this port is allowed from all desired source addresses. Only IPv4 addresses are supported as source addresses.

Note – The predefined Firewall Template allows access from any source to port 2543 on the Firewall. The IPv4 Access rules that require authentication must always also have specific Source and Destination criteria that limit access. See Editing Access Rules (page 672).

By default, the authentication dialog states “Stonesoft firewall <IP address>”. The IP address shown is the IP address of the interface that is selected as the **Identity for Authentication Requests**. You can replace the default message by creating a text file named `auth-message` in the `/data/config/base` directory on the Firewall node as explained below.

▼ To customize the authentication message

1. Connect to the engine command line using SSH or a serial cable.
   - You can enable and disable SSH access and change the password for the Root user through each node’s right-click menu under **Commands**.
2. Enter the following command to create a new text file for the message and launch a text editor: **`vi /data/config/base/auth-message`**
3. Type **a** to switch to editing mode.
4. Type in the message you want to show to your users.
   - Only the first line of the text file (excluding the newline) is used as the message.
   - There must be some text on the first line of the file for it to be used.
   - The message may contain the format string “**`%IP`**”, which is replaced with the IP address (as in the default message) when shown to users.
5. When you are done, press the `Esc` key on your keyboard and type **`ZZ`** to save the changes and exit the text editor.
6. Reboot the engine node to begin using the new message (for instructions, see Rebooting Nodes (page 219)).
7. (*Cluster only*) Copy the `auth-message` file to the `/data/config/base` directory on all nodes in the cluster (for example, using a USB memory stick or SFTP).

Related Tasks

- ▶ Authenticating to a Stonesoft Firewall or Virtual Firewall (page 879)

Monitoring and Testing User Authentication

Prerequisites: See Configuration Overview

Successful and failed user authentication attempts, as well as the system’s own connections to external authentication servers can be followed in the logs. You can also create reports based on this data. There is a separate view for checking the currently active authenticated users. See Monitoring Connections, Blacklists, VPN SAs, Users, and Routing (page 111) for more information.

If there are problems with integration with external components, you can activate more detailed diagnostics logging for authentication. See Enabling/Disabling Firewall/VPN Diagnostics (page 222).

Virtual Private Networks

In this section:

CHAPTER 53

BASIC POLICY-BASED VPN CONFIGURATIONS

The basic policy-based VPN configurations outline specific examples for common policy-based VPN configuration scenarios. For complete configuration instructions for policy-based VPNs and the Route-Based VPN, see Configuring IPsec VPNs (page 917).

The following sections are included:

Getting Started With Basic Policy-Based VPN Configuration

Prerequisites: None

The basic policy-based VPN configurations provide instructions for creating some common policy-based VPN scenarios. You may use these examples when setting up your own policy-based VPNs and then extend the configuration with other features as necessary once you have the basic scenario configured and working. Alternatively, you can follow the complete workflow for creating any type of IPsec VPN in Configuring IPsec VPNs (page 917).

Note – VPNs are not supported on Layer 2 Firewalls.

The following basic configurations are explained in this section:

- Configuration 1 is for creating a policy-based VPN between two or more Stonesoft Firewall/VPN engines that are managed through the same Management Server. A default set of VPN settings are used to simplify the configuration. See Configuration 1: Basic VPN Between Stonesoft Firewall/VPN Engines (page 889).
- Configuration 2 is for creating a policy-based VPN between a Stonesoft Firewall/VPN engine and an IPsec-compatible VPN gateway that is not managed through the same Management Server. A customized set of VPN settings is created, as this is typically mandatory. A pre-shared key is used for authentication. See Configuration 2: Basic VPN With a Partner Gateway (page 893).
- Configuration 3 is for creating a policy-based VPN between a Stonesoft Firewall/VPN engine and Stonesoft IPsec VPN Clients installed on individual computers. A default set of VPN settings are used to simplify the configuration. See Configuration 3: Basic VPN for Remote Clients (page 902).
- Configuration 4 is for creating a policy-based VPN in which several remote gateway connect to a hub gateway, which then forwards connections to the other remote gateways as necessary. A default set of VPN settings are used to simplify the configuration. See Configuration 4: Basic VPN Hub (page 910).

Configuration 1: Basic VPN Between Stonesoft Firewall/VPN Engines

Prerequisites: None

This basic configuration scenario walks you through creating a policy-based VPN between two or more Stonesoft Firewall/VPN engines managed through the same Management Center. This example VPN requires all the firewalls to have a fixed IP address (not DHCP- or PPPoE-assigned).

The address spaces protected by the different Gateways must not overlap within any single VPN. If you use the same IP addresses at the different locations, you must apply NAT to the communications and define the Sites using the translated IP addresses (the addresses that are actually used inside the VPN tunnels).

This scenario uses the default **VPN-A Suite** VPN Profile that contains the VPN settings specified for the VPN-A cryptographic suite in RFC 4308. The profile uses pre-shared keys for authentication.

What's Next?

▶ This scenario has three parts. Start the configuration in Creating Gateway Elements for Configuration 1.

Creating Gateway Elements for Configuration 1

This basic configuration scenario does not explain all settings related to Gateway elements. For more information, see Defining Security Gateways (page 923).

Following these instructions, your VPN will use all interfaces that contain the default "Any Network" element (the default route) as endpoints for the VPN.

▼ To create Gateway elements for configuration 1

1. Select **Configuration→Configuration→VPN**. The VPN Configuration view opens.
2. Right-click **Gateways** in the element tree and select **New→Internal Security Gateway**. The Internal Security Gateway Properties dialog opens.

Illustration 53.1 Internal Security Gateway Properties - General

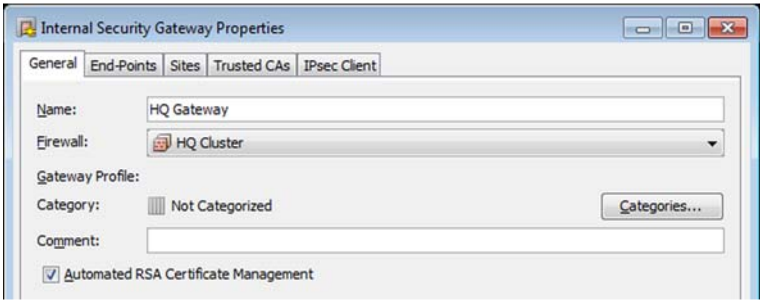

Name: HQ Gateway
Firewall: HQ Cluster
Gateway Profile:
Category: Not Categorized
Comment:
☑ Automated RSA Certificate Management

3. Select the **Firewall** engine the Gateway element represents in VPNs.
4. (*Optional*) Change the **Name** that was added based on the Firewall element's name.

Illustration 53.2 Internal Security Gateway Properties - Sites

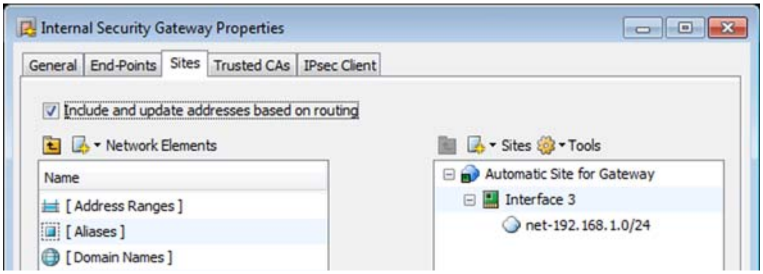

☑ Include and update addresses based on routing

[Address Ranges]
[Aliases]
[Domain Names]

Automatic Site for Gateway
Interface 3
net-192.168.1.0/24

5. If you plan to use NAT to translate the IP addresses of the hosts that make connections through this VPN, switch to the **Sites** tab and drag and drop the network(s) for the NAT address space on top of the (top-level) automatic Site element on the right.
   - The Sites must include only internal networks. Interfaces with the Any Network element are therefore not included by default and you must not add them in this type of VPN.
6. Click **OK** to close the Internal Security Gateway Properties.

To create an element for another Firewall/VPN engine that takes part in the VPN, repeat the steps from Step 2 (you must have at least two gateways configured).

What's Next?

▶ When you have at least two Gateways configured, proceed to Creating a VPN Element for Configuration 1.

Creating a VPN Element for Configuration 1

This basic configuration scenario does not explain all settings related to VPN elements. For more information, see Defining Policy-Based VPNs (page 949).

▼ To create a VPN element for configuration 1

1. Right-click **VPNs** in the element tree and select **New VPN**. The VPN Properties dialog opens.

Illustration 53.3 VPN Profile Properties

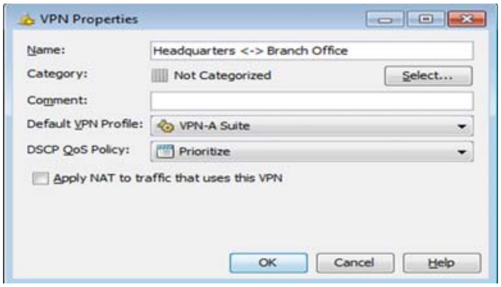

Name: Headquarters <-> Branch Office
Category: Not Categorized
Comment:
Default VPN Profile: VPN-A Suite
DSCP QoS Policy: Prioritize
☐ Apply NAT to traffic that uses this VPN

2. Give the new element a unique **Name**.
3. Select **VPN-A Suite** from the **Default VPN Profile** list.
   - The VPN Profile defines most of the IPsec settings. For instructions on creating your own VPN Profiles, see Defining VPN Profiles (page 939).

4. If you want to apply NAT rules to the communications that go through the VPN, select the **Apply NAT to traffic that uses this VPN** option. This does not affect the communications that the two internal gateways have with each other to set up and maintain the VPN; those are always matched to the NAT rules.

5. Click **OK**. The VPN editing view opens on the **Overall Topology** tab.

Illustration 53.4 VPN Editing View - Overall Topology

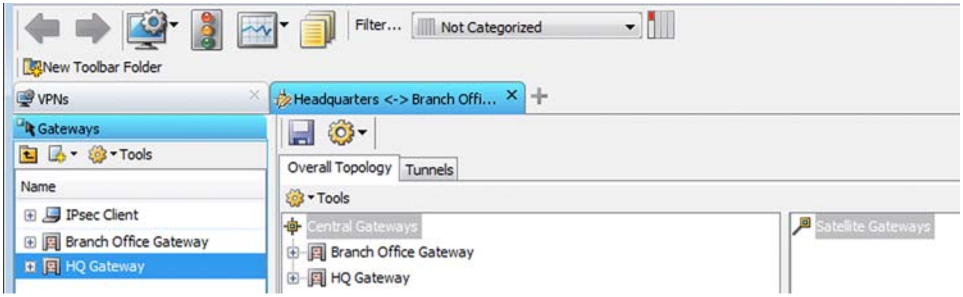

6. Drag and drop the Gateways you want to include in this VPN into either of the two panels to define which Gateways can create a VPN with each other.
   - If you add a Gateway under **Central Gateways**, the Gateway can establish a VPN with any other Gateway in this VPN (both Central and Satellite).
   - If you add a Gateway under **Satellite Gateways**, the Gateway can establish a VPN only with Gateways defined as Central in this VPN.
   - Add two or more Gateways in this view to create a VPN. You must add at least one of the Gateways under **Central Gateways**. You do not have to add any Gateways under **Satellite Gateways** (all Gateways can be Central).

Note – Be careful that you do not accidentally drop Gateways on top of other Gateways. This creates a hub topology where the top-level Gateway forwards connections from other components to the lower-level Gateway.

7. Switch to the **Tunnels** tab.

8. Check that the **Validity** column in the **Gateway<->Gateway** and the **End-Point<->End-Point** tables has a green check mark to indicate that there are no problems.
   - If the **Validity** column of a tunnel has a warning icon, see the **Issues** panel to check what the problem is (if the panel is not displayed, you can open it through the **View** menu or the right-click menu for the Validity cell of a rule). If Issues are displayed, correct them as indicated. Long issues are easiest to read by hovering over the issue text so that the text is shown as a tooltip.

9. Save the VPN.

Creating Rules for VPN Configuration 1

The rules in this example allow protected hosts to open connections both ways. Two rules are created here to allow the different directions of traffic separately. VPN rules are matched based on Source, Destination, and Service like any other rules.

This basic configuration scenario does not explain all settings related to VPN Access rules. For more information, see Creating Rules for Policy-Based VPNs (page 958).

▼ To create Access rules for configuration 1

1. Open the policy of one of the Firewall/VPN engines involved in the VPN for editing.
2. Add two IPv4 Access rules in a suitable location in the policy.
   - Make sure that the rules for sending traffic through the VPN are above any other rules that match the same traffic with Allow, Discard, or Refuse as their action.
   - Traffic that you do not want to send through the VPN must not match these rules. Traffic that is not routable through the VPN is dropped if it matches these rules.
3. Fill in the rules as outlined below. If NAT is enabled in the VPN, keep in mind that the Access rules are checked before the translations defined in the NAT rules are applied to the source and/or destination IP addresses in the packets.

Table 53.1 Example VPN Rules

Source	Destination	Service	Action
Local internal networks.	Remote internal networks.	Set as needed.	Select **Use VPN**, then switch **Action** to **Enforce** and click **Select** to add the VPN element you created.
Remote internal networks.	Local internal networks.	Set as needed.	Select **Use VPN**, then switch **Action** to **Enforce** and click **Select** to add the VPN element you created.

4. Save the policy.
5. Add the same rules in the policies of all firewalls involved in the VPN.

Caution – If you continue to use this VPN, change the pre-shared key periodically (for example, monthly) to ensure continued confidentiality of your data. See Renewing or Generating Pre-Shared Keys (page 990). Alternatively, you can switch to certificate-based authentication by creating a custom VPN profile.

Refresh the policies of all firewalls involved in the VPN to activate the new configuration. The VPN is established when some traffic matches the Access rules.

What's Next?

▶ If NAT is required for this VPN, add the NAT rules. See Creating NAT Rules for Policy-Based VPN Traffic (page 963).

Related Tasks

▶ Monitoring VPNs (page 969)

▶ Troubleshooting IPsec VPNs (page 1137)

Configuration 2: Basic VPN With a Partner Gateway

Prerequisites: None

This basic configuration scenario walks you through creating a policy-based VPN between one Stonesoft Firewall/VPN engine under your management and one external VPN gateway that is not managed through the same Management Center. To be able to configure this example VPN, your local firewall must have a fixed IP address (not DHCP- or PPPoE-assigned).

The address spaces protected by the different Gateways must not overlap within any single VPN. If you use the same IP addresses at the different locations, you must apply NAT to the communications and define the Sites using the translated IP addresses (the addresses that are actually used inside the VPN tunnels).

You can create VPNs with IPsec compliant gateway devices from many different manufacturers. This allows you to create VPNs with partner organizations that use a third-party VPN solution. The authentication and encryption options to use must be decided beforehand in co-operation with the administrator of the other gateway. See the VPN Configuration chapter in the *Firewall Reference Guide* for listings of the supported options for Stonesoft Firewall/VPN.

This scenario creates a new VPN Profile that contains the VPN settings that you match with settings defined on the external VPN gateway.

What's Next?

▶ This scenario has six parts. Start the configuration in Creating an Internal Gateway Element for Configuration 2.

Creating an Internal Gateway Element for Configuration 2

If you have already configured a VPN for the firewall engine, reuse the same element for this VPN. There is no need to change any of the settings.

This basic configuration scenario does not explain all settings related to Gateway elements. For more information, see Defining Security Gateways (page 923).

▼ To create an internal Gateway element for configuration 2

1. Select **Configuration→Configuration→VPN**. The VPN Configuration view opens.
2. Right-click **Gateways** in the element tree and select **New→Internal Security Gateway**. The Internal Security Gateway Properties dialog opens.

Illustration 53.5 Internal Security Gateway Properties - General

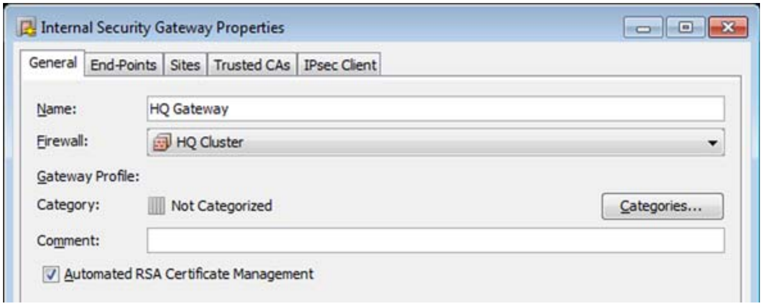

Name: HQ Gateway

Firewall: HQ Cluster

Gateway Profile:

Category: Not Categorized

Comment:

☑ Automated RSA Certificate Management

3. Give the new element a unique **Name**.

4. Use the **Firewall** list to select which Firewall/VPN engine the Gateway element represents in VPNs.

5. If the firewall has a dynamic IP address on an interface that faces the Internet, switch to the End-Points tab and double-click the interface in the table. The Internal End-Point Properties dialog opens:
   - In the **Phase 1 ID** section at the bottom of the dialog, type in an e-mail address in the **ID value** field and click **OK**. This can be any address that is not used as an ID in any of your other end-points.
   - The address entered here is used only as an identification, not for sending e-mail.

6. Switch to the **Sites** tab. A Site element is displayed with each internal network interface on the engine as the content.
   - The Sites represent the addresses that are routable through the VPN. This does not grant any host access directly; your Access rules define the allowed connections.
   - Leave the **Include and update addresses based on routing** option selected. This option automatically updates this information based on routing changes. You can exclude some interfaces while keeping the others automatically updated as explained in the next step.

Illustration 53.6 Internal Security Gateway Properties - Sites

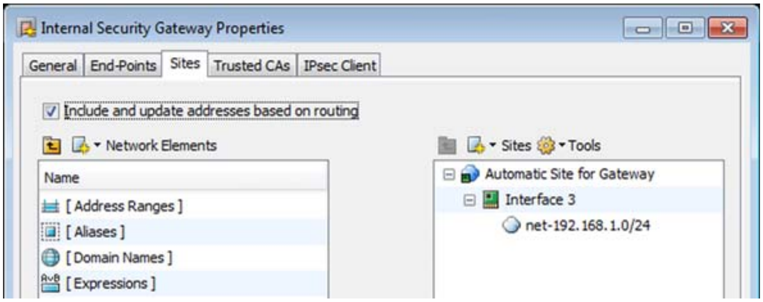

Internal Security Gateway Properties

☑ Include and update addresses based on routing

7. (*Optional*) Select the internal networks that you want to exclude from the VPN by disabling the interface they are under in the automatic site. Disabled interfaces are grayed-out.
   - If you want to include some individual network that is under an otherwise disabled interface, drag and drop it from under the disabled interface onto the Site element to copy the element to the higher level. The copied definition is not updated automatically.
   - The Sites must include only internal networks. Interfaces with the Any Network element are therefore not included by default and you must not add them in this type of VPN.

8. If you plan to use NAT to translate the IP addresses of the hosts that make connections through this VPN, drag and drop the network(s) for the NAT address space on top of the (top-level) automatic Site element on the right. Click **OK** to close the Gateway properties.

Creating an External Gateway Element for Configuration 2

This basic configuration scenario does not explain all settings related to Gateway elements. For more information, see Defining Security Gateways (page 923).

▼ To create an external Gateway element for Configuration 2

1. Right-click **Gateways** in the element tree and select **New→External Security Gateway**. The External Security Gateway Properties dialog opens.

Illustration 53.7 External Security Gateway Properties - General

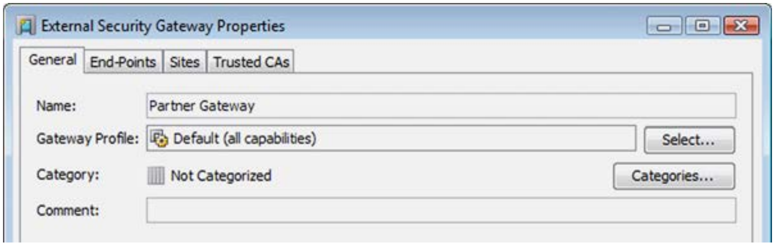

Name: Partner Gateway

Gateway Profile: Default (all capabilities)

Category: Not Categorized

Comment:

2. Give the new element a unique **Name**.
3. Click **Select** for **Gateway Profile** and choose the profile **Default (All Capabilities)** for third-party gateways, or the appropriate version-specific profile for Stonesoft Firewall/VPN engines managed by a different Management Server.
4. Switch to the **End-Points** tab.
5. Click the **Add** button below the table. The External End-Point Properties dialog opens.
6. Define the IP address for the end-point:
   - If the end-point has a static (manually defined) IP address, type in the **IPv4 Address**.
   - If the end-point has a dynamic (DHCP-assigned) IP address, select **Dynamic**.
7. If the external gateway has a dynamic IP address:
   - In the **Phase 1 ID** section at the bottom of the dialog, change the **ID Type** to **Email**.
   - Type in an e-mail address in the **ID value** field. This can be any address that is not used as an ID in any of your other end-points. The address entered here is used only as an identification, not for actually sending e-mail.
   - Click **OK** and repeat if you have other Internet-facing interfaces with a dynamic IP address.

Note – Make sure that the ID Type and ID Value match the identity configured on the external gateway device. If the device has a static IP address, make sure that the device is configured to use it as its identity in this VPN (the most common default configuration) or change the External Security Gateway element configuration as outlined above.

Leave the properties dialog open and continue to the next section.

Defining a Site for External Gateway in Configuration 2

▼ To define Site properties for external gateway in configuration 2

1. Switch to the **Sites** tab in the external Gateway's properties.
2. Double-click the **New Site** element on the right. The Site Properties dialog opens.

Illustration 53.8 Site Properties

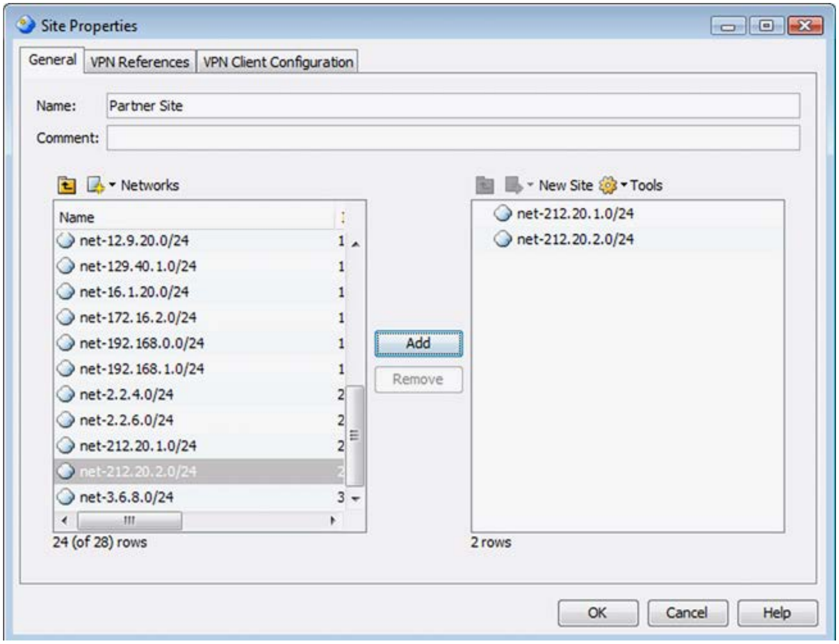

Name: Partner Site

Comment:

Name	
net-12.9.20.0/24	1
net-129.40.1.0/24	1
net-16.1.20.0/24	1
net-172.16.2.0/24	1
net-192.168.0.0/24	1
net-192.168.1.0/24	1
net-2.2.4.0/24	2
net-2.2.6.0/24	2
net-212.20.1.0/24	2
net-212.20.2.0/24	2
net-3.6.8.0/24	3

24 (of 28) rows

- net-212.20.1.0/24
- net-212.20.2.0/24

2 rows

3. Give the Site a descriptive **Name**.
4. Select/create the elements that represent the protected IP addresses behind the Gateway in the left panel and click the **Add** button to include them.
   - For a connection to or from the VPN tunnel to work, the internal IP address used as the source or destination address must be included in the Site of the Gateway. Other traffic cannot use the VPN.
5. Click **OK** in both open dialogs.

Creating a VPN Profile for Configuration 2

This basic configuration scenario does not explain all settings related to VPN Profiles. For more information, see Defining VPN Profiles (page 939).

▼ To create a VPN Profile for configuration 2

1. Expand the **Other Elements** branch in the element tree.
2. Right-click **Profiles** in the element tree and select **New→VPN Profile**. The VPN Profile Properties dialog opens.
3. Give the new element a unique **Name**.

▼ To define IKE SA settings for configuration 2

1. Switch to the **IKE SA** tab.

Illustration 53.9 VPN Profile Properties - IKE SA Tab

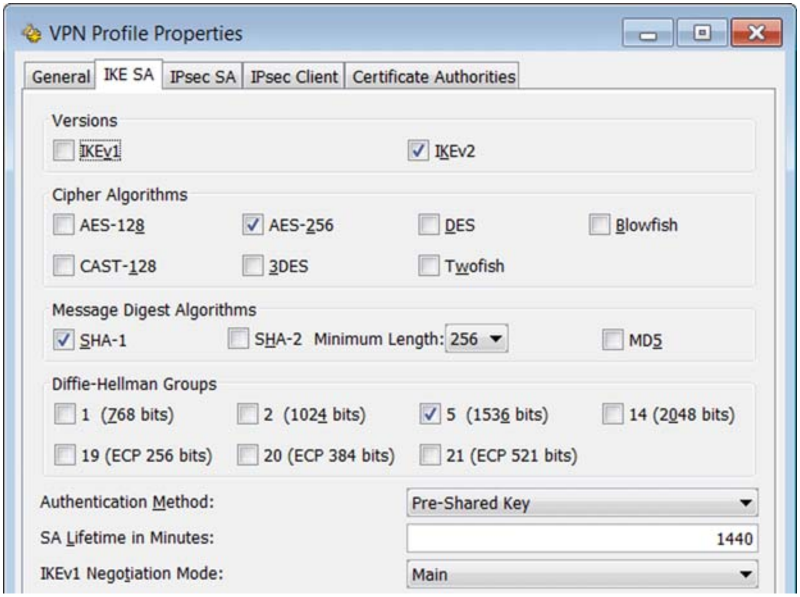

Versions

□ IKEv1 ☑ IKEv2

Cipher Algorithms

□ AES-128 ☑ AES-256 □ DES □ Blowfish

□ CAST-128 □ 3DES □ Twofish

Message Digest Algorithms

☑ SHA-1 □ SHA-2 Minimum Length: 256 □ MD5

Diffie-Hellman Groups

□ 1 (768 bits) □ 2 (1024 bits) ☑ 5 (1536 bits) □ 14 (2048 bits)

□ 19 (ECP 256 bits) □ 20 (ECP 384 bits) □ 21 (ECP 521 bits)

Authentication Method: Pre-Shared Key

SA Lifetime in Minutes: 1440

IKEv1 Negotiation Mode: Main

2. Select the **Version.**
   - You can select either **IKEv1** or **IKEv2** or both. If both versions are selected, IKEv2 is tried first in the negotiations, and IKEv1 is only used if the remote gateway does not support IKEv2.
3. Select the **Cipher Algorithm** (encryption method) that matches the settings of the external gateway device. We recommend that you limit the selection to as few choices as possible.[1]
   - Do not select **DES** unless you are required to do so. DES is no longer considered secure, since it is relatively easy to break DES encryption with modern computers.
   - **3DES** (Triple-DES) has a relatively high overhead compared to other protocols with a comparable level of security and is therefore not a good choice when high throughput is required.
4. Select the **Message Digest Algorithm** (integrity checking method) that matches the settings of the external gateway device.
5. Select the **Diffie-Hellman Group** (used for key exchange) that matches the settings of the external gateway device.
   - We recommend that you do not select group 1 if the external Gateway can use a higher level of protection. Group 1 is not considered sufficiently secure in all cases.
6. Select the **Authentication Method**.
7. If IKEv1 is selected as the Version, adjust the **SA Lifetime in Minutes** to match the settings of the external gateway device.
   - In IKEv2, lifetime is set locally, so it does not have to match the lifetime settings of the external gateway.
8. If one of the Gateways has a dynamic IP address, change the **IKEv1 Negotiation Mode** to **Aggressive**.

[1] The Russian product version has no strong encryption algorithms.

▼ To define IPsec SA settings for configuration 2

1. Switch to the **IPsec SA** tab.

Illustration 53.10 VPN Profile Properties - IPsec SA Tab

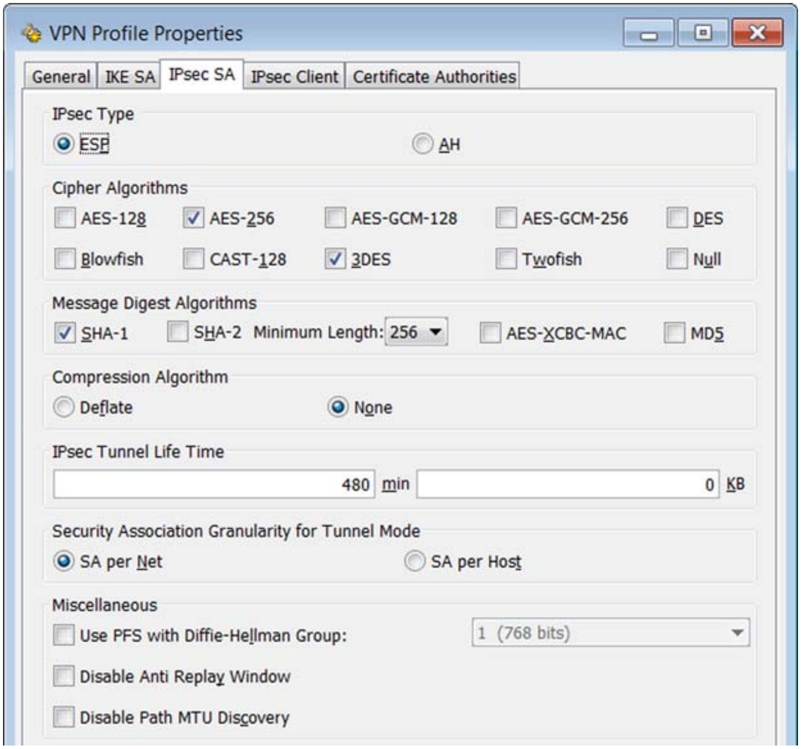

VPN Profile Properties

IPsec Type

● ESP ○ AH

Cipher Algorithms

□ AES-128 ☑ AES-256 □ AES-GCM-128 □ AES-GCM-256 □ DES

□ Blowfish □ CAST-128 ☑ 3DES □ Twofish □ Null

Message Digest Algorithms

☑ SHA-1 □ SHA-2 Minimum Length: 256 □ AES-XCBC-MAC □ MD5

Compression Algorithm

○ Deflate ● None

IPsec Tunnel Life Time

480 min 0 KB

Security Association Granularity for Tunnel Mode

● SA per Net ○ SA per Host

Miscellaneous

□ Use PFS with Diffie-Hellman Group: 1 (768 bits)

□ Disable Anti Replay Window

□ Disable Path MTU Discovery

2. Select the **IPsec Type**:
   - The recommended setting is **ESP** (the communications are encrypted).
   - In most cases, **AH** is not a valid option. The AH setting disables encryption for the VPN, fully exposing all traffic that uses the VPN to anyone who intercepts it in transit. You can use AH to authenticate and check the integrity of communications without encrypting them.
3. Select the **Cipher Algorithm** (encryption method) that matches the settings of the external gateway device.[2]
   - Do not select **Null**. This option disables encryption and allows anyone to view the data in transit.
   - Do not select **DES** option unless you are required to do so. DES is no longer considered secure, as it is relatively easy to break DES encryption with modern computers.
   - **3DES** (Triple-DES) has a relatively high overhead compared to other protocols with a comparable level of security. It is not a good choice when high throughput is required.
   - **AES-GCM-128** or **AES-GCM-256** are recommended for high speed networks.
4. Select the **Message Digest Algorithm** (integrity checking method) that matches the settings of the external gateway device.
5. Make sure **Compression Algorithm** is set to **None**. The external gateway must not use compression.
6. Adjust the **IPsec Tunnel Lifetime** to match the settings of the external gateway device.

[2] The Russian product version has no strong encryption algorithms.

7. Select the **Security Association Granularity for Tunnel Mode** that matches the settings of the external gateway device.
8. (*Recommended*) Select **Use PFS with Diffie-Hellman Group** if the external gateway device can use PFS, and select the Diffie-Hellman Group to use with PFS.
   - We recommend that you do not select group 1 if the external Gateway can use a higher level of protection. Group 1 is not considered sufficiently secure in all cases.
   - PFS is recommended especially when pre-shared keys are used, as in this example.
9. Click **OK**.

Creating a VPN Element for Configuration 2

This basic configuration scenario does not explain all settings related to VPN elements. For more information, see Defining Policy-Based VPNs (page 949).

▼ To create a VPN element for configuration 2

1. Right-click **VPNs** in the element tree and select **New VPN**. The VPN Profile selection dialog opens.
   - If you have not created VPNs before, the VPN Profile selection dialog may open first. If that happens, select your newly created profile to continue.

Illustration 53.11 VPN Properties

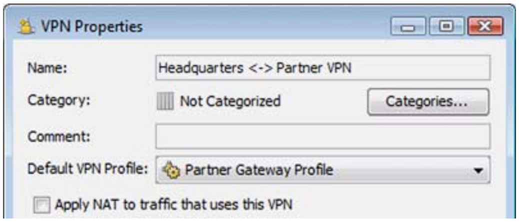
VPN Properties

Name: Headquarters <-> Partner VPN

Category: Not Categorized

Comment:

Default VPN Profile: Partner Gateway Profile

☐ Apply NAT to traffic that uses this VPN

2. Give the new element a unique **Name**.
3. Select the correct profile from the **VPN Profile** list.
4. If you want to apply NAT to the traffic going through the tunnels, select **Apply NAT Rules to Traffic That Uses This VPN**.
5. Click **OK**. The VPN editing view opens on the **Overall Topology** tab.
6. Drag and drop the Internal Gateway element under **Central Gateways**.
7. Drag an drop the External Gateway element under **Satellite Gateways**.

Illustration 53.12 Example Topology for Configuration 2

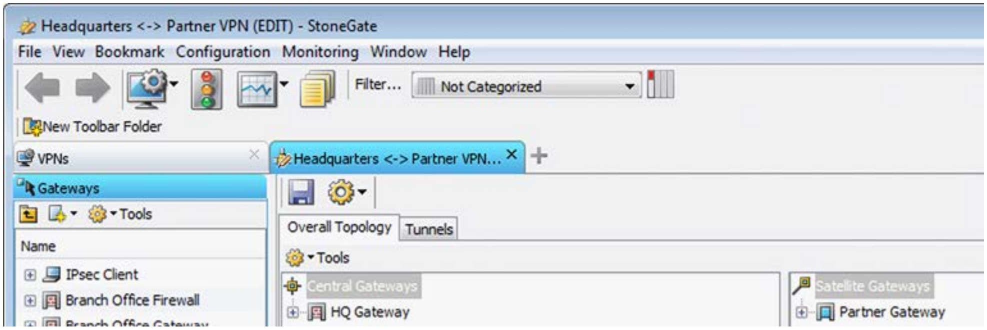

8. Switch to the **Tunnels** tab.

Illustration 53.13 Tunnel for Internal to External Gateway Connection

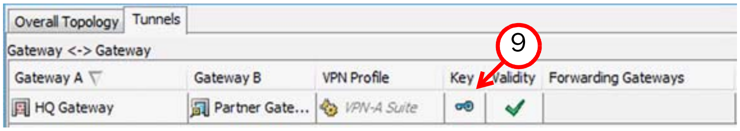

Gateway A	Gateway B	VPN Profile	Key	Validity	Forwarding Gateways
HQ Gateway	Partner Gate...	VPN-A Suite	⚷	✓	

9. Double-click the **Key** cell for the tunnel displayed in the **Gateway<->Gateway** panel. The Preshared Key dialog opens.

10. To match the pre-shared key between the two gateways:
    - To use the key that is automatically generated on the Management Server, click **Export** and transfer the key in the resulting text file securely to the external gateway.
    - To use a different key, replace the displayed key with the one that you have agreed on with the administrator of the external gateway device.

Caution – The pre-shared key must be long and random to provide a secure VPN. Change the pre-shared key periodically (for example, monthly).

11. Click **OK** to close the Pre-Shared Key dialog.

12. Check that the **Validity** column in the **Gateway<->Gateway** and the **End-Point<->End-Point** tables has a green check mark to indicate that there are no problems.
    - If the **Validity** column of a tunnel has a warning icon, see the **Issues** panel to check the problem (if the panel is not displayed, you can open it through the **View** menu or the right-click menu for the Validity cell of a rule). If Issues are displayed, correct them as indicated. Long issues are easiest to read by hovering over the issue text so that the text is shown as a tooltip.

13. Save the VPN.

Creating Rules for Configuration 2

The rules in this example allow protected hosts to open connections both ways. Two rules are created here to allow the different directions of traffic separately. VPN rules are matched based on Source, Destination, and Service like any other rules.

This basic configuration scenario does not explain all settings related to VPN Access rules. For more information, see Creating Rules for Policy-Based VPNs (page 958).

▼ To create Access rules for example 2

1. Open the policy of the Firewall/VPN engine involved in the VPN for editing.
2. Add two IPv4 Access rules in a suitable location in the policy.
    - Make sure that the rules for sending traffic through the VPN are above any other rules that match the same traffic with Allow, Discard, or Refuse as their action.
    - Traffic that you do not want to send through the VPN must not match these rules. Traffic that is not routable through the VPN is dropped if it matches these rules.

3. Fill in the rules as outlined below. If NAT is enabled in the VPN, remember that the Access rules are checked before the translations defined in the NAT rules are applied to the source and/or destination IP addresses in the packets.

Table 53.2 Example VPN Rules

Source	Destination	Service	Action
Local internal networks.	Remote internal networks.	Set as needed.	Select **Use VPN** and change the **Action** to **Enforce**, then click **Select** to add the VPN element you created.
Remote internal networks.	Local internal networks.	Set as needed.	Select **Use VPN** and change the **Action** to **Enforce**, then click **Select** to add the VPN element you created.

4. Save the policy.

Caution – If you continue to use this VPN, change the pre-shared key periodically (for example, monthly) to ensure continued confidentiality of your data. See Renewing or Generating Pre-Shared Keys (page 990). Alternatively, you can switch to certificate-based authentication by creating a custom VPN profile.

Refresh the policies of all firewalls involved in the VPN to activate the new configuration. The VPN is established when some traffic matches the Access rules created here.

What's Next?

- If NAT is required for this VPN, add the NAT rules as well. See Creating NAT Rules for Policy-Based VPN Traffic (page 963).

Related Tasks

- Monitoring VPNs (page 969)
- Troubleshooting IPsec VPNs (page 1137)

Configuration 3: Basic VPN for Remote Clients

Prerequisites: None

This basic configuration scenario walks you through creating a policy-based VPN between a Stonesoft Firewall/VPN engine and any number of Stonesoft IPsec VPN Clients. To be able to configure any client-to-gateway VPN, the firewall must have a fixed IP address (not DHCP- or PPPoE-assigned).

Depending on the desired configuration, you can add VPN client access to an existing gateway-to-gateway VPN as well, but in this example scenario, a separate VPN is created for VPN clients since this approach works in all cases.

This scenario assumes that automatic Site management is used in the VPN without any need for modifications in this VPN (or in other VPNs where the same Gateway is used).

In this scenario, the VPN settings are defined in a copy of the default **VPN-A Suite** VPN Profile that contains the VPN settings specified for the VPN-A cryptographic suite in RFC 4308.

What's Next?

▶ This scenario has six parts. Start the configuration in Managing VPN Client Addresses in Configuration 3.

Managing VPN Client Addresses in Configuration 3

VPN clients cannot use their local IP address in the internal corporate network. In this scenario, NAT is used to solve this problem. This address management method allows connection opening from the VPN client end only. This method is simpler to set up for testing, as it does not require an external DHCP server. However, this method has some restrictions:

- It does not allow connections to be opened from hosts in the internal network to VPN clients.
- It prevents the Stonesoft IPsec VPN Clients from using internal DNS servers.

You may want to change the IP address allocation method to Virtual IP after you have tested the basic VPN connectivity with the configuration explained here (as explained in Configuring Virtual IP Addressing for VPN Clients (page 1002)).

This basic configuration scenario does not explain all settings related to VPN client address management. For more information, see Getting Started With VPN Client Settings (page 996).

▼ To add address translation settings for VPN clients for configuration 3

1. Open the properties of the Firewall element for the Firewall/VPN gateway to which the VPN clients connect.
2. Switch to the **Advanced** tab.
3. Click **VPN Settings** in the **Traffic Handling** section. The VPN Settings dialog opens.

Illustration 53.14 VPN Settings in Firewall Properties

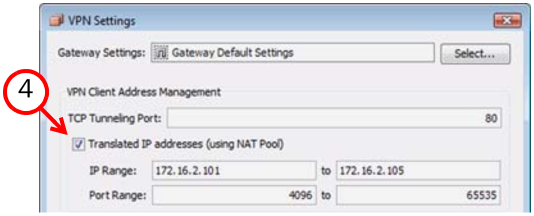

Gateway Settings: Gateway Default Settings

VPN Client Address Management

TCP Tunneling Port: 80

☑ Translated IP addresses (using NAT Pool)

IP Range: 172.16.2.101 to 172.16.2.105

Port Range: 4096 to 65535

4. Select **Translated IP address (using NAT Pool)** and enter the **IP Range** of addresses and the **Port Range** you wish to use for translating VPN client traffic.
   - The VPN clients use these IP addresses when they connect to services in your internal network. Make sure these IP addresses are not used elsewhere in your network.
   - The translation is dynamic, so the number of IP addresses you enter does not need to correspond to the number of clients connecting. Typically, each connection a VPN client user opens to a resource reserves one port from whichever IP address has unreserved ports within the configured range.
   - The reverse NAT for the reply packets is done automatically.
5. Click **OK** to both open dialogs.

Creating Gateway Elements for Configuration 3

This basic configuration scenario does not explain all settings related to Gateway elements. For more information, see Defining Security Gateways (page 923).

The same Gateway element can be used in several VPNs. If you have previously created a Gateway element for the Firewall/VPN engine to which the VPN clients connect, use the existing element instead of creating a new one. Open the Properties dialog for the existing gateway and check that the configuration of the existing Gateway is correct starting from Step 5 below. Changes in the Gateway element affect all VPNs where the same Gateway element is used.

▼ To create a Gateway element for configuration 3

1. Select **Configuration→Configuration→VPN** to switch to the VPN Configuration view.
2. Right-click **Gateways** in the element tree and select **New→Internal Security Gateway**. The Internal Security Gateway Properties dialog opens.

Illustration 53.15 Internal Security Gateway Properties - General

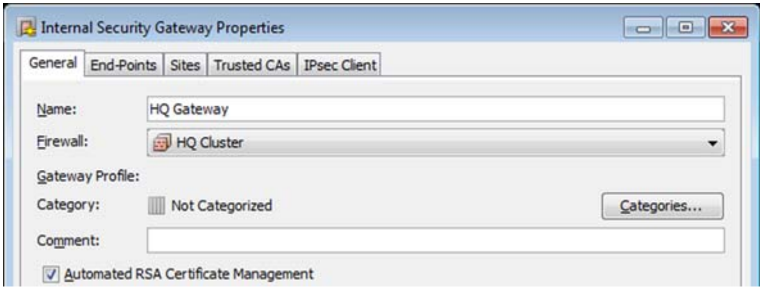

Name: HQ Gateway

Firewall: HQ Cluster

Gateway Profile:

Category: Not Categorized

Comment:

☑ Automated RSA Certificate Management

3. Give the new element a unique **Name**.
4. Select which **Firewall** engine the Gateway element represents in VPNs.

5. Make sure **Automatic Certificate Management** is selected. The gateway must have a certificate for a client-to-gateway VPN.

6. Switch to the **Sites** tab. A Site element is displayed with each internal network interface on the engine as the content.
   - The Sites represent the internal addresses that VPN clients can reach through the VPN. This definition alone does not grant access to any hosts; your Access rules define the allowed connections.
   - Leave **Include and update addresses based on routing** selected. This option automatically updates this information based on routing changes. You can exclude some interfaces while keeping the others automatically updated.
   - If you need to make changes in the Sites (add or remove destination addresses that VPN clients route through the VPN), see Defining Sites for VPN Gateways (page 933).

Illustration 53.16 Internal Security Gateway Properties - Sites

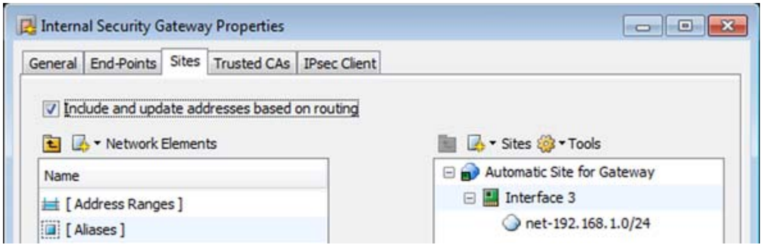

☑ Include and update addresses based on routing

Automatic Site for Gateway
Interface 3
net-192.168.1.0/24

7. Click **OK**.

Adding VPN Client Settings for Configuration 3

This basic configuration scenario does not explain all settings related to authenticating VPN client users. For more information, see Defining VPN Client Settings (page 946).

▼ To add the VPN Client Settings for Configuration 3

1. Expand the **Other Elements→Profiles** branch in the element tree and select **VPN Profiles**. The defined VPN Profiles are displayed.
2. Right-click **VPN-A Suite** and select **New→Duplicate**. The settings from the default profile are copied into the VPN Profile Properties dialog that opens.
3. Give this new VPN Profile a unique **Name**.
4. Switch to the **IKE SA** tab.[3]

Illustration 53.17 VPN-A Profile Properties - IKE SA Tab

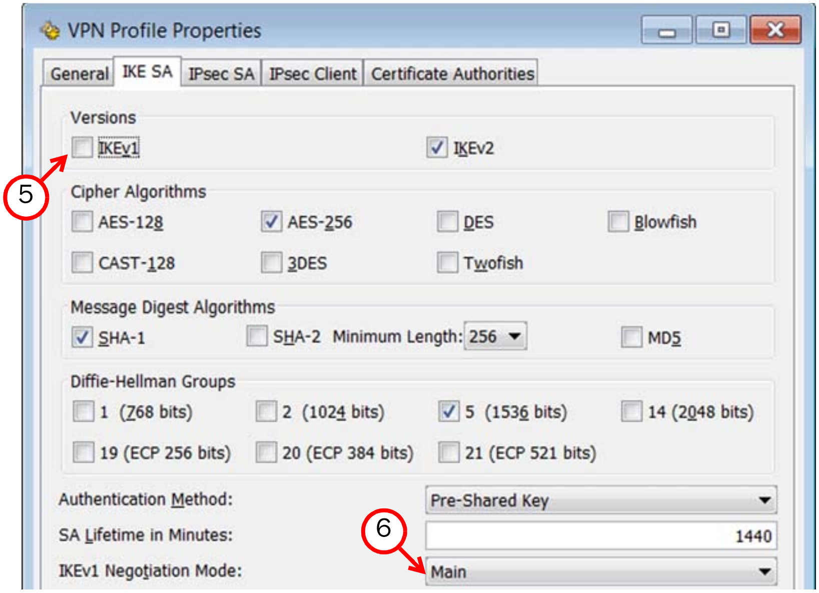

Versions

☐ IKEv1 ☑ IKEv2

Cipher Algorithms

☐ AES-128 ☑ AES-256 ☐ DES ☐ Blowfish

☐ CAST-128 ☐ 3DES ☐ Twofish

Message Digest Algorithms

☑ SHA-1 ☐ SHA-2 Minimum Length: 256 ☐ MD5

Diffie-Hellman Groups

☐ 1 (768 bits) ☐ 2 (1024 bits) ☑ 5 (1536 bits) ☐ 14 (2048 bits)

☐ 19 (ECP 256 bits) ☐ 20 (ECP 384 bits) ☐ 21 (ECP 521 bits)

Authentication Method: Pre-Shared Key

SA Lifetime in Minutes: 1440

IKEv1 Negotiation Mode: Main

5. Select the **Version.**
   - You can select **IKEv1** or **IKEv2** or both. If both versions are selected, IKEv2 is tried first in the negotiations, and IKEv1 is only used if the remote gateway does not support IKEv2.
6. (*Only if IKEv1 is selected as Version*) Make sure **IKEv1 Negotiation Mode** is set to **Main**. This helps ensure that the usernames and passwords of the VPN client users remain confidential.
7. Switch to the **IPsec Client** tab.

[3] The Russian product version has no strong encryption algorithms.

Illustration 53.18 VPN-A Profile Properties - IPsec Client Tab

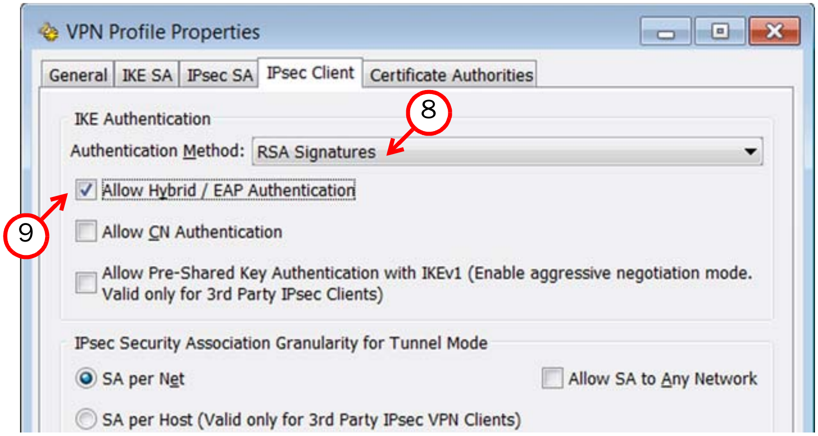

IKE Authentication

Authentication Method: RSA Signatures

☑ Allow Hybrid / EAP Authentication

☐ Allow CN Authentication

☐ Allow Pre-Shared Key Authentication with IKEv1 (Enable aggressive negotiation mode. Valid only for 3rd Party IPsec Clients)

IPsec Security Association Granularity for Tunnel Mode

● SA per Net ☐ Allow SA to Any Network

○ SA per Host (Valid only for 3rd Party IPsec VPN Clients)

8. Make sure the **Authentication Method** is set to **RSA Signatures**.
9. Select **Allow Hybrid/EAP Authentication**.
   - Hybrid authentication is used with IKEv1.
   - EAP (Extensible Authentication Protocol) is used with IKEv2.
10. Make sure **IPsec Security Association Granularity for Tunnel Mode** is set to **SA Per Net**.
11. Click **OK**.

Creating a VPN Element for Configuration 3

This basic configuration scenario does not explain all settings related to VPN elements. For more information, see Defining Policy-Based VPNs (page 949).

▼ To create a VPN element for configuration 3

1. Right-click **VPNs** in the element tree and select **New VPN**.
   - If you have not created VPNs before, the VPN Profile selection dialog may open first. If that happens, select your newly created VPN Profile to continue.

Illustration 53.19 VPN Properties

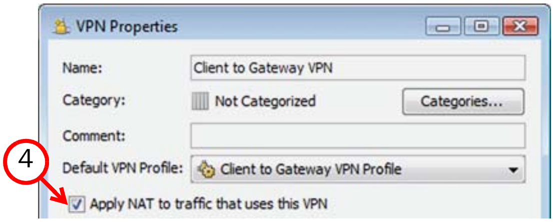

Name: Client to Gateway VPN

Category: Not Categorized

Comment:

Default VPN Profile: Client to Gateway VPN Profile

☑ Apply NAT to traffic that uses this VPN

2. Give the new element a unique **Name**.
3. Select your newly created VPN Profile from the **VPN Profile** list.
4. Select **Apply NAT Rules to Traffic That Uses This VPN**. This applies the NAT rules in the policy and, more importantly, the global NAT definition for the gateway (which you configured in the first part of this scenario).
5. Click **OK**. The VPN editing view opens on the **Overall Topology** tab.

6. Drag and drop the element for the internal gateway to **Satellite Gateways**.
   - Configuring the gateway as a satellite prevents gateway-to-gateway tunnels from being generated if you add several gateways that offer VPN client access, as this example VPN is meant for client-to-gateway access.
7. Drag and drop the default **IPsec Client** Gateway element to **Central Gateways**.

Illustration 53.20 Example Topology for Configuration 3

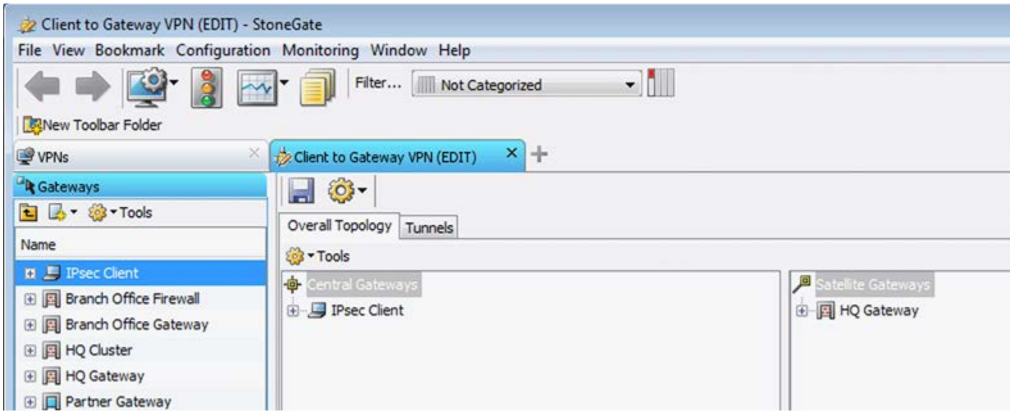

8. Switch to the **Tunnels** tab.
9. Check that the **Validity** column in the **Gateway<->Gateway** and the **End-Point<->End-Point** tables has a green check mark to indicate that there are no problems.
   - If the **Validity** column of a tunnel has a warning icon, check the **Issues** panel to see what the problem is (if the panel is not displayed, you can open it through the **View** menu or the right-click menu for the Validity cell of a rule). If Issues are displayed, correct them as indicated. Long issues are easiest to read by hovering over the issue text so that the text is shown as a tooltip.
10. Save the VPN.

Creating Users for VPN Configuration 3

This basic configuration scenario does not explain all settings related to user authentication. For more information, see Getting Started with User Authentication (page 856).

User authentication is configured in the same way for VPN client connections and normal, unencrypted connections. The same User elements (user accounts) can be used for both. The passwords entered in the VPN client are encrypted so that they remain confidential as they are transferred over the Internet.

▼ To create User elements for VPN client users for configuration 3

1. Select **Configuration→Configuration→User Authentication**. The User Authentication Configuration view opens.
2. Expand the **Users→InternalDomain→stonegate** branch in the element tree.
3. Right-click **stonegate** and select **New→User**. The User Properties dialog opens.

Illustration 53.21 User Properties - General

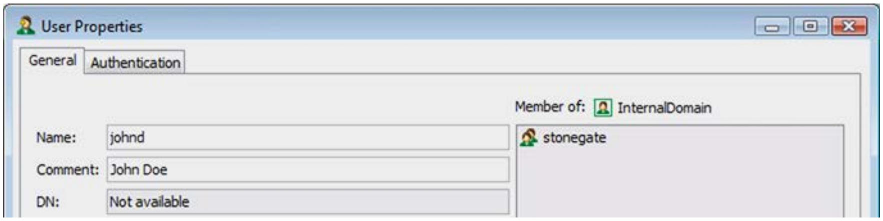

Name: johnd

Comment: John Doe

DN: Not available

Member of: InternalDomain

stonegate

4. Enter the user **Name** the end-user uses to authenticate to the VPN.
5. (*Optional*) Type the real name of the user in the **Comment** field.
6. Switch to the **Authentication** tab.

Illustration 53.22 User Properties - Authentication

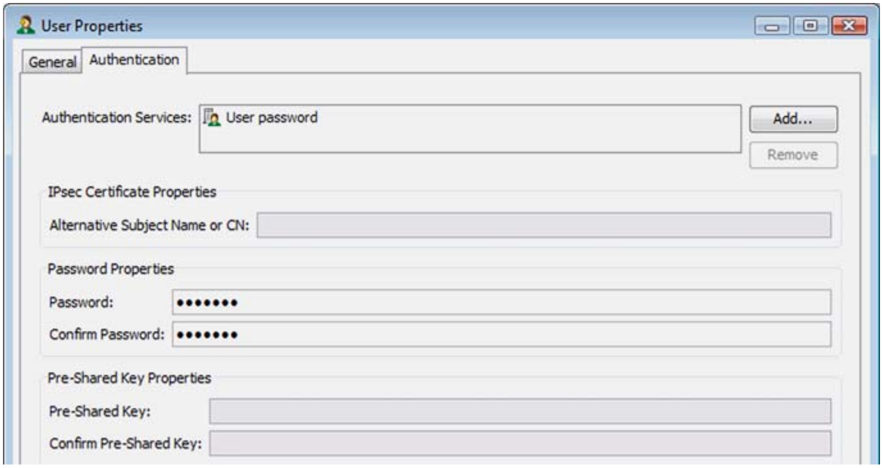

Authentication Services: User password

IPsec Certificate Properties

Alternative Subject Name or CN:

Password Properties

Password: •••••••

Confirm Password: •••••••

Pre-Shared Key Properties

Pre-Shared Key:

Confirm Pre-Shared Key:

7. Click **Add** in the **Authentication Service** section. The Select Element dialog opens.
8. Select **User Password** and click **Select**. This default element allows user password authentication for the internal LDAP database.

9. Enter the same password in the **Password** and **Confirm Password** fields in the User properties. Make a note of the password so that you can communicate it to the user.
   - The pre-shared key method is meant for third-party VPN clients. It is not supported by Stonesoft IPsec VPN Clients and it is not valid for this VPN.
10. Click **OK**. The information is added to the Management Server's internal LDAP user database.

Add additional users in the same way.

Creating Rules for VPN Configuration 3

The authentication connection from VPN clients is allowed in the Firewall Template. Authentication is always required to establish a VPN tunnel. VPN client connections are matched based on Source, Destination, and Service like any other traffic. The example rule matches only specific users and only after the users have already successfully authenticated. We recommend always adding the authentication requirement to rules that are specific to VPN clients.

After the VPN tunnel is established, any connection from the VPN clients to the internal network is matched against the Access rules as usual. The example rule that is created here allows these connections.

This basic configuration scenario does not explain all settings related to VPN Access rules. For more information, see Creating Rules for Policy-Based VPNs (page 958).

▼ To create Access rules for configuration 3

1. Open the policy of one of the Firewall/VPN engines involved in the VPN for editing.
2. Add an IPv4 Access rule in a suitable location in the policy and configure the rule as outlined below:
   - If NAT is active and applied to translate the destination, keep in mind that the Access rules are checked before the translations defined in the NAT rules are applied to the destination IP addresses in the packets.

Table 53.3 Example VPN Rule

Source	Destination	Service	Action	Authentication	Users
Set to ANY	Local internal networks.	Set as needed.	Select **Use VPN**, change the **Action** to **Enforce** and click **Select** to add the VPN element you created.	Set to ANY or to a particular method.	The **stonegate** User Group (in **InternalDomain**)

3. Save the policy.

Refresh the policies of all firewalls involved in the VPN to activate the new configuration. The VPN is established when some traffic matches the Access rules created here.

What's Next?

- Make sure that your NAT rules do not apply a second NAT operation to the VPN clients' IP addresses (defined the NAT Pool for VPN clients in Managing VPN Client Addresses in Configuration 3 (page 902)), see Editing Firewall NAT Rules (page 687).
- Install and configure the VPN clients as instructed in the VPN client's documentation. Many settings that affect how the client behaves are configured through the Management Center. To explore the available options you can define for VPN clients in the Management Center, see Getting Started With VPN Client Settings (page 996). The clients download a configuration file from each gateway they connect to.
- After testing the basic connectivity, you may want to change the IP address allocation method to use the Virtual Adapter to allow queries from the clients to your organization's internal DNS servers, see Configuring Virtual IP Addressing for VPN Clients (page 1002).

Related Tasks

- Monitoring VPNs (page 969)
- Getting Started with User Authentication (page 856)
- Troubleshooting IPsec VPNs (page 1137)

Configuration 4: Basic VPN Hub

Prerequisites: None

In a VPN hub configuration, a gateway is configured to forward VPN traffic between different VPN tunnels. The gateway that does this forwarding is called a *hub* gateway. The gateways that contact each other through a hub are called *spoke* gateways.

The hub gateway must be set up specifically as a hub. The hub configuration is reflected in the topology, the Site definitions, and the VPN rules. The spoke gateways do not require any hub-specific configuration. Following this example configuration, VPN tunnels are established from all spoke gateways to the hub gateway, and all networks of all gateways are configured as reachable through the hub (but actual connections are allowed only as defined in the Firewall Access rules).

Note – There must not be duplicate end-point-to-end-point tunnels in different VPNs. If there are existing VPNs between the Firewall/VPN engines involved in the hub, overlapping configurations must be removed.

This basic configuration scenario explains a configuration in which all connections are defined within the same VPN element, which is usually simpler to set up and maintain than forwarding traffic between VPN tunnels defined in different VPN elements. In this scenario, all Gateways are

Internal Gateways (Stonesoft Firewall/VPN engines controlled by the same Management Server). External Gateways can be added to this configuration even though their creation is not covered in detail in this workflow.

What's Next?

- ▶ This scenario has four parts. Start the configuration in Creating Gateway Elements for VPN Configuration 4.

Creating Gateway Elements for VPN Configuration 4

This basic configuration scenario does not explain all settings related to Gateway elements. For more information, see Defining Security Gateways (page 923).

If there is a previously configured Gateway element for a Firewall/VPN engine you plan to use in this configuration, use the existing element instead of creating a new one.

Following these instructions, the VPN will use all interfaces that contain the default "Any Network" element (the default route) as endpoints for the VPN.

▼ To create Gateway elements for configuration 4

1. Select **Configuration→Configuration→VPN** to switch to the VPN Configuration view.
2. Right-click **Gateways** in the element tree and select **New→Internal Security Gateway**. The Internal Security Gateway Properties dialog opens.

Illustration 53.23 Internal Security Gateway Properties - General

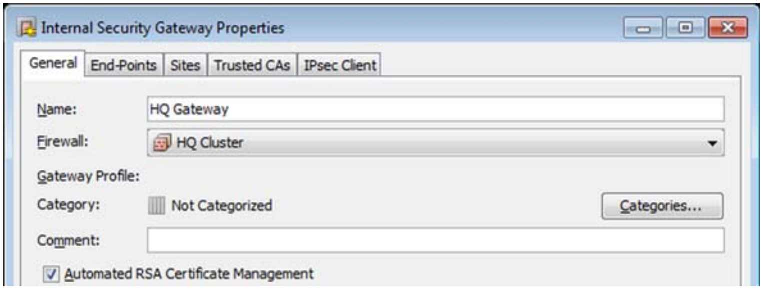

Internal Security Gateway Properties

Name: HQ Gateway

Firewall: HQ Cluster

Gateway Profile:

Category: Not Categorized

Comment:

☑ Automated RSA Certificate Management

3. Select the **Firewall** engine that the Gateway element represents in VPNs.
4. (*Optional*) Change the **Name** that was added based on the Firewall element's name.

Create further Gateway elements to represent each Firewall/VPN engine involved in this VPN.

Creating a VPN Element for VPN Configuration 4

This basic configuration scenario does not explain all settings related to VPN elements. For more information, see Defining Policy-Based VPNs (page 949).

▼ To create a VPN element for configuration 4

1. Right-click **VPNs** in the element tree and select **New VPN**. The VPN Properties dialog opens.

Illustration 53.24 VPN Profile Properties

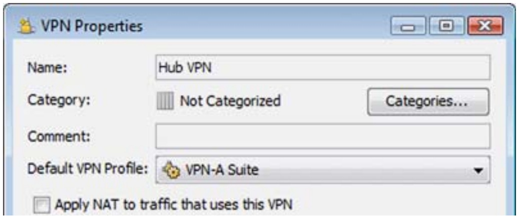

VPN Properties

Name: Hub VPN

Category: Not Categorized

Comment:

Default VPN Profile: VPN-A Suite

□ Apply NAT to traffic that uses this VPN

2. Give the new element a unique **Name**.
3. Select **VPN-A Suite** from the **VPN Profile** list.
4. If you want to apply NAT rules to the communications that go through the VPN, select **Apply NAT Rules to Traffic That Uses This VPN**. This does not affect the communications that the two internal gateways have with each other to set up and maintain the VPN. Those are always matched to the NAT rules.
5. Click **OK**. The VPN editing view opens on the **Overall Topology** tab.
6. Drag and drop the hub Gateway to **Central Gateways**.

Illustration 53.25 VPN Editing View - Overall Topology

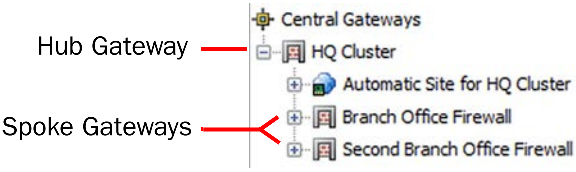

7. Drag and drop the other Gateways on top of the hub Gateway so that the Gateways are added as branches (spokes) under the hub Gateway as in the illustration above. These can include any other Internal or External Gateways.
8. Save the VPN, but do not close the VPN editing view yet. This intermediate save is required to store the changes in the database for the next operation.

Defining Site Properties for VPN Configuration 4

This basic configuration scenario does not explain all settings related to Site elements. For more information, see Defining Sites for VPN Gateways (page 933).

▼ To define a Hub Site for configuration 4

1. Right-click the hub Gateway and select **New→Site**. The Site Properties dialog opens.
2. Give the new element a unique **Name**.

Illustration 53.26 Site Properties - General Tab

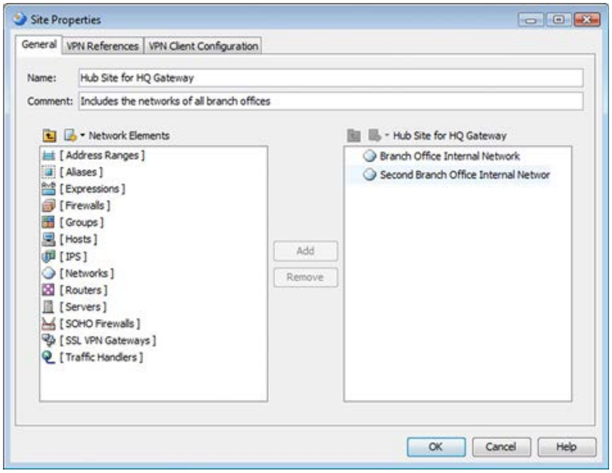

Name: Hub Site for HQ Gateway

Comment: Includes the networks of all branch offices

Network Elements
- [Address Ranges]
- [Aliases]
- [Expressions]
- [Firewalls]
- [Groups]
- [Hosts]
- [IPS]
- [Networks]
- [Routers]
- [Servers]
- [SOHO Firewalls]
- [SSL VPN Gateways]
- [Traffic Handlers]

Hub Site for HQ Gateway
- Branch Office Internal Network
- Second Branch Office Internal Networ

3. Add all networks protected by the spoke gateways to the Site contents on the right.
   - After you do this, the Site should contain all remote IP addresses that are used in spoke-to-hub traffic that is forwarded from the hub to other spokes.
   - The Site should not contain the hub gateway’s local networks. These are defined using the automatic Site management features in this example.
4. Switch to the **Site References** tab.
5. Select **Enable** for this VPN element and deselect it for all other VPNs. Note that the Site is still shown in all VPNs, but is grayed-out (disabled) and not included in the configuration.
6. Select **Hub** as the Mode. This activates VPN hub-related features for the Gateway.
7. Click **OK** to close the dialog and return to the main VPN editing view.
8. Switch to the **Tunnels** tab.
9. Check that the **Validity** column in the **Gateway<->Gateway** and the **End-Point<->End-Point** tables has a green check mark to indicate that there are no problems.
   - If the **Validity** column of a tunnel has a warning icon, see the **Issues** panel to check what the problem is (if the panel is not displayed, you can open it through the **View** menu or the right-click menu for the Validity cell of a rule). If Issues are displayed, correct them as indicated. Long issues are easiest to read by hovering over the issue text so that the text is shown as a tooltip.
10. Save the VPN.

Creating Rules for VPN Configuration 4

The rules in this example allow connections between hosts in protected networks of all Gateways to connect to all other protected networks. VPN rules are matched based on Source, Destination, and Service like any other rules.

This basic configuration scenario does not explain all settings related to VPN Access rules. For more information, see Creating Rules for Policy-Based VPNs (page 958).

▼ To create Access rules for configuration 4

1. Open the policy of the Firewall/VPN engine that is configured as the hub gateway in the VPN for editing.
2. Add an IPv4 Access rule in a suitable location in the policy.
   - Make sure that rules for sending traffic through the VPN are above any other rules that match the same traffic with Allow, Discard, or Refuse as their action.
   - Traffic that you do not want to send through the VPN must not match this rule. Traffic that is not routable through the VPN is dropped if it matches this rule.
3. Fill in the rule as outlined below.
   - If NAT is enabled in the VPN, keep in mind that the Access rules are checked before the translations defined in the NAT rules are applied to the source and/or destination IP addresses in the packets.
   - To ensure this rule does not match other traffic, you can add the VPN you created in the **Source VPN** cell.

Table 53.4 Example VPN Rule for Spoke-to-Spoke Forwarding

Source	Destination	Service	Action
Remote internal networks.	Remote internal networks.	Set as needed.	Select **Use VPN**, then change the **Action** to **Forward** and click **Select** to add the VPN element you created.

4. Add two more rules in a suitable location in the policy to allow traffic between the hub Gateway’s local protected networks to the spoke Gateway’s protected networks. Two rules are created here to allow the different directions of traffic.

Table 53.5 Example VPN Rules

Source	Destination	Service	Action
Local internal networks.	Remote internal networks.	Set as needed.	Select **Use VPN**, then change the **Action** to **Enforce** and click **Select** to add the VPN element you created.
Remote internal networks.	Local internal networks.	Set as needed.	Select **Use VPN**, then change the **Action** to **Enforce** and click **Select** to add the VPN element you created.

5. Save the policy.

6. Add rules in the policies of all Firewall/VPN engines involved as outlined in Step 4.

Caution – If you continue to use this VPN, change the pre-shared key periodically (for example, monthly) to ensure continued confidentiality of your data. See Renewing or Generating Pre-Shared Keys (page 990). Alternatively, you can switch to certificate-based authentication by creating a custom VPN profile.

Refresh the policies of all firewalls involved in the VPN to activate the new configuration. The VPN is established when traffic matches the Access rules.

What's Next?

- ▶ If NAT is required for this VPN, add the NAT rules as well. See Creating NAT Rules for Policy-Based VPN Traffic (page 963).

CHAPTER 54

CONFIGURING IPSEC VPNS

IPsec (Internet Protocol Security) VPNs (Virtual Private Networks) allow creating secure, private connections through networks that are not otherwise secure.

The following sections are included:

Getting Started With IPsec VPNs

Prerequisites: If you want to allow access to VPN clients: Getting Started with User Authentication

VPNs in Stonesoft Firewall/VPN are implemented according to the IPsec standard (an extension of the IP protocol). For more background information on IPsec and the VPN configuration, see the *Reference Guide*. For information about the Stonesoft SSL VPN product, see the *Stonesoft SSL VPN Administrator's Guide*. VPNs are not supported on Layer 2 Firewalls.

What VPNs Do

A virtual private network (VPN) extends a secured private network over public networks by *encrypting* connections so that they can be transported over insecure links without compromising confidential data. For this purpose, the devices that create the VPN check the identity of the other parties than by the way of *authentication*. A VPN also includes *integrity checking* to ensure the communications are not tampered with.

What Do I Need to Know Before I Begin?

There are two types of IPsec VPNs in Stonesoft Firewall/VPN. The main difference between the two is how traffic is selected to use the VPN:

- *Policy-based VPNs* are configured using **VPN** elements. The firewall IPv4 Access rules define which traffic is sent to the VPN and which traffic is allowed out of the VPN.
- *Route-based VPN* tunnels are configured using the **Route-Based VPN** element. Any traffic that is routed to firewall interfaces that are designated as end-points for a VPN tunnel is sent into the VPN tunnel. If the traffic is allowed by the Access rules, it is automatically sent through the tunnel to the peer end-point.

You can also use a policy-based VPN to provide encryption for the Route-Based VPN in Tunnel Mode. See Using the Route-Based VPN in Tunnel Mode (page 967).

You can create VPNs between gateway devices or between a VPN client and a gateway device:

- A *gateway-to-gateway VPN* is created between two or more gateway devices that provide VPN access to several hosts in their internal networks.
- A *client-to-gateway VPN* is created between a VPN client running on an individual computer and a gateway device. The recommended option is to use the Stonesoft IPsec VPN Client (available on Windows). Third-party IPsec-compatible VPN clients can also be used.

Note – VPN clients that are a part of a vendor-specific VPN gateway solution are generally incompatible with gateways from other vendors.

Limitations

- Client-to-gateway VPNs can only be created in Policy-Based VPNs. It is not possible to create a client-to-gateway tunnel using the Route-Based VPN.
- All VPNs you configure for VPN clients must be valid for Stonesoft IPsec VPN Clients even if you use only third-party VPN client software.
- It is only possible to use Multi-Link with the Route-Based VPN in Tunnel Mode. See Using the Route-Based VPN in Tunnel Mode (page 967).
- You cannot use the same End-Point IP addresses in both the Route-Based VPN and a policy-based VPN configured using VPN elements.
- You cannot use End-Points in standby or aggregate mode with the Route-Based VPN. Route-Based VPN tunnels are always active.

- If your Firewall/VPN engine installation is specifically configured in a restricted operating mode to comply with regulatory requirements, some VPN options are unavailable to you.
- Version-specific limitations in supported features for different Stonesoft Firewall/VPN versions are listed in the *Release Notes* for the versions you are using. The Management Center automatically prevents the use of unsupported settings based on engine version.

Proceed as explained in the Configuration Overview or follow one of the Basic Policy-Based VPN Configurations (page 887) to quickly create some types of basic VPNs.

Configuration Overview

Illustration 54.1 Elements in the VPN Configuration (Excluding Certificate-Related Elements)

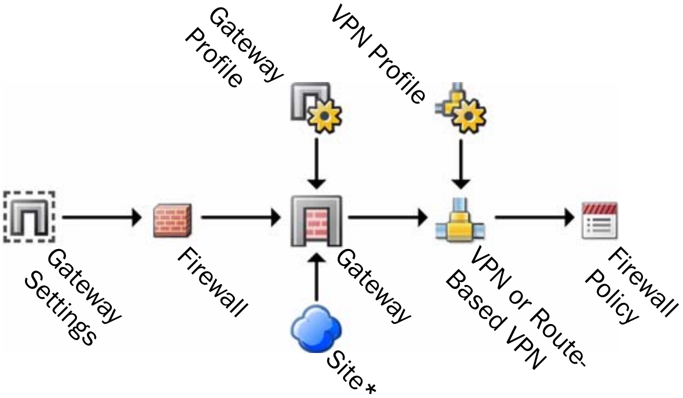

* Policy-based VPN only

The Gateway Settings element shown above is not part of the workflow below because the default settings should be used in most cases (see Advanced VPN Tuning (page 991) for more information). Otherwise, the following workflow covers all VPN configurations:

1. (*Route-Based VPN only*) Create Tunnel Interfaces on the Firewall/VPN to define end-points for tunnels in the Route-Based VPN as explained in Defining Tunnel Interfaces for Firewalls (page 435).
2. (*Route-Based VPN only*) Define which networks are reachable through each Tunnel Interface as explained in Defining Routing for the Route-Based VPN (page 604).
3. (*Optional*) If you are configuring a VPN with an external device, you may want to create a new Gateway Profile specific to the device. See Defining Gateway Profiles (page 921).
4. Add the necessary number of gateway elements to represent the physical VPN devices, see Defining Security Gateways (page 923). These define the VPN end-points (gateway IP addresses) and the Sites (see the next point).
5. (*Policy-based VPNs only*) Configure the Gateway's Sites. These define the IP addresses that can be made routable through VPNs, see Defining Sites for VPN Gateways (page 933). The Sites can be adjusted in different VPNs that the Gateway establishes.
6. (*Optional*) If the existing VPN Profiles do not have suitable settings for your new VPN, create a new one as explained in Defining VPN Profiles (page 939). This defines the IPsec settings (authentication, encryption, and integrity checking).

7. Define the VPN in one of the following ways:
    - Create a new VPN element as explained in Defining Policy-Based VPNs (page 949). This defines the topology (which gateways create tunnels with each other).
    - Configure the Route-Based VPN element as explained in Editing the Route-Based VPN (page 964).
8. Create certificates, if necessary. See Getting Started With VPN Certificates (page 972).
9. Add the necessary Access rules according to the type of VPN:
    - (*Policy-Based VPNs*) Add the IPv4 Access rules and, if necessary, the IPv4 NAT rules for VPN traffic. See Creating Rules for Policy-Based VPNs (page 958). This also activates the VPN on the engines.
    - (*Route-Based VPN*) Add Access rules to allow traffic between the internal network and the networks that are reachable through the Route-Based VPN. See Editing Access Rules (page 672).

What's Next?

▶ Start by Configuring IPsec VPNs.

Configuring IPsec VPNs

This workflow contains steps for all kinds of IPsec VPN configurations. Alternative next steps are included as necessary to achieve a particular type of configuration. Alternatively, you may want to follow a simplified workflow for building a particular type of VPN, see Basic Policy-Based VPN Configurations (page 887).

Continue the configuration as explained below, in the first section that applies to you.

What's Next?

▶ If you are configuring a VPN with an external gateway, you may want to create a new Gateway Profile if the existing profiles are not suitable in this case. See Defining Gateway Profiles (page 921).

▶ Otherwise, begin the configuration in Defining Security Gateways (page 923).

Related Tasks

▶ Configuration Overview (page 919)

▶ Getting Started With VPN Client Settings (page 996)

Defining Gateway Profiles

Prerequisites: None

Internal Gateways always use a default profile that is assigned according to the currently installed software version and cannot be changed manually. Gateway profiles can be used with external gateways to set certificate-related options (used with some gateways that do not support some operations) and to restrict the options to a supported set to prevent configuration errors. If you do not see a need to utilize these settings, you can use the unrestricted **Default (all capabilities)** profile.

What's Next?

- ▶ To add or modify a custom Gateway Profile, continue in Defining a Custom Gateway Profile.
- ▶ To configure a VPN using the existing Gateway profiles, continue in Defining Security Gateways (page 923).

Defining a Custom Gateway Profile

The properties on the **General** tab are meant for advanced users. The default values are the recommended values. These options affect the VPN directly.

The **IKE Capabilities** and **IPsec Capabilities** are not directly used in a VPN (the settings are selected for use in the VPN Profile element), but instead define a set of options that the Gateway supports, so that the system can automatically check for misconfigured settings.

▼ To define a custom Gateway Profile

1. Select **Configuration→Configuration→VPN**. The VPN Configuration view opens.

Illustration 54.2 Gateway Profiles in the VPN Configuration Tree

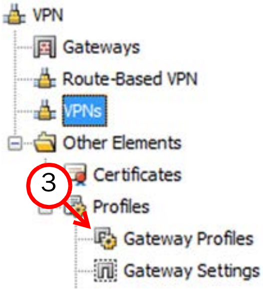

2. Browse to **Other Elements→Profiles**.
3. Right-click the **Gateway Profiles** branch and select **New Gateway Profile**. The Gateway Profile Properties dialog opens.

Illustration 54.3 Gateway Profile Properties

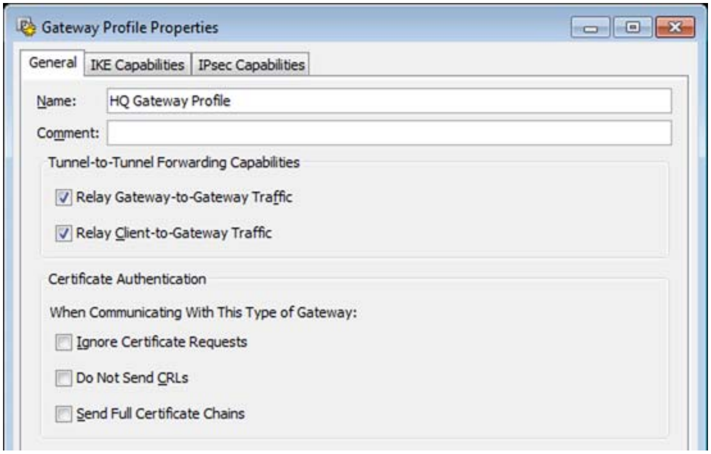
Gateway Profile Properties

General | IKE Capabilities | IPsec Capabilities

Name: HQ Gateway Profile

Comment:

Tunnel-to-Tunnel Forwarding Capabilities

☑ Relay Gateway-to-Gateway Traffic

☑ Relay Client-to-Gateway Traffic

Certificate Authentication

When Communicating With This Type of Gateway:

□ Ignore Certificate Requests

□ Do Not Send CRLs

□ Send Full Certificate Chains

4. Enter a unique **Name** and optionally a **Comment** for the element.
5. Select options for Tunnel-to-Tunnel Forwarding Capabilities:

Setting	Description
Relay Gateway-to-Gateway Traffic	Select this to indicate whether the Gateways using the profile are capable of forwarding gateway-to-gateway VPN traffic to other gateway-to-gateway VPNs. This reduces the number of tunnels created by default for VPNs involving this Gateway when you define forwarding from one VPN to another in the VPN element.
Relay Client-to-Gateway Traffic	This option is shown only because the setting is used in the default profiles for different versions of Stonesoft Firewall/VPN. This setting is not relevant to custom configurations.

6. Select options for Certificate Authentication:

Setting	Description
Ignore Certificate Requests	Selecting this option prevents the Firewall/VPN from processing certificate request payloads in IKE messages with Gateways that use this profile. Gateways use certificate requests to ask the other gateway to forward its certificates or related information.
Do Not Send CRLs	Selecting this option prevents the Firewall/VPN from sending certificate revocation lists in IKE messages with Gateways that use this profile. If the CRLs are available to the other gateway through other ways, it is not necessary to send them through IKE messages.
Send Full Certificate Chains	Selecting this option makes the Firewall/VPN send the full certificate chain (from the node’s certificate to the root Certificate Authority certificate) in IKE messages with Gateways that use this profile.

7. Switch to the **IKE Capabilities** tab and select the options that the device supports for IKE SA negotiations.

8. Switch to the **IPsec Capabilities** tab and select the options that the device supports for IPsec SA negotiations.

What's Next?

- Continue the configuration by Defining Security Gateways.

Defining Security Gateways

Prerequisites: None

The physical devices that establish the VPN are represented by Security Gateway elements in the configuration. Only one element per device is needed in the system, even if there are many VPNs. You can add more than one Gateway element to represent the same firewall, but each Gateway element reserves an end-point IP address that cannot be used by other elements. The following types of Gateway elements can be configured:

Gateway Type	Description
Internal Security Gateway	Stonesoft Firewall/VPN engines that are managed by the same Management Server (and administrative Domain) to which you are currently connected with your Management Client.
External Security Gateway	All other security gateways, including Stonesoft Firewall/VPN devices that are managed by some different Management Server than the one to which you are connected (or that are configured under a different administrative Domain).
IPSec Client Gateway (*policy-based VPNs only*)	All Stonesoft and third-party IPsec VPN clients are represented by the default IPsec Client Gateway element.

Proceed to the first section below that applies to you.

What's Next?

- If you do not already have any Gateway elements, start by Creating a New Security Gateway Element (page 924).
- To edit internal Gateway IP address settings, see Defining End-Points for Internal Security Gateways (page 924).
- To edit external Gateway IP address settings, see Defining End-Points for External Security Gateways (page 927).
- To change certificate acceptance settings, see Defining Trusted CAs for a Gateway (page 929).
- To configure settings for VPN clients that connect to a gateway, see Defining Gateway-Specific VPN Client Settings (page 930).

Creating a New Security Gateway Element

▼ To create a New Gateway

1. Select **Configuration→Configuration→VPN**. The VPN Configuration view opens.
2. Right-click the **Gateways** branch in the element tree and select one of the following options:
   - **New→Internal Gateway** to represent a firewall managed by this Management Server.
   - **New→External Gateway** to represent a third-party VPN device or a Stonesoft Firewall managed by a different Management Server.
3. Enter the **Name** and optionally a **Comment** for the element.
4. (*Optional, Internal Gateways only*) Deselect **Automated RSA Certificate Management** if you prefer to create and renew certificates manually when certificates are needed.
   - Automated certificate management is only for RSA certificates.
   - Even when the option is selected, certificates are only created when needed.
5. Select the device this Gateway represents:

Gateway Type	Configuration
Internal Security Gateway	Select the Firewall engine that the Gateway element represents in VPNs.
External Security Gateway	Click **Select** and select the Gateway Profile that contains information on the capabilities of the external Gateway.

Continue the configuration in the correct section according to the type of Gateway:

What's Next?

- ▶ Defining End-Points for Internal Security Gateways
- ▶ Defining End-Points for External Security Gateways (page 927)
- ▶ Defining VPN Profiles (page 939)

Defining End-Points for Internal Security Gateways

Each end-point is dedicated for one Gateway element. Any IP address that is already selected as an end-point for some other Gateway element is not shown on the End-Points list for subsequent Gateways you create for the same Firewall/VPN engine. Each Gateway element can be used in several VPNs. However, you cannot use the same end-point in a Route-Based VPN tunnel and a policy-based VPN tunnel. If you want to create both policy-based VPNs and Route-Based VPN tunnels using the same Gateway, you must define unique end-points for each type of VPN.

▼ To define end-points for Internal Gateways

1. Switch to the **End-Points** tab in the Gateway properties. The IP addresses available for use as end-points are displayed.

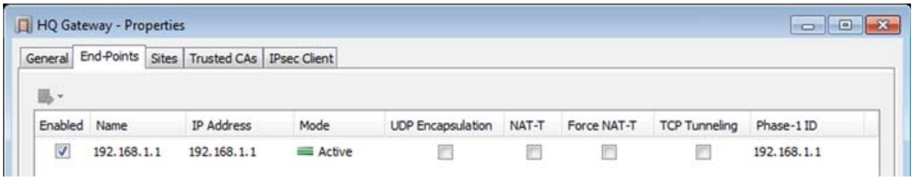

Enabled	Name	IP Address	Mode	UDP Encapsulation	NAT-T	Force NAT-T	TCP Tunneling	Phase-1 ID
☑	192.168.1.1	192.168.1.1	Active	□	□	□	□	192.168.1.1

2. (*Optional*) Change the selection of IP address(es) that you want to use as end-points in VPNs.
   - Typically, these are IP address(es) that belong to interface(s) towards the Internet, which are automatically selected based on the firewall's default routing table.
   - If loopback IP addresses are defined for the Internal Security Gateway, you can select a loopback IP address as the End-Point IP address. See Configuring Loopback IP Addresses for Firewalls (page 450) for more information about loopback IP addresses.

Note – VPN clients can only connect to end-points that have an IPv4 address.

   - On clustered firewalls, the IP addresses are CVIs.
   - (*Policy-based VPN only*) If you have more than one Internet connection, select an IP address from each ISP to make Multi-Link load balancing and failover possible for VPNs.

3. Double-click the end-point that you selected for this Gateway. The Internal End-Point Properties dialog opens.

Illustration 54.4 Internal End-Point Properties

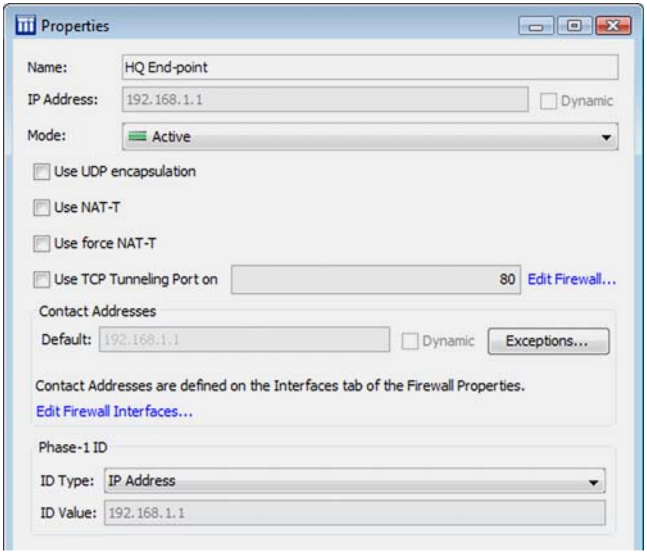

Properties

Name: HQ End-point

IP Address: 192.168.1.1 □ Dynamic

Mode: 🟩 Active

□ Use UDP encapsulation

□ Use NAT-T

□ Use force NAT-T

□ Use TCP Tunneling Port on 80 Edit Firewall...

Contact Addresses

Default: 192.168.1.1 □ Dynamic Exceptions...

Contact Addresses are defined on the Interfaces tab of the Firewall Properties.

Edit Firewall Interfaces...

Phase-1 ID

ID Type: IP Address

ID Value: 192.168.1.1

4. (*Optional*) Give the end-point a **Name**.

5. (*Optional, policy-based VPNs only*) Select the **Mode** to define how the end-point is used in a Multi-Link configuration. You can override these settings in each individual VPN.
   - **Active**: use the tunnel(s) of this end-point whenever possible and balance the traffic across the tunnels based on a performance measurement or based on the links' relative bandwidths.
   - **Aggregate**: use the tunnel(s) of this end-point whenever possible and balance each connection between the tunnels in round robin fashion.
   - **Standby**: use the tunnel(s) of this end-point only if the Active or Aggregate end-points cannot be used.

6. (*Optional, Firewall/VPN versions prior to 5.2 only*) Select **Use UDP Encapsulation** if you want to encapsulate IPsec communications in UDP packets using Stonesoft Firewall/VPN's proprietary method.

7. (*Optional*) Select one of the Use NAT-T options to activate encapsulation for NAT traversal in gateway-to-gateway VPNs, which may be needed to traverse a NAT device at the local or at the remote gateway end.

Option	Description
Use NAT-T	Select this option if you if you want to allow encapsulating the IPsec communications in standard NAT-T UDP packets in gateway-to-gateway VPNs when the gateways detect that a NAT operation is applied to the communications. If both gateways do not support this option, the option is ignored.
Use force NAT-T	Select this option to force NAT-T even when the gateways do not detect a NAT operation being applied to the communications. If both gateways do not support this option, the VPN fails to establish.

- The gateway always allows VPN clients to use NAT-T regardless of these settings.
- NAT-T always uses the standard UDP port 4500.

Note – If the external IP address is private and translated to a public IP address by an external NAT device, make sure Contact Addresses and Locations are correctly defined in the Firewall properties.

8. (*Optional*) Select **Use TCP Tunneling Port on** if you want to tunnel Stonesoft IPsec VPN Client communications with this Gateway end-point in a TCP connection to bypass a traffic filtering device that does not allow standard IPsec ports to pass or to traverse a NAT device.
   - This option may not be supported by all external gateways. Support is required at both ends of each tunnel.
9. If necessary, change the default **Contact Address** and/or add **Exceptions** for the Location(s) of other Gateway(s) involved in the VPN.
   - The two contacting Gateways must be in different Locations (note that external gateways are always in the Default Location).
   - In clusters, VPNs use CVIs, not NDIs.
   - For more information on Locations and Contact Addresses, see Configuring System Communications (page 63).

Example **The internal gateway is behind a NAT device. The real address is defined as the End-Point address, because the IP address is also used for identification inside the encrypted traffic. Contact must be made using the translated address, so it is defined as a Contact Address.**

10. In the Phase-1 settings, change the **ID Type** to according to your environment.
    - The ID identifies the Gateways during the IKE SA negotiations.
    - The **Distinguished Name** type is only valid in certificate-based authentication.
    - The **IP Address** type is not valid for end-points with a dynamic IP address.
11. Enter an **ID Value** if you selected **DNS Name** or **Email** as the type. The value for IP address is filled in automatically.

Continue the configuration as explained below in the first section that applies to you.

What's Next?

- If you want to use certificate authentication and you want to limit which of the Certificate Authorities you define in the system this Gateway considers trusted, continue in Defining Trusted CAs for a Gateway (page 929).
- If you plan to allow VPN clients to establish VPNs with this Gateway, continue in Defining Gateway-Specific VPN Client Settings (page 930).
- If you are configuring the Route-Based VPN and need to create a VPN Profile, proceed to Defining VPN Profiles (page 939).
- If you are configuring the Route-Based VPN and already have a suitable VPN Profile, proceed to Defining Route-Based VPN Tunnels (page 965).
- Otherwise, continue in Defining Sites for VPN Gateways (page 933).

Defining End-Points for External Security Gateways

To define end-points for internal Gateways

1. Switch to the **End-Points** tab in the Gateway properties. The end-point table is displayed.
2. Click the **Add** button below the table. The External End-Point Properties dialog opens.

Illustration 54.5 External End-Point Properties

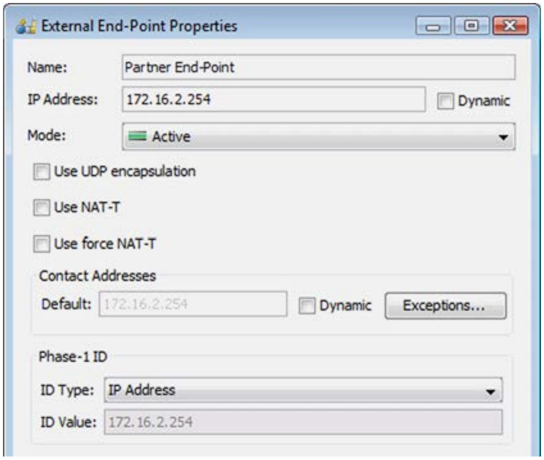

Name: Partner End-Point

IP Address: 172.16.2.254 □ Dynamic

Mode: Active

□ Use UDP encapsulation

□ Use NAT-T

□ Use force NAT-T

Default: 172.16.2.254 □ Dynamic

ID Type: IP Address

ID Value: 172.16.2.254

3. (*Optional*) Enter a **Name** for the end-point.
4. Define the IP address for the end-point:
   - If the end-point has a static (manually defined) IP address, enter the **IP Address**. This must be the IP address that is configured for the external device in its configuration.
   - If the end-point has a dynamic (DHCP-assigned) IP address, select **Dynamic**.

5. (*Optional*) Select the **Mode** to define how the system treats the end-point in a Gateway with multiple end-points. This is a default setting for the tunnels that are generated for VPNs that use this Gateway. You can override the Mode setting in each VPN.
   - **Active**: use the tunnel(s) to this end-point whenever possible and balance the traffic across the tunnels based on a performance measurement or based on the links' relative bandwidths.
   - **Aggregate**: use the tunnel(s) of this end-point whenever possible and balance each connection between the tunnels in round robin fashion.
   - **Standby**: use the tunnel(s) to this end-point only if the Active or Aggregate end-points cannot be used.
6. (*Optional, Firewall/VPN versions prior to 5.2 only*) Select **Use UDP Encapsulation** if you want to encapsulate IPsec communications in UDP packets using Stonesoft Firewall/VPN's proprietary method.
   - This method is proprietary to Stonesoft Firewall/VPN. Do not use it for external gateway devices from other vendors. Support is required at both ends of each tunnel.
7. (*Optional*) Select one of the Use NAT-T options to activate encapsulation for NAT traversal in gateway-to-gateway VPNs, which may be needed to traverse a NAT device at the local or at the remote gateway end.

Option	Description
Use NAT-T	Select this option if you if you want to allow encapsulating the IPsec communications in standard NAT-T UDP packets in gateway-to-gateway VPNs when the gateways detect that a NAT operation is applied to the communications. If both gateways do not support this option, the option is ignored.
Use force NAT-T	Select this option to force NAT-T even when the gateways do not detect a NAT operation being applied to the communications. If both gateways do not support this option, the VPN fails to establish.

   - This option may not be supported by all external gateways, but support is required at both ends of each tunnel. If both gateways do not support this option, the option is ignored.
   - Stonesoft IPsec VPN Clients are always allowed to use this encapsulation regardless of this setting.
   - NAT-T always uses the standard UDP port 4500

Note – If the external IP address is private and translated to a public IP address by an external NAT device, make sure Contact Addresses and Locations are correctly defined in the Firewall properties.

8. If necessary, change the default **Contact Address** and/or add **Exceptions** for the Location(s) of other Gateway(s) involved in the VPN.
   - The Contact Address must be defined if the IP address for contacting this Gateway is different from the IP address the Gateway actually has on its interface (for example, because of NAT).
   - For more information on Locations and Contact Addresses, see Getting Started with System Communications (page 64).

Example **An external gateway is behind a NAT device. The real address is defined as the End-Point address, because the IP address is also used as the Phase 1 ID inside the encrypted traffic. Contact must be made using the translated address, so it is defined as a Contact Address.**

9. In the Phase-1 settings, select the **ID Type** for your preferred option.
   - The ID identifies the Gateways during the IKE SA negotiations.
   - The **Distinguished Name** type is only valid in certificate-based authentication.
   - The **IP Address** may not work as an ID if the address is translated using NAT.
10. Enter an **ID Value** of the correct type.
    - If the end-point has a dynamic IP address, you must enter a specific IP address as the value for the **IP Address** type.
    - If the end-point has a static IP address, the value for the **IP Address** type is filled in automatically according to the IP address you defined for this end-point.
    - Make sure that the **ID Value** matches the identity configured on the external gateway device.

What's Next?

- ▶ If you want to use certificate authentication and you want to limit which of the Certificate Authorities you define in the system this Gateway considers trusted, continue in Defining Trusted CAs for a Gateway.
- ▶ If you are configuring the Route-Based VPN and need to create a VPN Profile, proceed to Defining VPN Profiles (page 939).
- ▶ If you are configuring the Route-Based VPN and already have a suitable VPN Profile, proceed to Defining Route-Based VPN Tunnels (page 965).
- ▶ Otherwise, continue in Defining Sites for VPN Gateways (page 933).

Defining Trusted CAs for a Gateway

Certificate Authorities (CA) verify certificate authenticity with their signatures. Gateways accept certificates only from specifically configured trusted CAs. By default, the Gateways trust all VPN CAs that are currently defined in the system, but you can restrict the trusted CAs as instructed here. You can alternatively or additionally restrict trusted CAs also in the VPN Profiles.

For external gateways, the system uses the trusted CA definition in the External Gateway element to check that all gateways have the necessary certificates.

▼ To define CAs trusted by a Gateway

1. Switch to the **Trusted CAs** tab in the Gateway properties.
2. Select the Limit Trusted Certificate Authorities options:

Option	Description
Trust All Defined	The gateway accepts any valid CA that is configured in your system, unless restricted in the VPN element.
Trust Only Selected	Restricts the trusted CAs and activate the controls below and select Enabled for the CAs that the Gateway needs to trust.

What's Next?

- If you are creating an Internal Security Gateway and you plan to use VPN clients to establish VPNs with this Gateway, continue in Defining Gateway-Specific VPN Client Settings.
- Otherwise, continue in Defining Sites for VPN Gateways (page 933).

Defining Gateway-Specific VPN Client Settings

The Internal Security Gateway element contains settings for assigning valid IP addresses to VPN clients for connections through the VPN to the internal network. If you use Stonesoft IPsec VPN Clients, you should configure the Virtual Adapter, since the alternative NAT Pool method does not allow the Stonesoft IPsec VPN Client computers to use your organization's internal DNS servers. Virtual IP addressing works with all Stonesoft IPsec VPN Client versions and with third-party VPN clients that support this feature.

For full information on all configuration options, see Managing VPN Client IP Addresses (page 1000).

Note – The Virtual Adapter IP addresses must be assigned by a DHCP server. It is not possible to define the IP addresses in the VPN client or in the VPN gateway configuration. When using the firewall's internal DHCP server in a Single Firewall installation, use the IP address of the firewall interface on which the internal DHCP server is enabled as the IP address of the DHCP Server element.

If an Internal Security Gateway that contains IPsec Client settings is used in the Route-Based VPN, the IPsec Client settings are ignored.

▼ To configure the VPN client settings (for Stonesoft IPsec VPN Clients)

1. Switch to the **IPsec Client** tab in the Gateway properties.

Illustration 54.6 Internal Security Gateway Properties - IPsec Client Tab

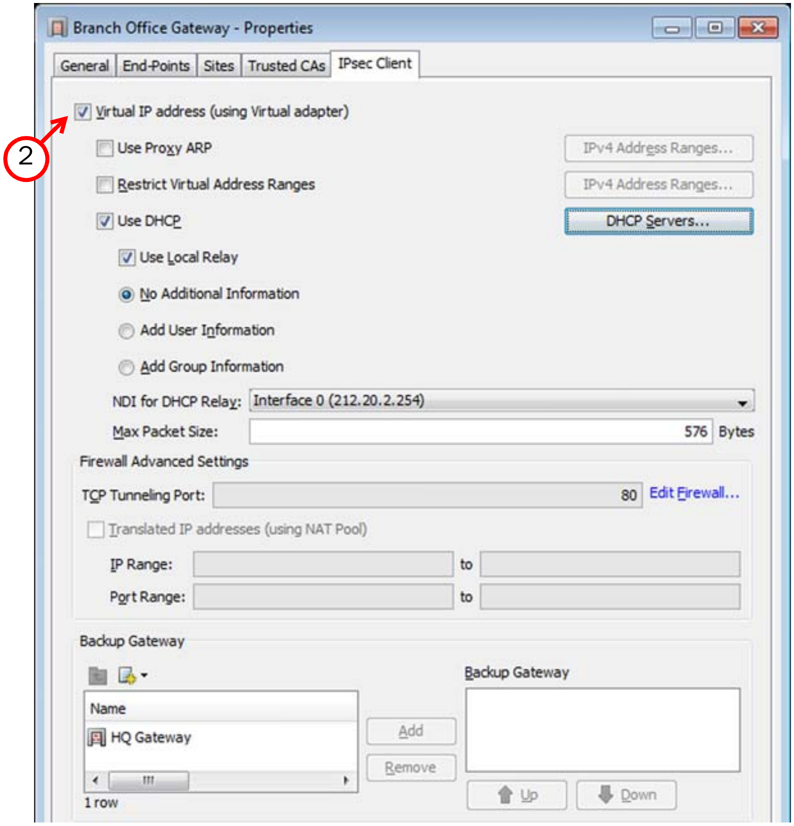

Branch Office Gateway - Properties

☑ Virtual IP address (using Virtual adapter)

□ Use Proxy ARP

□ Restrict Virtual Address Ranges

☑ Use DHCP

☑ Use Local Relay

● No Additional Information

○ Add User Information

○ Add Group Information

NDI for DHCP Relay: Interface 0 (212.20.2.254)

Max Packet Size: 576 Bytes

Firewall Advanced Settings

TCP Tunneling Port: 80

□ Translated IP addresses (using NAT Pool)

IP Range: to

Port Range: to

Backup Gateway

Name
HQ Gateway

Backup Gateway

2. Select **Virtual IP Address (using Virtual adapter)**. The other options in the dialog are enabled.
3. Configure the settings as explained in the table below:

Table 54.1 Virtual IP Address Settings

Setting	Configuration
Use Proxy ARP	Select this option to make the Firewall/VPN engine respond to ARP requests for the virtual IP address range. Click **IPv4 Address Ranges** to select the address range to define the scope for this option (all virtual IP addresses).
Restrict Virtual Address Ranges	(*Optional*) Select this option and click **IPv4 Address Ranges** to select the addresses that the DHCP server is allowed to assign.
Use DHCP	Select this option to define DHCP settings for the VPN clients. Click the **DHCP Servers** button and select the correct external DHCP Server that assigns the IP addresses.
Firewall Advanced Settings	Deselect **Translated IP Addresses (using NAT Pool)** if the option is selected. This disables the NAT Pool feature on the firewall.

Table 54.1 Virtual IP Address Settings (Continued)

Setting	Configuration
Use Local Relay	(*Optional*) Select this option to force the use of unicast DHCP relay messages even if the DHCP server is in a directly connected network in relation to the firewall engine. By default, the Firewall/VPN engine sends a normal DHCP client broadcast message to a DHCP server that is in a directly connected network. Select the option for what additional information to add to the messages. Select the **NDI for DHCP Relay** to be used as the source address for the DHCP packets when querying the DHCP server (the interface towards the DHCP server), and adjust the **Max Packet Size** if necessary in your network.
Backup Gateway	(*Optional*) Select the **Backup Gateways** that Stonesoft IPsec VPN Client version 5.1 and higher use if this gateway is unavailable and organise them in the order you want them to be contacted. This removes the need for the user to launch new connections to different gateways manually. Each switchover launches a prompt that allows the user to confirm the switchover. If the backup gateway's certificate authority is not trusted, the user can manually approve the certificate fingerprint and continue.

What's Next?

▶ Continue by Defining Sites for VPN Gateways (page 933).

Defining Sites for VPN Gateways

Prerequisites: Creating a New Security Gateway Element

You must define Sites for all Internal and External Gateways that are used in policy-based VPNs. Sites are not used in the configuration of the Route-Based VPN. If an Internal Security Gateway that contains a Site is used in the Route-Based VPN, the Site-related settings are ignored. The Site element defines the internal IP addresses that send or receive traffic through the VPN:

- Connections are only allowed to use the VPN if the source and destination IP addresses are included in the Sites of the Gateways involved. Other traffic is rejected.
- If you activate NAT for tunneled traffic, you must add the translated addresses in the Site definitions to make those IP addresses valid to be used in the VPN tunnel.
- The addresses must be unique at each end of a VPN tunnel. Use NAT to create unique address spaces if the addresses overlap.
- If you use a Gateway in several VPNs, only complete Site elements can be included or excluded on a VPN-to-VPN basis. All addresses in a Site are always included and excluded together.
- For internal Gateways, an automatic Site element is available and enabled by default. The IP address definitions of the automatic Site are created and updated based on routing definitions. However, NAT addresses must always be added manually.

The IP address information is also checked in the VPN establishment phase. When creating VPNs with external Gateways, make sure the IP address spaces of both gateways are defined identically in the SMC and on the external device, or the VPN establishment may fail in one or both directions. With internal Gateways, make sure to update the policies of both firewalls when there are (automatic or manual) changes in the Sites at either end.

What's Next?

- ▶ To turn automatic Site management off or on for an internal gateway, see Disabling/Re-Enabling Automatic VPN Site Management
- ▶ To continue with the automatic SIte (Internal Gateways) configuration, proceed as explained in Adjusting Automatic VPN Site Management (page 935).
- ▶ To define the IP addresses for an external security gateway element, see Defining Protected Networks for VPN Sites (page 936).

Related Tasks

- ▶ Adding a New VPN Site (page 935).
- ▶ Disabling a VPN Site Temporarily in All VPNs (page 938)

Disabling/Re-Enabling Automatic VPN Site Management

Automatic Site management is active by default on all new Internal Gateways. Automatic Site management copies the internal IP addresses from the routing view (all interfaces except those with the "Any Network" element attached) and continuously keeps the information up-to-date as you make changes to routing. If you prefer not to update the information automatically for any interface, you can disable this feature completely as instructed below. Alternatively, you can disable this feature just for selected interfaces as explained in Adjusting Automatic VPN Site Management (page 935).

▼ To disable/re-enable automatic Site management

1. In the internal Gateway element properties, switch to the **Sites** tab.
2. Deselect/select the **Include and Update Addresses Based on Routing** option.
   - When the option is not selected, you must manually define the addresses that you want to be routable through the VPN.
   - When the option is selected, the Site content updates automatically according to changes made in the corresponding firewall's Routing view (for those interfaces that are not disabled).

When you disable the automatic site management, the automatic Site is completely removed. There must be some other Site configured for the Gateway for it to be valid in a VPN.

What's Next?

▶ To define Sites manually, continue by Adding a New VPN Site (page 935).

Adjusting Automatic VPN Site Management

By default, Automatic Site Management is active for all new Internal Gateways. This adds a Site element for your Gateway that copies interfaces and networks from the Routing view to your VPN Gateway definition and updates them whenever the Routing view is changed. You can modify this automatic Site in the following ways:

- You can disable individual interfaces through their right-click menu. This way, you can exclude some of the internal interfaces from VPNs.
- You can add addresses to the automatic Site at the top level (at the same level with the Interface elements that hold the automatic content, not inside them) by dragging and dropping the correct Networks or other elements.
- You can add additional Sites for the Gateway alongside the automatic Site. See Adding a New VPN Site (page 935).
- You can mark the automatic Site as Private in some VPNs. See Adjusting VPN-Specific Site Settings (page 937).

What's Next?

- ▶ If you plan to translate the IP addresses of the local hosts that communicate through the VPN, add the addresses used in NAT in the Gateway's Site definition and mark the Sites that contain the original addresses as private (see Adjusting VPN-Specific Site Settings (page 937)).
- ▶ To continue configuring a new VPN without defining further settings for the Site, add all necessary Site and Gateway elements, then either configure a new set of IPsec settings as explained in Defining VPN Profiles (page 939) or configure the VPN with an existing set of settings as explained in Defining Policy-Based VPNs (page 949).

Adding a New VPN Site

You can add as many Site elements as you need to a Gateway. By default, each Site is included in all VPNs where the Gateway is used, but individual Sites can be manually disabled in any VPN without affecting the other VPNs. It is not possible to partially disable Sites; if the IP address space must vary between different VPNs, you need several Sites.

▼ To add a new Site element

1. Switch to the **Sites** tab in the Gateway element properties.
2. Right-click in the panel on the right and select **New→Site**. The Site Properties dialog opens.
3. Enter a **Name** and optionally a **Comment** for your reference.

What's Next?

- ▶ Continue the configuration in Defining Protected Networks for VPN Sites (page 936).

Defining Protected Networks for VPN Sites

The IP addresses configured for a Site define the addresses that are allowed to communicate through the VPN tunnel (the *encryption domain* of each gateway). If traffic in the tunnel is subject to NAT, you must add the NAT addresses to the Site. On Internal Gateways, you must add both the NAT addresses and any untranslated IP addresses that are not automatically added to the Site. External Gateway Sites only require the translated address space that the Internal Gateway actually contacts.

The local and remote Site definitions must match the same information on the other gateways involved in the VPN because the gateways verify this information. It may also make a difference whether addresses are entered as individual IP addresses, address ranges, or networks.

Note – Site definitions are applied globally to every VPN in which the Gateway is used unless you specifically adjust this in the VPN-specific Site settings.

▼ To add networks to a Site element

1. Check that you are on the **General** tab in the Site element’s properties.
2. Select the elements that represent the protected IP addresses behind the Gateway in the left panel and click **Add** to include them in this Site.
   - Addresses outside the Gateway’s local networks should generally not be included in the Site in most configurations. There is no need to include the Gateways’ own IP addresses in the Sites, but there is usually no need to specifically exclude those addresses if they happen to be in the networks you add to the Site.
   - IP address ranges may be interpreted differently from lists of IP addresses and networks depending on the VPN device. The system converts Group or Expression elements into address ranges, networks, or individual IP addresses depending on the IP addresses included. Other VPN devices may treat same types of values differently.

If you edited a previously configured VPN, make sure that the configuration of any external VPN gateway device involved contains the same IP address information and refresh the policy on all affected gateways to transfer the changes.

What’s Next?

▶ Continue the Site configuration by Adjusting VPN-Specific Site Settings (page 937).

Adjusting VPN-Specific Site Settings

Site elements allow you to adjust how the Site is used in each VPN.

▼ To adjust the VPN references of a Site

1. In the Site element properties, switch to the **VPN References** tab.
2. In the table, select/deselect the **Enable** option for the existing VPNs displayed to include/exclude the Site from the configuration. When the Site is disabled, it is grayed out.

Example **You can disable a Site that contains translated address in VPNs in which NAT is not used, or in which a different address space is used for translation.**

3. Select the **Mode** for the Site for each VPN in which it is enabled.
   - **Normal** mode is the default. Use this for all active Site elements that do not require one of the other two modes.
   - **Hub** mode is used on a hub Gateway in tunnel-to-tunnel forwarding. Hub mode Sites contain the IP addresses of the networks that are behind the remote spoke gateways (the networks between which the hub Gateway forwards traffic). The automatically generated Site cannot be used as a Hub Site.
   - *(Internal Gateways only)* **Private** mode is used for the local untranslated addresses when addresses are translated using NAT in the VPN. You must include the translated IP addresses (the addresses that the other end sees) as a Normal-mode Site element in these types of VPNs. If NAT is disabled in the VPN, any Sites in the Private mode are ignored.

What's Next?

▶ Add all necessary Site and Gateway elements, then proceed to the next relevant section below:

- Configure a new set of IPsec settings as explained in Defining VPN Profiles (page 939).
- Configure the VPN with existing settings as explained in Defining Policy-Based VPNs (page 949).
- Continue the configuration of the Route-Based VPN in Tunnel Mode as explained in Using the Route-Based VPN in Tunnel Mode (page 967).

Disabling a VPN Site Temporarily in All VPNs

These instructions are for disabling a Site that has been manually added to the Gateway. The Site is disabled globally in all VPNs. To disable the automatic Site for an Internal Gateway, see Disabling/Re-Enabling Automatic VPN Site Management (page 934).

▼ To disable a Site element globally

1. Open the Gateway element's properties and switch to the **Sites** tab.
2. Right-click the Site (the top-level element) and select **Disable**. The Site is grayed out to indicate that it is disabled in all VPNs in which this Gateway is used.

If you edited a previously configured VPN, refresh the policy on all affected gateways to transfer the changes. The configurations of external Gateways may also require an update.

What's Next?

- There must be at least one enabled site. To add a new Site, see Adding a New VPN Site (page 935).
- To continue without modifying the SItes, proceed as explained in the What's Next section in Defining Protected Networks for VPN Sites (page 936).

Related Tasks

- Adjusting VPN-Specific Site Settings (page 937)
- Defining VPN Topology for Policy-Based VPNs (page 951)

Removing a VPN Site Permanently from All VPNs

These instructions are for removing a Site that has been manually added to the Gateway. The Site is removed from all VPNs where the Gateway is used. To remove the automatic Site from an Internal Gateway, see Disabling/Re-Enabling Automatic VPN Site Management (page 934).

▼ To remove a Site element from a Gateway element

1. Open the properties of the Gateway element and switch to the **Sites** tab.
2. Right-click the Site and select **Remove**. The Site is removed from this Gateway and deleted from all VPNs.

If you edited a previously configured VPN, refresh the policy on all affected gateways to transfer the changes. The configurations of external Gateways may also require an update.

Defining VPN Profiles

Prerequisites: None

Each VPN refers to a VPN Profile. The VPN profile is the main point of configuration for IPsec VPN integrity checking, authentication, and encryption settings. The VPN Profile also contains some settings for VPN clients.

Several VPNs can use the same VPN Profile. There are predefined profiles in the system, which are mostly useful for VPNs between internal Gateways. Client-to-gateway VPNs generally require a custom profile.

What's Next?

- To create a new VPN profile, continue in Creating a New VPN Profile.
- To modify an existing custom-made profile, continue in Modifying an Existing VPN Profile (page 940).
- If you want to use an existing VPN Profile without editing it, continue in Defining Policy-Based VPNs (page 949) or Defining Route-Based VPN Tunnels (page 965).

Creating a New VPN Profile

▼ To create a new VPN Profile

1. Select **Configuration→Configuration→VPN**. The VPN Configuration view opens.

Illustration 54.7 VPN Profiles in the Element Tree

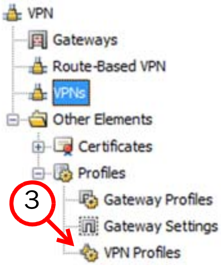

2. Expand the **Other Elements→Profiles** branch in the element tree.
3. Right-click **VPN Profiles** and select **New VPN Profile**. The VPN Profile Properties dialog opens.
4. Give the new profile a **Name** and optionally a **Comment**.

What's Next?

- Continue by Defining IKE SA Settings for a VPN (page 941).

Modifying an Existing VPN Profile

Before editing a VPN Profile that is used in active VPNs, we recommend backing up the settings (by duplicating the element, exporting it, or creating a Management Server backup). After editing a VPN profile that is used in active VPNs, check all VPNs that use the profile for any possible Issues that the changes may have caused.

Note – Only the IPsec SA settings of the VPN Profile are used in the Route-Based VPN. The other setting are ignored.

▼ To modify a VPN Profile

1. Select **Configuration→Configuration→VPN**. The VPN Configuration view opens.
2. Expand the **Other Elements→Profiles→VPN Profiles** branch in the element tree.
3. Double-click the profile you want to edit. The VPN Profile Properties dialog opens.
4. (*Optional*) Change the profile's **Name** and optional **Comment**. If you change the name, it is changed in all configurations without any need for further action.
5. Edit the other properties as instructed in the sections listed below.

What's Next?

- ▶ Defining IKE SA Settings for a VPN (page 941)
- ▶ Defining IPsec SA Settings for a VPN (page 943)
- ▶ Defining VPN Client Settings (page 946)
- ▶ Defining Trusted CAs for a VPN (page 948)

Defining IKE SA Settings for a VPN

The options you choose are a balance between performance and security. A higher level of security generally requires more processing power. If external gateways are involved, you must make sure that all settings match between the gateways.

▼ To define IKE SA Settings for a VPN

1. Switch to the **IKE SA** tab in the VPN Profile Properties dialog.

Illustration 54.8 VPN Profile Properties - IKE SA Tab

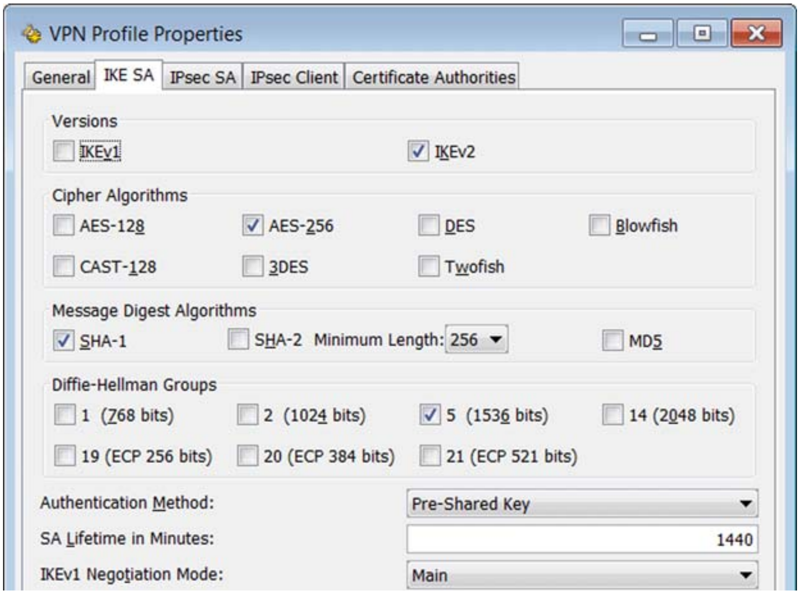

VPN Profile Properties

Versions

□ IKEv1 ☑ IKEv2

Cipher Algorithms

□ AES-128 ☑ AES-256 □ DES □ Blowfish

□ CAST-128 □ 3DES □ Twofish

Message Digest Algorithms

☑ SHA-1 □ SHA-2 Minimum Length: 256 □ MD5

Diffie-Hellman Groups

□ 1 (768 bits) □ 2 (1024 bits) ☑ 5 (1536 bits) □ 14 (2048 bits)

□ 19 (ECP 256 bits) □ 20 (ECP 384 bits) □ 21 (ECP 521 bits)

Authentication Method: Pre-Shared Key

SA Lifetime in Minutes: 1440

IKEv1 Negotiation Mode: Main

2. Select the **Version.**
   - You can select either **IKEv1** or **IKEv2** or both. If both versions are selected, IKEv2 is tried first in the negotiations, and IKEv1 is only used if the remote gateway does not support IKEv2.
3. Select the **Cipher Algorithm**(s) (encryption method) to use in the VPN. We recommend that you limit the selection to as few choices as possible, preferably only one. If you make multiple choices, multiple proposals are sent in IKE negotiations.[1]
   - Choose the method(s) according to your company’s requirements and the algorithms supported by the gateways. Consider the sensitivity of the transferred information and any regulations that you may have to follow.
   - Do not select the **DES** option unless you are required to do so. DES is no longer considered secure, since it is relatively easy to break DES encryption with modern computers.
   - **3DES** (Triple-DES) has a relatively high overhead compared to other protocols with a comparable level of security and is therefore not a good choice when high throughput is required.
4. Select the **Message Digest Algorithm**(s) (integrity checking method) to use in the VPN. We recommend that you select just one of these options if you have no specific reason to select more.

[1] The Russian product version has no strong encryption algorithms.

5. Select the **Diffie-Hellman Group** (for key exchange) to use in the VPN. We recommend that you select either **2 (1024 bits)** or **5 (1536 bits)**. Group 1 is not a secure option in all configurations.

6. Select the **Authentication Method**:
   - **DSS Signatures** and **RSS Signatures** use certificates for authentication and require that each Gateway has a valid certificate.
   - **Pre-Shared Key** requires that you periodically change the pre-shared keys for each tunnel in the VPN elements to be secure.

> **Note – The authentication method you select here is used for gateway-to-gateway VPNs. Client-to-gateway VPNs have separate settings on the IPsec Client tab.**

7. Adjust the **SA Lifetime** if you have a particular need to do so. The default lifetime is 1440 minutes.
   - If IKEv1 is selected as the Version, adjust the **SA Lifetime in Minutes** to match the settings of the external gateway device.
   - This setting affects tunnels that carry traffic continuously. Tunnels that are not used are closed after a short delay regardless of the lifetime set here.
   - If the VPN is continuously used and the lifetime passes, the IKE SA negotiations are done again.
   - Renegotiations improve security, but may require heavy processing.
   - There is a separate setting for IPsec Tunnel Lifetime on the IPsec SA tab. The IKE SA Lifetime must be longer than the IPsec Tunnel Lifetime.

8. (*Only if IKEv1 is selected as the Version*) Select the **IKEv1 Negotiation Mode**:

Mode	Description
Main	Main negotiation mode (*recommended*) protects the identity information of the Gateways so that malicious parties cannot gain information on the Gateway's identity by launching IKE negotiations with the gateway.
Aggressive	Aggressive negotiation mode skips some steps that are included in the main mode, resulting in somewhat quicker negotiations. For security reasons, we recommend that you do not use the aggressive negotiation mode if you use pre-shared keys for authentication. You must select Aggressive mode for VPNs that involve a gateway with a dynamic IP address. In this case, we recommend you use certificates for authentication rather than pre-shared keys

What's Next?

▶ Continue by Defining IPsec SA Settings for a VPN (page 943).

Defining IPsec SA Settings for a VPN

The options you choose are a balance between performance and security. A higher level of security generally requires more processing power. If external gateways are involved, you must make sure that all settings match between the gateways. For detailed descriptions of the Supported Cipher Algorithms, see the *Firewall Reference Guide*.

▼ To define IPsec SA Settings for a VPN

1. Switch to the **IPsec SA** tab in the VPN Profile Properties dialog.

Illustration 54.9 VPN Profile Properties - IPsec SA Tab

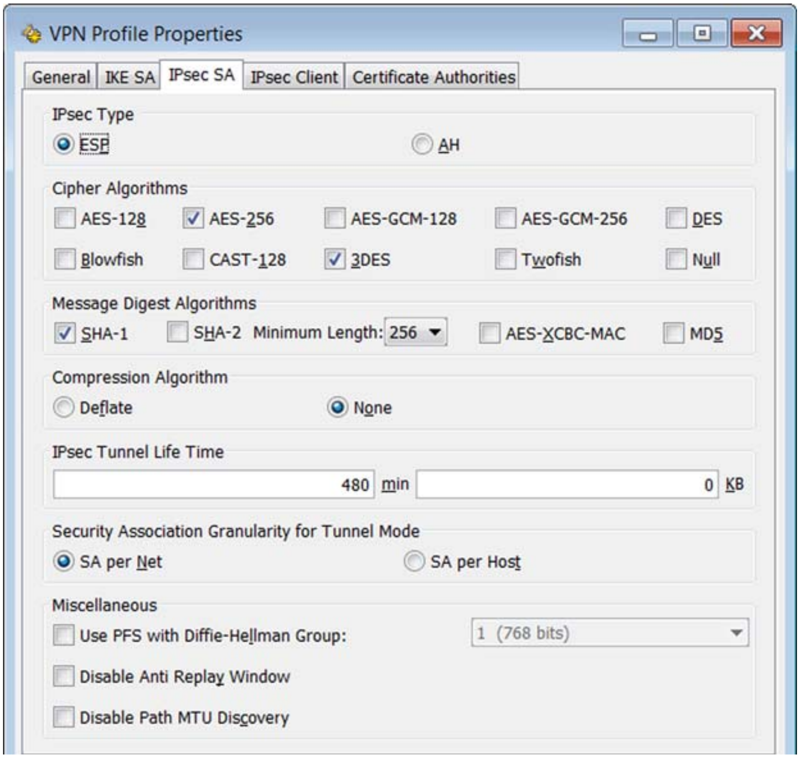

IPsec Type
● ESP ○ AH

Cipher Algorithms
□ AES-128 ☑ AES-256 □ AES-GCM-128 □ AES-GCM-256 □ DES
□ Blowfish □ CAST-128 ☑ 3DES □ Twofish □ Null

Message Digest Algorithms
☑ SHA-1 □ SHA-2 Minimum Length: 256 □ AES-XCBC-MAC □ MD5

Compression Algorithm
○ Deflate ● None

IPsec Tunnel Life Time
480 min 0 KB

Security Association Granularity for Tunnel Mode
● SA per Net ○ SA per Host

Miscellaneous
□ Use PFS with Diffie-Hellman Group: 1 (768 bits)
□ Disable Anti Replay Window
□ Disable Path MTU Discovery

2. Select the **IPsec Type**:

- The recommended setting is **ESP** (the communications are encrypted).
- In most cases, **AH** is not a valid option. The AH setting disables encryption for the VPN, fully exposing all traffic that uses the VPN to anyone who intercepts it in transit. You can use AH to authenticate and check the integrity of communications without encrypting them.

3. Select the **Cipher Algorithm**(s) (encryption method) to use in the VPN. We recommend that you limit the selection to as few choices as possible, preferably only one.[2]

- Choose the method(s) according to your company’s requirements and the algorithms supported by the gateways. Consider the sensitivity of the transferred information and any regulations that you may have to follow.
- Do not select the **Null** option unless you want to disable encryption. This option fully exposes all traffic that uses the VPN to anyone who intercepts it in transit. You can use

[2] The Russian product version has no strong encryption algorithms.

Null encryption to authenticate and check the integrity of communications without encrypting them.

- Do not select the **DES** option unless you are required to do so. DES is no longer considered secure, since it is relatively easy to break DES encryption with modern computers.
- **3DES** (Triple-DES) has a relatively high overhead compared to other protocols with a comparable level of security and is therefore not a good choice when high throughput is required.
- **AES-GCM-128** or **AES-GCM-256** are recommended for high speed networks.

4. Select the **Message Digest Algorithm**(s) (integrity checking method) to use in the VPN. We recommend that you select just one of these options if you have no specific reason to select more.

5. Select the **Compression Algorithm** to use in the VPN:
   - **Deflate** (*not supported in Firewall/VPN versions 4.2 and 4.3*) compresses the data to reduce the bandwidth use on congested links. This requires processing and memory resources, which increases latency. Latency may increase also for non-VPN traffic. Do not select this option if the resource utilization is very high to begin with. Gateways at both ends of each tunnel involved must support the option.
   - **None** (*recommended for most environments*) sends the data without compressing it. Provides better performance when bandwidth congestion for VPN traffic is not a constant issue or if there is significant processor load to begin with. The only supported option for Firewall/VPN versions 4.2 and 4.3.

6. (*Optional*) Adjust the **IPsec Tunnel Lifetime** if you have a particular need to do so. The default is 480 minutes with no limit on amounts of transferred data.
   - Reaching either the time or data amount limits triggers new IPsec SA negotiations, which must happen at regular intervals for the sake of security.
   - This setting affects tunnels that carry traffic continuously. Tunnels that are not used are closed after a short delay regardless of the lifetime set here.
   - IPsec SA negotiations are lighter on the processor than IKE SA negotiations, but still require some processing. Too frequent renegotiations can reduce performance down to unacceptable levels.
   - There is a separate setting for the SA Lifetime on the IKE SA tab. The SA Lifetime must be longer than the IPsec Tunnel Lifetime.

7. Select the **Security Association Granularity for Tunnel Mode** for VPN traffic. The IPsec SAs can be created for each communicating IP network or host.
   - **SA per Net** creates a security association (SA) for each network from which connections are made through the VPN. This setting reduces the overhead when there are a large number of hosts making connections through the VPN.
   - **SA per Host** creates a SA for each host that makes connections through the VPN. This setting may provide more even load balancing in clusters than the Per Net setting, but increases the overhead, as Per Host usually requires that more SAs are negotiated.

8. (*Optional*) Select **Use PFS with Diffie-Hellman Group** to use Perfect Forward Secrecy (PFS) and select the Diffie-Hellman Group you want to use with PFS.
   - When you use this option, the gateways calculate completely new values for key negotiations when renegotiating the SAs instead of deriving the values from previously negotiated keying material. This increases security if a key is compromised.
   - We recommend using either Group **2 (1024 bits)** or Group **5 (1536 bits)**.

9. (*Optional*) Select **Disable Anti-Replay Window** if you have a particular need to do so.
   - The anti-replay window feature provides protection against attacks in which packets are replayed. When the anti-replay window is enabled, the gateway keeps track of the

sequence numbers of the arriving packets, and discards any packet whose number matches the number of a packet that has already arrived.

- It is usually recommended to leave the anti-replay window enabled. However, if QoS is applied to ESP/AH traffic, some of the ESP packets (for the same SA) may be delayed due to the classification and arrive at the destination so late that the anti-replay window has moved too far. This causes the packets to be dropped. In this case, it may be necessary to disable the anti-replay window.

10. (*Optional*) Select **Disable Path MTU Discovery** to prevent the gateway from sending ICMP "Fragmentation needed" messages to the originator when the packet size (including the headers added for IPsec) exceeds the Ethernet-standard 1500 bytes.
    - If this option is selected, packets may be fragmented for transport across the VPN and reassembled at the receiving gateway.
    - Selecting the option may be necessary if ICMP messages do not reach the other gateway or the other gateway does not react to them correctly.

What's Next?

▶ If this Profile is used for client-to-gateway VPNs, continue in Defining VPN Client Settings (page 946).

▶ If you use certificates for authentication and you want to restrict trusted certificate authorities at the VPN level, continue in Defining Trusted CAs for a VPN (page 948).

▶ Otherwise, click **OK** and continue in Defining Policy-Based VPNs (page 949) or Defining Route-Based VPN Tunnels (page 965).

Defining VPN Client Settings

VPN client settings are selected on the IPsec Client tab in VPN Profiles. These options affect all IPsec VPN client connections and override any overlapping options on other tabs. If a VPN Profile that contains IPsec Client Settings is used in the Route-Based VPN, the IPsec Client settings are ignored.

▼ To define VPN client settings

1. Switch to the **IPsec Client** tab in the VPN Profile Properties dialog.

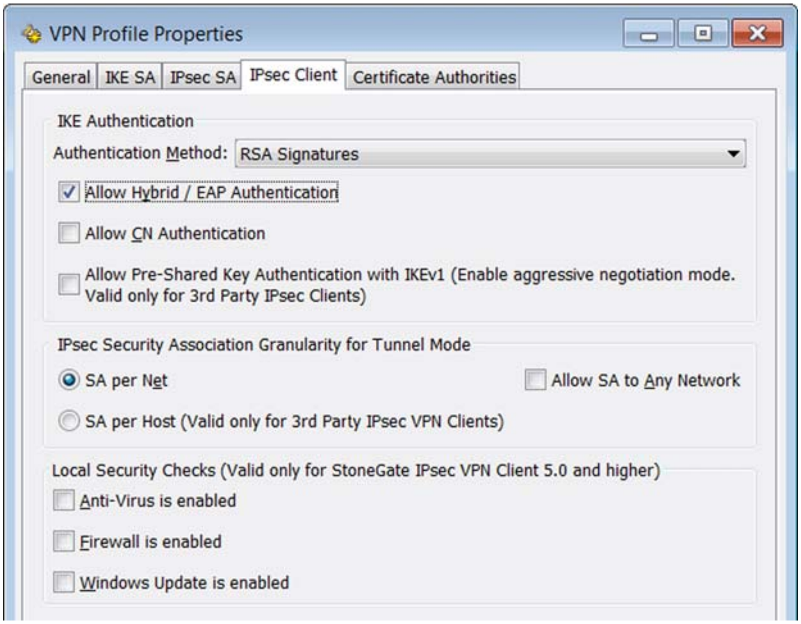

VPN Profile Properties

General | IKE SA | IPsec SA | IPsec Client | Certificate Authorities

IKE Authentication

Authentication Method: RSA Signatures

☑ Allow Hybrid / EAP Authentication

□ Allow CN Authentication

□ Allow Pre-Shared Key Authentication with IKEv1 (Enable aggressive negotiation mode. Valid only for 3rd Party IPsec Clients)

IPsec Security Association Granularity for Tunnel Mode

● SA per Net

□ Allow SA to Any Network

○ SA per Host (Valid only for 3rd Party IPsec VPN Clients)

Local Security Checks (Valid only for StoneGate IPsec VPN Client 5.0 and higher)

□ Anti-Virus is enabled

□ Firewall is enabled

□ Windows Update is enabled

2. Select the **Authentication Method** for certificate-based authentication.
   - This option is always used for the Gateway certificates for the Gateways involved in client-to-gateway VPNs, and if certificate authentication is used, also for the client.
   - Certificate authentication does not need separate activation. However, you must configure the issuing authority separately as trusted and you must create certificates for the VPN clients in a manual process.
3. (*Optional, Stonesoft IPsec VPN Clients only*) Select **Allow Hybrid/EAP Authentication** if you want to allow users of the Stonesoft IPsec VPN Client to authenticate by filling in a username/password combination or a similar authentication scheme provided by an external authentication server. The gateway still authenticates itself to the IPsec VPN Clients using a certificate.
4. (*Optional, certificate authentication only*) Select **Allow CN Authentication** to allow using the common name of the certificates for authentication. The CN is checked against a value entered in the User elements.
5. (*Optional*) Select **Allow Pre-Shared Key Authentication** if you have third-party VPN clients that use a pre-shared key for authenticating the VPN clients and the gateway. The pre-

shared key is defined at the VPN tunnel level in the VPN properties. Stonesoft IPsec VPN clients do not support this method.

Caution – The pre-shared key option requires aggressive mode IKE negotiations in the client-to-gateway VPN. In aggressive mode, user information is not protected, so we recommend you take precautions such as not using the same username for the users as they have when they access other services in your internal network.

6. Select the **IPsec Security Association Granularity for Tunnel Mode** for VPN traffic. The Stonesoft IPsec VPN Client supports only the **SA per Net** option.

Option	Description
SA per Net	Creates a security association (SA) for each network from which connections are made through the VPN. This setting reduces the overhead when there are a large number of hosts making connections through the VPN.
SA per Host	Creates a SA for each host that makes connections through the VPN. This setting may provide more even load balancing in clusters than the Per Net setting if there are many clients connecting from the same network, but increases the overhead significantly if there are connections from many IP addresses.
Allow SA to Any Network	Select this option together with **SA per Net** to support both Stonesoft IPsec VPN Clients and any third-party VPN clients that only support SAs negotiated per Host.

7. (*Optional*) Activate any combination of basic **Local Security Checks** for Stonesoft IPsec VPN Client version 5.0 and above.
   - Activating these options does not affect older versions of Stonesoft IPsec VPN Clients or third-party VPN clients; their connections are allowed as usual.
   - The selected types of external security software must be operational as reported by the Windows Security Center on the client computer; otherwise the connection attempt will fail. The check is performed after the user has successfully authenticated.
   - All client security checks are for on/off status of the external security software only; the checks do not include the update status of virus definitions or if any Windows updates have actually been installed.
   - If the security check fails, the IPsec VPN Client notifies the user on which checks have failed in a balloon message in the Windows Task Bar.

What's Next?

▶ If you use certificates for authentication and you want to restrict trusted certificate authorities at the VPN level, continue in Defining Trusted CAs for a VPN (page 948).

▶ Otherwise, click **OK** and continue in Defining Policy-Based VPNs (page 949).

Defining Trusted CAs for a VPN

If you want to use certificates signed by a particular CA (certificate authority), you must define the CA as an element in the system. By default, all VPN CAs defined in the system are considered trusted, but you can restrict the trusted CAs for particular VPNs as instructed here.

You can alternatively (or additionally) restrict trusted CAs also in the Gateway elements. If you restrict trusted CAs in both the Gateway and the VPN Profile, make sure that any two Gateways that form a VPN tunnel have a common CA to trust after all defined restrictions are applied (this is also automatically validated).

▼ To define trusted certificate authorities

1. Switch to the **Certificate Authorities** tab in the VPN Profile Properties dialog.
2. Select the **Trust Only Selected** option to restrict the trusted CAs and activate the controls below.
3. Select the CAs that you want to be considered as trusted in the VPNs that use this profile.
4. Click **OK**.

What's Next?

▶ Continue by Defining Policy-Based VPNs (page 949).

Defining Policy-Based VPNs

Prerequisites: Defining Security Gateways

Policy-based VPNs are defined using VPN elements. The VPN element collects together a set of other VPN-related elements to define settings for a particular VPN instance. The main configuration for the VPN consists of defining which Gateways are in the VPN and which of the Gateways form tunnels with each other. This is also where you can enter and renew pre-shared keys if you use them for authentication in this VPN.

Consider the following when creating new VPN elements:

- Check whether you can utilize an existing VPN element instead. Most settings can be set individually for each gateway-to-gateway pair even within a single VPN. The VPN Profile, pre-shared key, and Multi-Link settings can all be selected separately for each VPN tunnel. Site definitions are the only major exception to this rule; each Gateway's Sites are fixed within each VPN element.
- There must not be duplicate tunnels (two tunnels between the same two end-points) in the configuration of any Firewall/VPN engine. Duplicate tunnels cause a policy installation failure. The easiest way to avoid duplicate tunnels is to define all VPNs between your internal gateways in the same VPN element.
- If you are creating VPNs with partner organizations, you may only want to include a subset of the internal IP address space in the VPN definitions to avoid having to reveal all internal addresses to your partner. Any cases where Site definitions must be different for different VPN tunnels requires creating separate VPN elements.

When you configure the VPN element, the validity of the VPN is automatically checked. If problems are found, they are shown in the **Issues** view. While useful in many cases, the automatic check does not detect all problems, especially regarding external gateways or interference between several separate VPN elements.

What's Next?

- ▶ To add a new VPN to the system, see Creating a New VPN Element (page 950).
- ▶ To modify the contents of an existing predefined or custom VPN element, see Modifying an Existing VPN Element (page 951).

Creating a New VPN Element

The configuration of a new VPN element has two stages: first you define some basic properties for the element as explained here, and then you can proceed to populating the element with Gateways and adjusting the tunnels in the VPN editing view.

▼ To create a new VPN element

1. Select **Configuration→Configuration→VPN**. The VPN Configuration view opens.
2. Right-click the **VPNs** branch in the tree and select **New VPN**. The VPN Properties dialog opens.

Illustration 54.10 VPN Properties

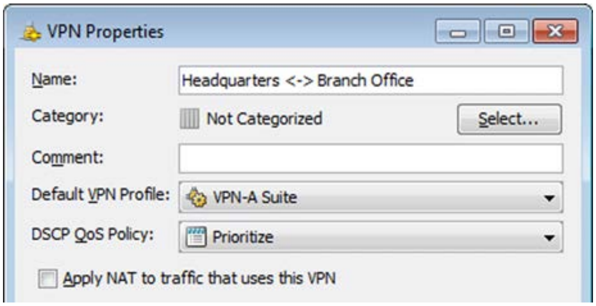

VPN Properties

Name: Headquarters <-> Branch Office

Category: Not Categorized

Comment:

Default VPN Profile: VPN-A Suite

DSCP QoS Policy: Prioritize

□ Apply NAT to traffic that uses this VPN

3. Enter the **Name** and optionally a **Comment** for the VPN element.
4. Select the **Default VPN Profile** for the VPN. By default, this profile is used for all tunnels, but you can override the selection for individual tunnels.
5. (*Optional*) Select the **DSCP QoS Policy** that defines how DSCP matching and/or marking is done for VPN traffic in one of the following ways:
   - Select an existing QoS Policy from the list.
   - Select **Other** and select an existing QoS Policy or click the New icon to create a new QoS Policy as explained in Defining QoS Policies (page 785).
6. (*Optional*) Select **Apply NAT Rules to Traffic That Uses This VPN** if you want the NAT rules in the FIrewall/VPN engine's policy to apply to traffic that it sends into or receives from the VPN, or if you want to use the NAT Pool feature to translate VPN client connections.
   - The option affects the traffic that is transported inside the tunnels.
   - This option does not affect the tunnel negotiations or the encrypted packets between gateways. These communications are always matched to NAT rules.
7. Click **OK**. The VPN opens in the editing view.

What's Next?

▶ Continue by Defining VPN Topology for Policy-Based VPNs (page 951).

Modifying an Existing VPN Element

The VPN element can be configured in two ways: the basic properties for the element including the VPN Profile and a NAT option are defined in the VPN element's properties. All other settings, including the included Gateways, Sites and tunnels are configured in the VPN editing view.

▼ To modify an existing VPN element

1. Select **Configuration→Configuration→VPN**. The VPN Configuration view opens.
2. Expand the **VPNs** branch in the element tree.
3. Open the correct view for the settings you want to edit:
   - To edit the basic properties, right-click the VPN element and select **Properties**. See the explanations in Creating a New VPN Element (page 950) if you need help with the options. Note that you will need to open the editing manually (see below) after changing the basic settings if necessary - the editing view does not open automatically in this case.
   - To adjust the other settings, right-click the VPN element and select **Edit**. Continue the configuration in Defining VPN Topology for Policy-Based VPNs (page 951) or Defining VPN Tunnel Settings for Policy-Based VPNs (page 953) according to which settings you want to change.

Defining VPN Topology for Policy-Based VPNs

The VPN element editing view has two tabs. The Gateway selection on the **Overall Topology** tab determines which Gateways are included in the VPN, which of them form tunnels with each other, and which gateways contact each other through a hub gateway instead of contacting each other directly. The general VPN topology is defined by classifying Gateways as *Central* or *Satellite*. This classification defines which tunnels are generated on the Tunnels tab, on which you can then further disable any unnecessary tunnels that may be generated.

For a valid VPN, you must have at least two Gateways included in the VPN and at least one of the Gateways must be listed as a Central Gateway. The Satellite Gateways panel can be left empty (for a full-mesh topology).

IPv4 Access rules control which connections use the VPN tunnels. Always remember to check the Access rules after you add or remove tunnels.

Note – Each end-point-to-end-point tunnel can only exist in one active VPN. If you use the same two Gateway elements in more than one VPN, make sure the topology does not create duplicate tunnels and/or disable any duplicates of existing tunnels on the Tunnels tab.

▼ To define the VPN topology

1. Check that you have the **Overall Topology** tab active in the VPN element in the editing mode (see Modifying an Existing VPN Element (page 951) for instructions on how to open the VPN in editing mode).

Illustration 54.11 VPN Editing View - Overall Topology

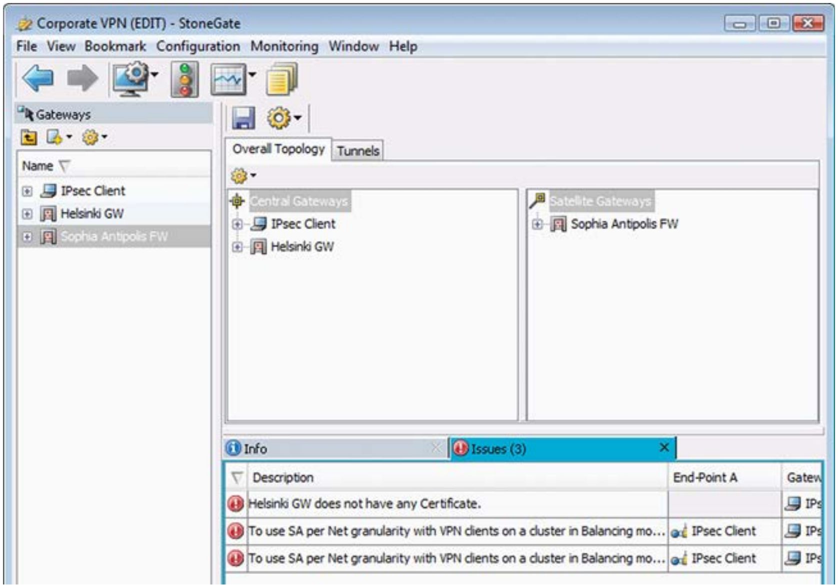

Description	End-Point A	Gatew
Helsinki GW does not have any Certificate.		IPs
To use SA per Net granularity with VPN clients on a cluster in Balancing mo...	IPsec Client	IPs
To use SA per Net granularity with VPN clients on a cluster in Balancing mo...	IPsec Client	IPs

2. Drag and drop the Gateways you want to include in this VPN into either of the two panels for the VPN topology.
   - If you add a Gateway under **Central Gateways**, the Gateway can establish a VPN with any other Gateway in the VPN and the **Tunnels** tab is populated with tunnels between the endpoints of the Gateway you add and the endpoints of all other Gateways in the VPN.
   - If you add a Gateway under **Satellite Gateways**, the Gateway can establish a VPN only with Central Gateways in this VPN and the Tunnels tab is populated with tunnels between the endpoints of the Gateway you add and the endpoints of the Central Gateway(s).
   - The **Issues** panel alerts you to any incompatible or missing settings you must correct.

Note – Be careful to not unintentionally drop Gateways on top of other Gateways. This indicates a forwarding relationship on a hub gateway (see the next step below).

3. (*Optional*) If you want to forward some connections from one VPN tunnel into another through a hub gateway, drag and drop a Gateway on top of another Gateway. The gateway is added under the other gateway at the same level as the Sites.
   - Full support for this feature requires that the hub gateway is running Firewall/VPN version 5.0 or higher.
   - See Configuration 4: Basic VPN Hub (page 910) for a VPN hub configuration example.
   - The Gateway used as a hub requires a special Site configuration, see Defining Protected Networks for VPN Sites (page 936) and Adjusting VPN-Specific Site Settings (page 937).

Example **In a setup in which VPN client users can connect to networks behind both Gateway A and Gateway B when they connect to Gateway A, you would drop Gateway B on top of Gateway A in the Central Gateways list and add the Client Gateway to Satellite Gateways.**

4. (*Optional*) To allow VPN client access within this VPN, drag and drop the default **IPsec Client** Gateway element to **Satellite Gateways** or **Central Gateways**.
5. (*Optional*) If you want to exclude a Gateway's Site (some IP addresses) from this VPN, right-click the Site element under the Gateway and select **Disable**.

What's Next?

▶ Continue by Defining VPN Tunnel Settings for Policy-Based VPNs.

Defining VPN Tunnel Settings for Policy-Based VPNs

The Tunnels tab in the in the VPN element editing view allows you to define settings particular to individual tunnels or disable some tunnels altogether. The topology of the VPN (defined on the Overall Topology tab) determines which tunnels are shown on the Tunnels tab.

If an Internal Gateway has a Multi-Link VPN configuration, you can select whether to use tunnels as backups or actively balance traffic between them. Multi-Link is specific to Stonesoft Firewall/VPN, and is not part of the IPsec standard. You may not be able to use Multi-Link with third-party gateways. Satisfactory results can be achieved if the third-party gateway allows ICMP probes, RTT ICMP probes, and supports DPD. You can disable redundant tunnels to the third-party gateway on this Tunnels tab if required.

This is also where you can view the *link summary* (a summary of addresses and settings that have been configured for individual tunnels), which you may want to check especially when there are complex setups involving external components (such as a VPN hub configuration).

Before modifying a VPN element that is used in active VPNs, we recommend making a backup of the Management Server as instructed in Creating Backups (page 1009).

▼ To define VPN tunnel settings

1. Switch to the **Tunnels** tab in a VPN element in the editing mode (see Modifying an Existing VPN Element (page 951) for instructions on how to open the VPN in editing mode). The list of tunnels is displayed.

Illustration 54.12 VPN Editing View - Tunnels

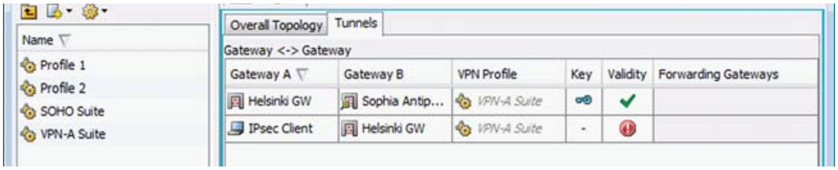

Gateway A	Gateway B	VPN Profile	Key	Validity	Forwarding Gateways
Helsinki GW	Sophia Antip...	VPN-A Suite	🔑	✓	
IPsec Client	Helsinki GW	VPN-A Suite	-	!	

- If no tunnels are listed, see Defining VPN Topology for Policy-Based VPNs (page 951).
- The **Gateway<->Gateway** list shows connections between pairs of gateways.
- The **End-Point<->End-Point** list shows the individual connections that form the tunnels in the **Gateway<->Gateway** list. There can be several connections at this level for any Gateway pair if one or both of the Gateways have multiple endpoints (Multi-Link). If both Gateways have only one endpoint, there is only one tunnel also at this level for the Gateway pair.
- If you have set up connection forwarding between the Gateways on the **Overall Topology** tab, the number of generated tunnels is reduced according to the relationships configured and the capabilities of the Gateway that forwards the traffic. The forwarding relationships are shown under **Forwarding Gateways**.

2. (*Optional*) If there are tunnels listed that are not needed, right-click the tunnel and select **Disable**.
   - Duplicate tunnels are not allowed between VPNs. If some other VPN already defines a tunnel between the same end-points as some tunnel in this VPN, you must disable the duplicate tunnel in one of the VPNs.
3. If you use pre-shared keys for authentication with external gateways, either set the key agreed with your partner or export the keys that have been automatically generated for your partner to use.
   - To view, change, or export the pre-shared key for a particular tunnel, double-click the key icon in the **Key** column in the **Gateway<->Gateway** list.
   - This pre-shared key is used only with gateway devices. Set pre-shared keys for third-party VPN clients in the User elements (Stonesoft IPsec VPN Clients do not allow pre-shared key authentication).

Caution – The pre-shared key must be long and random to provide a secure VPN. Change the pre-shared key periodically (for example, monthly). Make sure that it is not possible for outsiders to obtain the key while you transfer it to other devices.

4. (*Optional*) Change the **VPN Profile** used at the tunnel level to override the profile selected for the VPN element:
   - If you change a profile for a tunnel on the **Gateway<->Gateway** list, both IKE SA and IPsec SA settings are overridden from what is default for the VPN.
   - If you change a profile for a tunnel on the **End-Point<->End-Point** list, only the IPsec SA settings are overridden from what is selected for the main tunnel on the Gateway level.

5. (*Optional*) If you have multiple tunnels (network links) between two Gateways (Multi-Link configuration), you can select the Mode in which **End-Point<->End-Point** links are used. The Mode that you select for a link overrides the Mode setting in the End-Point properties.
   - Select a tunnel on the **Gateway<->Gateway** list.
   - Right-click the **Mode** column for a link on the **End-Point<->End-Point** list and select the mode from the right-click menu.

Option	Description
Active	The links are used at all times. If there are multiple links in Active mode between the Gateways, the VPN traffic is load-balanced between the active links based on a performance measurement or on the links' relative bandwidths. This means that VPN traffic is directed to the link that has the lowest load.
Aggregate	The links are used at all times. If there are multiple links in Aggregate mode, each VPN connection is load-balanced on a packet-by-packet basis between the links in round robin fashion. For example, if there are two links in Aggregate mode, individual packets of a new VPN connection are directed to both links. Use of this Aggregate mode is likely to cause packet reordering.
Standby	The links are used only when all Active or Aggregate links are unusable.

Note – The Active and Aggregate modes are mutually exclusive. All the VPN end-points and the links between two gateways must be either in Active and Standby modes or in Aggregate and Standby modes.

   - The Mode you select directly in the link's right-click menu is used for all traffic that is directed to the link. You can also define that the link's Mode is automatically calculated based on the Mode defined for the end-points. In addition, you can define QoS Exceptions to specify that the link's Mode depends on the QoS class of the traffic that is directed to the link. See Editing VPN Link Modes in Policy-Based VPNs (page 956).

6. (*Optional*) Review the IP addresses and settings used in the individual tunnels by right-clicking the tunnels on the **End-Point<->End-Point** list and selecting **View Link Summary**. This is especially useful in complex configurations that involve external components to check the IP address details and other settings that must match with the external configuration.

7. After making all changes, check the **Validity** column for all tunnels.
   - If a tunnel has a warning icon in the **Validity** column, right-click the tunnel and select **View Issues**. You must resolve all problems indicated in the messages shown.
   - If all tunnels are shown as valid, the VPN is correctly configured, although the Management Server cannot check all possible problems at this point, so additional issues can be shown at policy installation. Any validation and issues that are shown for external gateways are based only on the definitions that have been entered manually into the related elements.

8. Click the **Save** icon above the tunnel lists.

The VPN is now configured, but to direct outgoing traffic to the VPN and allow incoming traffic from the VPN, you must add VPN Access rules and possibly also NAT rules.

What's Next?

- If you need to add a trusted certificate authority or certificates that are not generated automatically, proceed to Getting Started With VPN Certificates (page 972) before adding VPN rules.
- Otherwise, continue by Creating Rules for Policy-Based VPNs (page 958).

Editing VPN Link Modes in Policy-Based VPNs

Prerequisites:

The Mode of a VPN link determines how the link is used for VPN traffic. You can select the Mode(s) in which **End-Point<->End-Point** tunnels are used if there are multiple links between two Gateways (Multi-Link configuration). The Mode you select is the default mode for the link.

To edit link modes

1. On the Tunnels tab in the VPN view, select a tunnel on the **Gateway<->Gateway** list. The links between the gateways are displayed in the **End-Point<->End-Point** list.
2. Right-click the **Mode** column for a link on the **End-Point<->End-Point** list and select **Edit Mode**. The Link Mode Properties dialog opens.

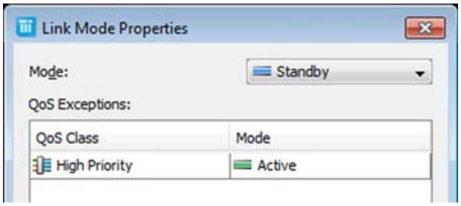

Mode: Standby

QoS Class	Mode
High Priority	Active

3. Select the **Mode** for the link.

Table 54.2 Link Modes

Option	Description
<option (*default*)>	The mode is automatically calculated based on the Mode selected for the end-points. If the end-points' Mode changes, the link's Mode is automatically updated. The (*default*) mode is calculated in the following way: - If both end-points are in the Active mode, the link's Mode is Active. - If both end-points are in the Aggregate mode, the link's Mode is Aggregate. - If one of the end-points is in the Standby mode, the link's mode is Standby.
Active	The link is used at all times. If there are multiple links in Active mode between the Gateways, the VPN traffic is load-balanced between the links based on the links' load. This means that VPN traffic is directed to the link that has the lowest load.

Table 54.2 Link Modes (Continued)

Option	Description
Aggregate	The link is used at all times and each VPN connection is load-balanced in round-robin fashion between all the links that are in the Aggregate mode. For example, if there are two links in Aggregate mode, a new VPN connection is directed to both links.
Standby	The link is used only when all Active or Aggregate mode links are unusable.

- The link's Mode setting overrides the Mode defined in the end-point properties.

Note – The Active and the Aggregate modes are mutually exclusive. All the VPN end-points and the links between two gateways must be either in Active and Standby modes or in Aggregate and Standby modes.

4. (*Optional*) Add a **QoS Exception** that defines the tunnel's Mode for VPN traffic that has been assigned a particular QoS class.
   - Click **Add** and select the QoS Class in the dialog that opens or create a new QoS Class through the New icon at the top of the dialog. A new row is added to the QoS Exceptions table.
   - Click the **Mode** column to select the Mode for the QoS Class (the Modes are explained in the table above).
   - For more information on QoS Classes, see Creating QoS Classes (page 784).

Note – Each QoS Exception definition is link-specific. If you want to direct traffic that has a particular QoS Class to more than one link in a Multi-Link VPN configuration, you must define a QoS Exception for each link through which you want to allow the traffic.

5. (*Optional*) Repeat step 4 to add more QoS Exceptions.
6. Click **OK**.

Creating Rules for Policy-Based VPNs

The firewall IPv4 Access rules define which traffic is sent to the policy-based VPN and which traffic is allowed out of the policy-based VPN. These checks are made in addition to the enforcement of the Site definitions of the Gateways, which define the allowed source and destination addresses for each VPN.

There are three options for the **Use IPsec VPN** rule action, which all behave identically for connections that originate in the local protected network, but which each have a special meaning for connections coming in from external sources.

- **Apply VPN** and **Enforce VPN** both direct traffic from protected local networks into the VPN tunnel and allow traffic that arrives through a VPN. However:
    - **Enforce VPN** drops any non-VPN traffic from external networks to the local network if it matches the rule.
    - **Apply VPN** does not match non-VPN traffic from outside networks into the protected networks; matching continues from the next rule.
- **Forward**: Directs traffic from protected local networks or from a VPN tunnel into a VPN tunnel. Useful for forwarding connections from one VPN tunnel into another (VPN hub configuration) or connections from local networks to currently connected VPN client computers.

There is also a matching cell for **Source VPN** in firewall IPv4 Access rules. The cell can be used to match traffic based on whether the traffic is coming from a VPN tunnel. When the Source VPN cell is set to match VPNs, the rule only matches traffic from the selected VPNs. The cell can also be set to only match non-VPN traffic. Access rules that do not have any Source VPN definition can match any traffic, including traffic that is received through a VPN. However, the Access rules must contain at least one rule that refers to the VPN (in the Action or Source VPN cell) for any of the VPN's settings to be included in the engine's local configuration.

Note – We recommend that you activate logging for the VPN rules for initial testing even if you do not plan to log the connections that use the VPN later on. VPN negotiations between the gateways are always logged.

This subject is covered in the following topics:

- Creating Rules for Gateway Connections in Policy-Based VPNs (page 959)
- Creating Rules for VPN Client Connections in Policy-Based VPNs (page 960)
- Creating Forwarding Rules on Hub Gateways for Policy-Based VPNs (page 961)
- Preventing Other Access Rules from Matching Policy-Based VPN Traffic (page 963)
- Creating NAT Rules for Policy-Based VPN Traffic (page 963)

Creating Rules for Gateway Connections in Policy-Based VPNs

In addition to the traffic that flows through the tunnels, the VPN traffic to form and maintain the tunnels must be allowed in the policy. The Firewall Template policy allows this traffic, but if you use a custom top-level template, make sure that at least the **ISAKMP (UDP)** Service is allowed between the Gateways (further ports may need to be opened if encapsulation is used).

▼ To create rules for incoming gateway-to-gateway VPN traffic

1. Insert the following type of rule to allow traffic from a single VPN with an Apply or Enforce action:

Table 54.3 Basic Rule for Allowing Incoming VPN Traffic from a Single VPN

Source	Destination	Service	Action
Remote networks.	Local networks.	Set as needed.	Select **Use IPsec VPN**, then change the **Action** to **Apply** or **Enforce** and add the VPN element.

2. (*Optional*) Define the following types of rules when you want to match the rule based on whether traffic is using a VPN:

Table 54.4 Rule for Allowing Incoming VPN Traffic from Any Number of Different VPNs

Source	Destination	Service	Action	Source VPN
Remote networks.	Local networks.	Set as needed.	Select **Allow**.	(Double-click the cell to edit.) To ignore this rule when processing non-VPN traffic, activate the **Match traffic based on source VPN** option and add one or more specific VPN elements according to where the traffic is coming from. This rule does not match traffic from other sources.

▼ To create rules for outgoing VPN traffic

↳ Insert the following types of rule:

Table 54.5 Basic Rule for Outgoing VPN Traffic

Source	Destination	Service	Action
Local networks.	Remote networks.	Set as needed.	Select **Use IPsec VPN** and add the VPN element (the Apply, Enforce, and Forward actions are all identical in this use).

Note – If Access rules send traffic into a VPN, but the source and/or destination IP address(es) are not included in the Site definitions, the traffic is dropped. This configuration error is shown as the message “tunnel selection failed” in the logs.

Creating Rules for VPN Client Connections in Policy-Based VPNs

In addition to the traffic that flows through the tunnels, the VPN traffic to form and maintain the tunnels must be allowed in the policy. The Firewall Template policy allows this traffic, but if you use a custom top-level template, make sure that at least the **ISAKMP (UDP)** Service is allowed (further ports may need to be opened if encapsulation is used) from addresses that VPN client hosts use when contacting the Gateway (ANY or some group of specific known addresses). VPN client user authentication is also allowed as part of this VPN connection establishment process.

▼ To create rules for incoming VPN client traffic

1. Insert the following type of rule to allow DHCP queries from VPN clients to a DHCP server (mandatory for the Virtual Adapter feature):

Table 54.6 Rule for Allowing DHCP Queries from VPN Clients

Source	Destination	Service	Action	Source VPN
ANY	DHCP server.	BOOTPC (UDP) BOOTPS (UDP)	Select **Jump** and select the default **DHCP Relay** sub-policy.	(Double-click the cell to edit.) Activate the **Match traffic based on source VPN** option. Select **Rule matches traffic from any VPN client** from the list.

Note – The Virtual Adapter IP addresses must be assigned by an external DHCP server if you use a Firewall Cluster as a VPN gateway. If the VPN gateway is a Single Firewall you can alternatively use the internal DHCP server on the Single Firewall to assign IP addresses to Virtual Adapters.

2. Insert the following type(s) of rule(s) to allow incoming connections from VPNs with VPN clients.

Table 54.7 Rule for Allowing Incoming Traffic from VPN Clients

Source	Destination	Service	Action	Authentication
VPN clients' Virtual Adapter address space or NAT Pool address space.	Local networks.	Set as needed.	Select **Use IPsec VPN**, then change the **Action** to **Apply** or **Enforce** and add the VPN element.	(Double-click the cell to edit.) Add User or User Group elements and allowed Authentication Method(s).

- When a specific VPN and specific Authentication Method(s) are used somewhere in the installed policy (such as in a rule similar to the one above), the corresponding configurations are activated on the firewall. But note that all other rules are also matched to the VPN client user's connections.
- Any users known to the system who can authenticate using the specified authentication method can connect with a VPN client. Any such connected users can then access resources if there is a matching rule that allows connections without specific Users

defined. You can also use the **Source VPN** cell to prevent unwanted matches in Access rules.

- When filled in, User and Authentication cells are equal to Source, Destination, and Service as rule matching criteria; matching continues from the next rule if the defined User and Authentication Method do not match the connection that is being examined. You can, for example, define that the same user has access to different resources depending on the authentication method used.

3. (*Optional*) To allow internal hosts to open connections to the VPN client computers when the VPN is active (for example, to allow administrators' remote desktop connections to the hosts):

Table 54.8 Rule for Sending Traffic to VPN Clients

Source	Destination	Service	Action
Local networks.	VPN clients' Virtual Adapter address space or NAT Pool address space.	Set as needed.	Select **Use IPsec VPN**. Switch the **Action** to **Forward**, then add a specific VPN element or select the **$ Client-to-gateway IPsec VPNs** option to match any VPN client connection.

- To use the VPN, the connecting hosts' IP addresses must be included in the gateway's Site definition. See Defining Sites for VPN Gateways (page 933).
- Remember to also add a NAT rule for this traffic direction if NAT is used inside this VPN's tunnels. See Creating NAT Rules for Policy-Based VPN Traffic (page 963).

Creating Forwarding Rules on Hub Gateways for Policy-Based VPNs

For a configuration example of creating a VPN hub, see Configuration 4: Basic VPN Hub (page 910).

▼ To create rules for forwarding VPN traffic from one tunnel to another

➥ Insert the following type of rule:

Table 54.9 Basic Rule for Forwarding VPN Traffic

Source	Destination	Service	Action
One/some/all addresses in remote (spoke) networks as needed.	One/some/all addresses in remote (spoke) networks as needed.	Set as needed.	Select **Use IPsec VPN**, then change the **Action** to **Forward** and add the VPN element for the VPN into which matching traffic is forwarded.

▼ To create rules for forwarding tunneled traffic to the Internet

1. Insert the following type of rule:

Table 54.10 Rule for Allowing Traffic Except if it Arrives through VPNs

Source	Destination	Service	Action	Source VPN
Remote networks or VPN clients' Virtual Adapter address space/NAT Pool address space.	ANY or an Expression element that excludes the local internal networks.	Set as needed.	Select **Allow**.	(Double-click the cell to edit.) Activate the **Match traffic based on source VPN** option and add one or more specific VPN elements according to where the traffic is coming from. This rule does not match traffic from other sources.

2. In most cases, the source addresses are from a private space and you must add a NAT rule to translate them to publicly routable IP addresses. Make sure NAT is enabled in the Properties dialog of the VPN element(s) and add a NAT rule for the VPN traffic if a suitable NAT rule does not exist already.

3. Configure the remote Gateway or VPN clients to forward all traffic to the VPN:
   - VPN clients: configure the DHCP server to assign the hub Gateway as the default gateway for the VPN clients.
   - Stonesoft Firewall/VPN gateways: create a VPN rule that directs all traffic to the VPN with the hub gateway, see Creating Rules for Gateway Connections in Policy-Based VPNs (page 959).

Note – For the traffic to be allowed into the VPN, the destination IP address must be part of the Site definition of the hub Gateway. When you forward Internet traffic, the hub's Site must usually include the Any Network element. This Site may interfere with other VPN configurations, so you should generally disable it in other VPNs. See Defining Protected Networks for VPN Sites (page 936).

Preventing Other Access Rules from Matching Policy-Based VPN Traffic

▼ To create rules that do not match connections opened through a VPN

➥ Insert the following type of rule:

Table 54.11 Rule for Allowing Traffic Except if it Arrives through VPNs

Source	Destination	Service	Action	Source VPN
Set as needed.	Set as needed.	Set as needed.	Set as needed.	(Double-click the cell to change settings.) Activate the **Match traffic based on source VPN** option. Select **Rule does not match traffic from any VPN** from the list.

Creating NAT Rules for Policy-Based VPN Traffic

NAT rules are always applied to the VPN tunnel communications (the encrypted communications that have the Gateways as their source and destination), but not to the traffic that is inside the VPN tunnel (the traffic that uses the VPN tunnel). If you want to apply NAT to traffic going into or coming out of the VPN tunnel, you must specifically allow this with an option in the VPN element's basic properties (see Modifying an Existing VPN Element (page 951)).

Observe the following guidelines:

- You must define Sites (encryption domains) that contain the translated IP addresses that the packets use when they are inside the VPN tunnel. Set the Sites that contain the real IP addresses to **Private** mode in the VPN. See Adjusting VPN-Specific Site Settings (page 937).

Example **If you translate IP addresses of traffic going into the VPN, you must add a Site that includes the translated IP addresses to your internal Gateway to make those addresses valid in the VPN. The Sites that contain the internal addresses are set to Private mode.**

- If you have configured address translation for VPN clients in the Firewall element's properties (NAT Pool), the NAT Pool translation is applied before the NAT rules are applied. NAT rules cannot match traffic to which NAT pool translation is applied. NAT Pool is the preferred method for translating VPN client addresses.
- If you want to forward traffic originating from VPN clients to the Internet, you must typically have at least two NAT rules: the first rule to translate internal IP addresses to an external IP address for the Internet connections and the second rule (above the first one) for connections to internal resources to prevent NAT from being applied or to translate to an internal IP address as necessary.

The order of processing for traffic going into a VPN tunnel is:
Access Rules→NAT Rules→VPN tunnel.

The order of processing for traffic coming out of a VPN tunnels is:
Access Rules→(VPN client NAT Pool)→NAT Rules→Internal Network.

Other than what is stated above, there are no further VPN-specific issues with NAT rules. As usual, the first matching NAT rule is applied to those connections that are matched against the NAT rules and the rest of the NAT rules are ignored. For detailed instructions on creating NAT rules, see Editing Firewall NAT Rules (page 687).

Editing the Route-Based VPN

Prerequisites: None

The Route-Based VPN element is a predefined element. There can only be one Route-Based VPN element on each Management Server. Rather than creating a new Route-Based VPN element for each tunnel, you edit the default Route-Based VPN element to define all of the tunnels in the Route-Based VPN.

What's Next?

- To define the default encryption settings for new Route-Based VPN tunnels, proceed to Selecting the Default Encryption for the Route-Based VPN.
- To use the Route-Based VPN on its own, proceed to Defining Route-Based VPN Tunnels (page 965).
- To use the a policy-based VPN to encrypt Route-Based VPN tunnels, proceed to Using the Route-Based VPN in Tunnel Mode (page 967).

Selecting the Default Encryption for the Route-Based VPN

The Default Encryption setting in the Route-Based VPN properties defines the encryption mode that is selected by default when you create new Route-Based VPN tunnels. You can change the Default Encryption setting for individual tunnels as explained in Defining Route-Based VPN Tunnels (page 965).

▼ To select the default encryption for the Route-Based VPN

1. Select **Configuration→Configuration→VPN**. The VPN Configuration view opens.
2. Browse to **Route-Based VPN**. The Route-Based VPN tree opens in the right panel.
3. Right-click **Route-Based VPN** in the Route-Based VPN tree and select **Properties**. The Route-Based VPN Properties dialog opens.
4. Select the Default Encryption:

Encryption Mode	Description
Transport Mode	Sets the Encryption cell of new Route-Based VPN tunnels to Transport Mode and uses the encryption settings defined in the selected VPN Profile.
Tunnel Mode	Sets the Encryption cell of new Route-Based VPN tunnels to Tunnel Mode and uses the encryption settings defined in the selected policy-based VPN.
No Encryption	Sets the Encryption cell of new Route-Based VPN tunnels to No Encryption. **Caution!** This option defines a tunnel in which traffic is encapsulated but not encrypted. It is provided to allow encapsulation of traffic when the traffic does not need to be secured by a VPN. The No Encryption option is recommended only for the following uses: - creating unencrypted tunnels entirely within protected networks - testing and troubleshooting routing and connectivity without using VPN protection

5. Click **OK**.

What's Next?

- To use the Route-Based VPN on its own, proceed to Defining Route-Based VPN Tunnels (page 965).
- To use the a policy-based VPN to encrypt Route-Based VPN tunnels, proceed to Using the Route-Based VPN in Tunnel Mode (page 967).

Defining Route-Based VPN Tunnels

▼ To define Route-Based VPN tunnels

1. Select **Configuration→Configuration→VPN**. The VPN Configuration view opens.
2. Browse to **Route-Based VPN**. The Route-Based VPN tree opens in the right panel.
3. Right-click **Route-Based VPN** in the Route-Based VPN tree and select **Edit VPN**. The Route-Based VPN opens for editing.
4. Click **Add** to add a tunnel.
5. Configure the settings for the tunnel as instructed below:

Setting	Configuration
Gateway A	Drag and drop an Internal Security Gateway element to the **Gateway A** cell.
Gateway B	Drag and drop an Internal Security Gateway element or an External Security Gateway element to the **Gateway B** cell.
End-Point A	Select the End-Point IP addresses. You cannot use the same End-Point in a Route-Based VPN tunnel and a policy-based VPN tunnel. If loopback IP addresses are defined for an Internal Security Gateway, you can select a loopback IP address as the End-Point IP address. See Configuring Loopback IP Addresses for Firewalls (page 450) for more information about loopback IP addresses.
End-Point B	
Tunnel Interface A	Select the Tunnel Interface(s) on the Firewall(s) though which Route-Based VPN traffic is routed. You must always define **Tunnel Interface A**. Define **Tunnel Interface B** only if Gateway B is an Internal Security Gateway. Because the Management Server does not have any information about the interface configuration of Firewalls managed by a different Management Server, selecting a Tunnel Interface for an External Security Gateway causes configuration errors.
Tunnel Interface B	

Setting	Configuration
Encryption	Select the encryption mode for the tunnel. You can select the encryption mode individually for each tunnel.
Transport Mode	Select the VPN profile that defines the encryption settings for the tunnel. Settings in the VPN Profile that do not apply to the Route-Based VPN, such as IPsec Client settings, are ignored.
Tunnel Mode	Select the policy-based VPN that provides the encryption for the encapsulated Route-Based VPN tunnel. See Using the Route-Based VPN in Tunnel Mode (page 967) for the complete configuration workflow.
No Encryption	**Caution!** This option defines a tunnel in which traffic is encapsulated but not encrypted. It is provided to allow encapsulation of traffic when the traffic does not need to be secured by a VPN. The No Encryption option is recommended only for the following uses: - creating unencrypted tunnels entirely within protected networks - testing and troubleshooting routing and connectivity without using VPN protection
Key (*Optional*, *Transport Mode only*)	If you use pre-shared keys for authentication with external gateways, either set the key agreed with your partner or export the keys that have been automatically generated for your partner to use. To view, change, or export the pre-shared key, double-click the Key cell. See Renewing or Generating Pre-Shared Keys (page 990) for more information. **Caution!** The pre-shared key must be long and random to provide a secure VPN. Change the pre-shared key periodically (for example, monthly). Make sure that it is not possible for outsiders to obtain the key while you transfer it to other devices.
Mode	Select the encapsulation mode for the tunnel. You can select the encapsulation mode individually for each tunnel. The following encapsulation modes are supported: **GRE** (*Generic Routing Encapsulation*) is the industry standard, and is the recommended encapsulation mode in most cases. **IP-IP** (*IP in IP*) is for use with third-party gateways that only support IP-IP. **SIT** (*Simple Internet Transition*) is for use with IPv6 addresses.
TTL (*Optional*)	Specify the initial time-to-live (TTL) value that is inserted into the encapsulation header of packets that enter the tunnel. The TTL setting is needed when dynamic routing is used. The default TTL value is 64. You can usually use the default value, but you may need to configure the TTL setting according to your environment.
MTU (*Optional*)	Enter a Maximum Transmission Unit (MTU) value to specify the largest unit of data that can be transmitted without fragmenting a packet. The MTU size should be as large as possible but not so large that it causes packets to be fragmented. You can usually use the default value, but you may need to configure the MTU setting according to your environment.
PMTU Discovery (*Optional*)	Enable Path MTU (PMTU) Discovery if you use dynamic routing and want to automatically determine the Maximum Transmission Unit (MTU) size on the network path to avoid IP fragmentation.

6. Repeat Step 4-Step 5 to define additional tunnels.

7. Check the **Validity** column for all tunnels.
   - If a tunnel has a warning icon in the **Validity** cell, right-click the tunnel and select **View Issues**. You must resolve all problems indicated in the messages shown.
   - If all tunnels are shown as valid, the Route-Based VPN tunnels are correctly configured, although the Management Server cannot check all possible problems at this point. Additional issues can be shown at policy installation. Any validation and issues that are shown for external gateways are based only on the definitions that have been entered manually into the related elements.
8. Click the Save icon in the toolbar. The Route-Based VPN tunnel configuration is saved.

What's Next?

▶ Add Access rules to allow traffic between the internal network and the networks that are reachable through the Route-Based VPN. See Editing Access Rules (page 672).

Using the Route-Based VPN in Tunnel Mode

When you use the Route-Based VPN in Tunnel Mode, the encapsulation is provided by the Route-Based VPN, and the encryption is provided by a policy-based VPN. For detailed information about IPsec modes and VPNs, see the *Firewall Reference Guide*.

Using the Route-Based VPN in Tunnel Mode allows you to do the following:

- Encrypt multiple encapsulated tunnels in the same VPN tunnel. This improves compatibility with third-party devices and cloud-based services that do not support multiple, separately encrypted tunnels, or that require the use of IPsec Tunnel Mode.
- Create multiple tunnels between remote and local sites when only one public IP address is available.
- Use Multi-Link with Route-Based VPN tunnels for high availability.

▼ To use the Route-Based VPN in Tunnel Mode

1. Create an Internal Security Gateway element as instructed in Creating a New Security Gateway Element (page 924) and define at least two End-Points: one for the policy-based VPN and one for the Route-Based VPN.
2. Create a Host element with the same IP address as the End-Point you use in the Route-Based VPN.

Note – You may receive a warning that the IP address of the Host element is not unique. You can safely ignore the warning and save the element.

3. Add the Host element to a new Site in the Internal Security Gateway as instructed in Adding a New VPN Site (page 935).
4. Use the Internal Security Gateway you created in Step 1 in the policy-based VPN that provides the encryption. See Defining Policy-Based VPNs (page 949).

5. Define one or more Route-Based VPN tunnels. Configure the settings for each tunnel as instructed in Defining Route-Based VPN Tunnels (page 965), except for the following settings:

Setting	Configuration
Gateway A	Use the same Internal Security Gateway that you used in the policy-based VPN in Step 4.
End-Point A	Select the End-Point IP address for which you created a Host element in Step 2.
Encryption	Select **Tunnel Mode** and then select the policy-based VPN that you defined in Step 4.

6. Add Firewall Access rules to allow traffic between the internal network and the networks that are reachable through the Route-Based VPN. See Editing Access Rules (page 672).

Note – The Access rules that direct the Route-Based VPN traffic into the policy-based VPN are automatically generated for the Firewalls associated with the Internal Security Gateway elements. The rules are not visible in the Firewall policy, and cannot be edited. If a policy that contains the automatically-generated rules is installed on a Firewall that is not involved in the VPN, the rules are not included in the configuration that is transferred to the engine.

Monitoring VPNs

Prerequisites: Defining Security Gateways, Defining Policy-Based VPNs, Creating Rules for Policy-Based VPNs

You can monitor the current status of the VPN in the System Status view. The overall status of the VPN elements and the tunnels they contain is shown in the tree of monitored elements. See Reading Component Statuses (page 100) for more information on the colors used.

Logging for VPNs is separate for the tunnels and the traffic that uses the tunnels:

- VPN negotiations are always logged (regardless of the logging options in Access rules) as informational messages. For information on these messages, see Reading IPsec VPN-Related Logs (page 1138) and IPsec VPN Log Messages (page 1241).
- More detailed logging is available when you activate IPsec diagnostic logging for the Firewall/VPN engine for troubleshooting purposes. See Enabling/Disabling Firewall/VPN Diagnostics (page 222).
- The traffic using the VPN tunnels is logged according to the logging options you set for the rule that allows the traffic in or out of the VPN. See Defining Access Rule Logging Options (page 684).
- The System Status view provides shortcuts to logs filtered for the specific VPN or Security Gateway(s) referenced in the log event.
    - Right-click a VPN in the Status tree or connectivity diagram and select **Monitoring→Logs by VPN**.
    - Right-click a Security Gateway in the Status tree or connectivity diagram and select **Monitoring→Logs by Security Gateway**.
    - Right-click the connection between two Security Gateways in the connectivity diagram and select **Monitoring→Logs by Security Gateways** to view logs of traffic between the two Security Gateways.

Log pruning filters that have been activated in your system may delete some (or even all) of the generated messages.

CHAPTER 55

MANAGING VPN CERTIFICATES

A digital certificate is a proof of identity. Stonesoft Firewall/VPN supports using certificates for authenticating gateways and IPsec VPN clients in IPsec VPNs. VPNs are not supported on Layer 2 Firewalls.

The following sections are included:

Getting Started With VPN Certificates

Prerequisites: (*To certify internal Gateways*) Defining Security Gateways

What VPN Certificates Do

You can use certificates for authentication in all types of VPNs as required. Certificates are issued by a certificate authority (CA) as a proof of identity. Gateways that form a VPN tunnel are configured to trust the CA that signed the other gateway's certificate. All certificates issued by a trusted CA are accepted as valid, so certificates can be flexibly added, renewed, and changed without affecting the VPN as long as the actual identity information is correct.

Certificates are an available option in all VPNs, and certificates are always necessary for Gateways to which Stonesoft IPsec VPN Clients connect. Certificates are an available option on the VPN client side as well, but not mandatory.

Certificates reduce the required maintenance work, because they need to be changed much less frequently than pre-shared keys. All certificates are created with an expiration date, after which the certificate is not valid anymore. Certificates signed by the internal IPsec CA are valid for three years from their creation. When a certificate expires, a new certificate is needed.

Certificate Management

Certificate-related tasks in the SMC mostly involve internal VPN Gateways. RSA certificates can be created and renewed automatically, but if DSA certificates or an external CA are involved, some manual steps are required.

There are two options for signing the internal gateway certificates:

- The Management Server includes a dedicated internal CA for signing VPN certificates, the "Internal IPsec CA", which you use through the Management Client. This CA is used in automatic certificate management.
- You can create certificate requests in the Management Client, export them, sign them using an external CA, and then import the signed certificate back into the SMC.

The internal VPN certificate authority can also sign certificate requests created by external components. This feature is mainly meant to support VPN client deployments. The internal VPN CA does not allow you to cancel issued certificates, so consider how widely you can use it for signing external certificate requests within your organization.

Configuration Overview

1. (*Optional*) If you plan to use certificates that are signed by some external certificate authority (CA), define the CA in the system. See Defining a VPN Certificate Authority (page 973).
2. To use an externally-signed certificate, DSA certificate, or if automatic certificate management is disabled, start by Creating a VPN Certificate or Certificate Request for an Internal Gateway (page 975).
3. (*For externally signed certificates*) When the certificate is signed, import it as explained in Importing a VPN Gateway Certificate (page 981).
4. Select a certificate-based **Authentication Method** on the **IKE SA** tab of the VPN Profile for VPNs. See Defining IKE SA Settings for a VPN (page 941).

Defining a VPN Certificate Authority

Prerequisites: You need the root certificate (or a valid certificate) from the certificate authority

For a certificate to be valid, a mutually trusted Certificate Authority (CA) must sign the certificate. Only the internal VPN CA of your system is configured as a trusted VPN CA in VPNs by default. You must define additional CAs in the following cases:

- If you create a VPN with an external gateway and you do not want to use the internal VPN CA to create a certificate for the external gateway. The external gateway must also be configured to trust the issuer of the certificate of your internal gateway.
- If you want to use a certificate signed by an external CA on an internal VPN gateway or on a VPN client.

You can configure the CA as trusted by importing its root certificate or a valid certificate signed by the CA. The certificates must be X.509 certificates in PEM format (Base64 encoding). It may be possible to convert between formats using, for example, OpenSSL or the certificate tools included in Windows.

The CAs you add can be either private (for self-signed certificates) or public (commercial certificate issuers). When you define a CA as trusted, all certificates signed by that CA are considered as valid until their expiration date (or until the CA's certificate expires). Optionally, you can also set up the SMC to check the certificate revocation status from certificate revocation lists (CRLs) or through the OCSP protocol. The CA may cancel a certificate, for example, because it is compromised.

▼ To define a new Certificate Authority

1. Select **Configuration→Configuration→VPN** from the menu to switch to the VPN Configuration view.
2. Expand the **Other Elements** branch in the element tree.
3. Right-click **Certificates** and select **New VPN Certificate Authority**. The VPN Certificate Authority Properties dialog opens.
4. Type a **Name** for the element. This name is only for your reference.

Note – All fields but the Name on the General tab are grayed out. The grayed out fields are always filled in automatically based on information contained in the certificate you import and you cannot change them (information is shown when you close and reopen the element after importing the information).

5. Switch to the **Certificate** tab and do one of the following:
   - Click the **Import** button and import a certificate file.
   - Copy-paste the information into the field on the tab (including the “Begin Certificate” header and “End Certificate” footer).

Tip – The certificate information for many public certificate authorities can be copied and pasted from the default Trusted Certificate Authority elements available in the Security Engine Configuration view under Other Elements→Engine Properties→Certificates→Trusted Certificate Authorities.

Illustration 55.1 Certificate Authority Properties - Certificate Tab

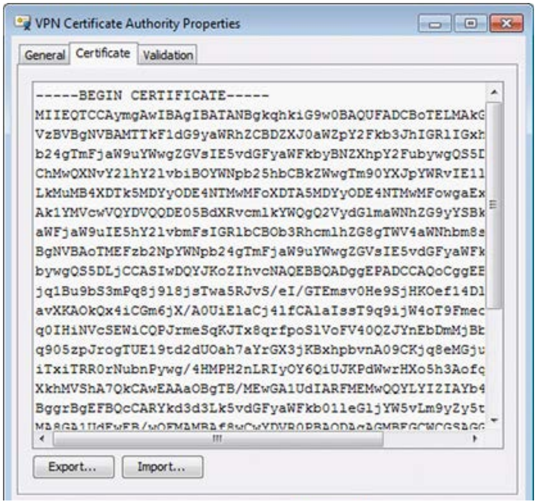

```
-----BEGIN CERTIFICATE-----
MIIEQTCCAymgAwIBAgIBATANBgkqhkiG9w0BAQUFADCBoTELMAkG
VzBVBgNVBAMTTkF1dG9yaWRhZCBDZXJ0aWZpY2Fkb3JhIGR1IGxh
b24gTmFjaW9uYWwgZGVsIE5vdGFyaWFkbyBNZXhpY2FubywgQS5D
ChMwQXNvY21hY21vbiBOYWNpb25hbCBkZWwgTm90YXJpYWRvIE11
LkMuMB4XDTk5MDYyODE4NTMwMFoXDTA5MDYyODE4NTMwMFowgaEx
Ak1YMVcwVQYDVQQDE05BdXRvcmlkYWQgQ2VydGlmaWNhZG9yYSBk
aWFjaW9uIE5hY21vbmFsIGR1bCBOb3RhcmlhZG8gTWV4aWNhbm8s
BgNVBAoTMEFzb2NpYWNpb24gTmFjaW9uYWwgZGVsIE5vdGFyaWFk
bywgQS5DLjCCASIwDQYJKoZIhvcNAQEBBQADggEPADCCAQoCggEE
jq1Bu9bS3mPq8j918jsTwa5RJvS/eI/GTEmsv0He9SjHKOef14Dl
avXKAOkQx4iCGm6jX/A0UiElaCj41fCAlaIssT9q9ijW4oT9Fmec
q0IHiNVcSEWiCQPJrmeSqKJTx8qrfpoSlVoFV40QZJYnEbDmMjBk
q905zpJrogTUE19td2dU0ah7aYrGX3jKBxhpbvnA09CKjq8eMGju
iTxiTRR0rNubnPywg/4HMPH2nLRIyOY6QiUJKPdWwrHXo5h3Aofq
XkhMVShA7QkCAwEAAaOBgTB/MEwGA1UdIARFMEMwQQYLYIZIAYb4
BggrBgEFBQcCARYkd3d3Lk5vdGFyaWFkb01leGljYW5vLm9yZy5t
MA8GA1UdEwEB/wQFMAMBAf8wCwYDVR0PBAQDAgGMBEGCWCGSAGG
```

6. (*Optional*) If you want the Firewall/VPN engines to check the revocation status of certificates signed by this CA, switch to the **Validation** tab and select the options as follows:
   - Select the **Check Validity on Certificate-Specified CRLs** option to activate CRLs for certificate status checking.
   - Select the **Check Validity on Certificate-Specified OCSP Servers** option to activate OCSP certificate status checking.
7. (*Optional*) To define more CRL servers to check in addition to those defined in the certificates, click the corresponding **Add** button and select:
   - **LDAP Server Elemen**t to choose an existing element
   - **Manual LDAP Server Address** to type in the address in a dialog that opens.

Example `ldap://`*`example.com`*`:`*`389`*

8. (*Optional*) To define more OCSP servers to check in addition to those defined in the certificates, click the corresponding **Add** button and type in an address in the dialog that opens.

Example `http://ocsp.example.com`

Caution – When certificate checking is defined, all certificates signed by the CA are considered as invalid if the validity check cannot be performed (for example, due to incorrectly entered address or connectivity problems).

9. Click **OK**. If you see an invalid certificate error, the certificate you imported may be in an unsupported format. Try converting the certificate to an X.509 certificate in PEM format (Base64 encoding) using OpenSSL or the certificate tools included in Windows.

If you use the Default policy template, both LDAP (port 389) and HTTP (port 80) connections from the engine are allowed. If your firewall or server configuration differs from these standard definitions, edit the firewall policy to allow the necessary connections from the engines.

You can define several Certificate Authority servers in the system. By default, all CAs you have defined in the system are trusted by all Gateways and in all VPNs. If necessary, you can limit trust to a subset of the defined CAs when you configure the Security Gateway and VPN Profile elements.

What's Next?

- By default, all Gateways trust all configured VPN certificate authorities in all VPNs. The trust relationships can be changed at the gateway level (Defining Trusted CAs for a Gateway (page 929)) and in the VPN Profiles (Defining Trusted CAs for a VPN (page 948)).
- To obtain a certificate from an external certificate authority, first create a certificate request as explained in Creating a VPN Certificate or Certificate Request for an Internal Gateway (page 975) and then import the signed certificate as explained in Importing a VPN Gateway Certificate (page 981).

Creating and Signing VPN Certificates

Prerequisites: (*To certify internal Gateways*) Defining Security Gateways

Creating a VPN Certificate or Certificate Request for an Internal Gateway

Following the instructions below, you can create a certificate request and sign it either using the internal or an external certificate authority (CA). If the Internal Gateway has automatic certificate management activated in its properties, the steps below are necessary only when you want to use DSA certificates or create a certificate request to be signed by an external CA.

▼ To create a VPN certificate or certificate request for an internal gateway

1. Select **Configuration→Configuration→VPN** from the menu to switch to the VPN Configuration view.
2. Select **Gateways**. The defined Gateways are displayed.
3. Right-click the Gateway that needs a VPN certificate and select **Tools→Generate Certificate**. The Generate Certificate dialog opens.

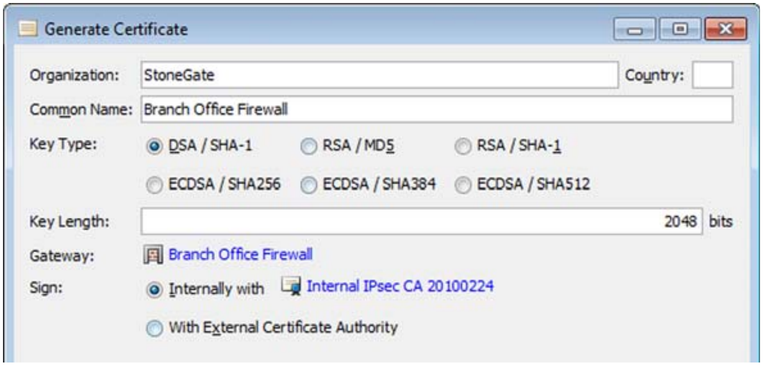

Generate Certificate

Organization: StoneGate
Country:
Common Name: Branch Office Firewall
Key Type: ● DSA / SHA-1 ○ RSA / MD5 ○ RSA / SHA-1 ○ ECDSA / SHA256 ○ ECDSA / SHA384 ○ ECDSA / SHA512
Key Length: 2048 bits
Gateway: Branch Office Firewall
Sign: ● Internally with Internal IPsec CA 20100224
○ With External Certificate Authority

4. Enter the certificate information:
   - **Organization**: the name of your organization as it should appear in the certificate.
   - **Country**: (*optional*) standard two-character country code for the country of your organization.
   - **Common Name**: a descriptive name for the Gateway as it should appear in the certificate.
   - **Key Length**: length of the key for the generated public-private key pair. The default is 2048 bits.
   - **Gateway**: internal reference to the Gateway element. Not editable.
   - **Key Type**: select the type of key you want to use according to the requirements of your organization and, if you use an external certificate authority, possible restrictions in supported key types.
5. Select how you want to **Sign** the certificate:
   - Select **Internally** to sign the certificate using the SMC's internal VPN Certificate Authority. If the internal VPN Certificate Authority is in the process of being renewed and there are temporarily two valid internal VPN Certificate Authorities, select the one you want to use from the list of internal VPN Certificate Authorities.
   - Select **Externally** if you want to create a certificate request that some other Certificate Authority signs.
6. Click **OK**. The signed certificate or unsigned certificate request is added under the Gateway in the gateway list.
   - Note that there may be a slight delay while the request is generated.
   - If you signed the certificate using the internal VPN CA, the certificate is automatically transferred to the Firewall/VPN engine and no further action is needed.
7. With external certificate authorities, further action is required: Right-click the certificate request and select **Export Certificate Request**. A file save dialog opens for saving the exported certificate request.
   - To generate certificates for an Internal Gateway, the CA must support PKCS#10 certificate requests in PEM format (Base64 encoding) and the signed certificates must also be in the PEM format. It may be possible to convert between formats using, for example, OpenSSL or the certificate tools included in Windows.
   - The CA must be able to copy all attributes from the certificate request into the certificate; especially X.509 extension Subject Alternative Name must be copied as it is in the request because the value is used for authentication.

What's Next?

▶ When you receive the signed certificate, import it as explained in Importing a VPN Gateway Certificate (page 981).

Signing External Certificate Requests Internally

The SMC's internal VPN Certificate Authority can be used to sign external certificate requests. This feature is mainly intended for use with VPN clients, but it can be used to sign any certificate request that is in the supported format (PKCS#10 certificate requests in PEM format). An alternative is to configure the Internal Gateway to accept an externally signed certificate by defining the external certificate issuer as trusted. See Defining a VPN Certificate Authority (page 973).

Make sure that the date, time, and time zone are all set correctly on the Management Server and on the external component that uses the certificate. Certificates are valid for three years starting from the date and time they are created. The validity start and end date and time are written in the certificate and are enforced in the authentication.

Caution – The internal CA does not support certificate revocation lists, so it is not possible to cancel an internally signed certificate before it expires.

▼ To sign an external certificate request

1. Select **Configuration→Configuration→VPN** from the menu to switch to the VPN Configuration view. A VPN-specific toolbar is shown at the top edge of the element panel.

Illustration 55.2 VPN Tools Menu

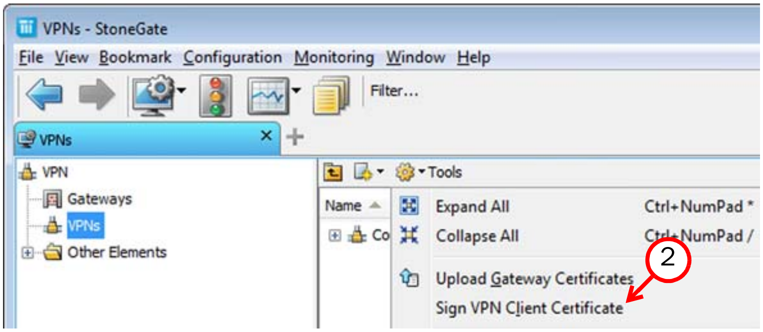
VPNs - StoneGate
File View Bookmark Configuration Monitoring Window Help
Filter...
VPNs
VPN
Gateways
VPNs
Other Elements
Tools
Name
Co
Expand All Ctrl+NumPad *
Collapse All Ctrl+NumPad /
Upload Gateway Certificates
Sign VPN Client Certificate
2

2. Click the Tools icon in the toolbar and select **Sign VPN Client Certificate**. The Sign VPN Client Certificate dialog opens. The same dialog works for all X.509 certificate requests in PEM format and is not limited to VPN client certificate request signing.

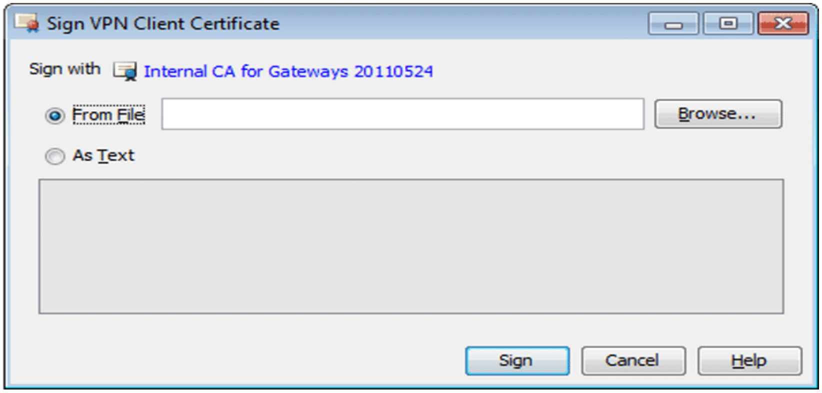
Sign VPN Client Certificate
Sign with Internal CA for Gateways 20110524
● From File
Browse...
○ As Text
Sign
Cancel
Help

3. If the internal VPN Certificate Authority is in the process of being renewed and there are temporarily two valid internal VPN Certificate Authorities, select the internal VPN Certificate Authority for signing the certificate request from the list at the top of the dialog.
4. Either browse to the certificate request file on your local workstation or copy-paste the content of the certificate request into the dialog (with the “begin certificate request” header and “end certificate request” footer).
5. Click **Sign**. The certificate is signed and the dialog switches to show information on the new certificate. Note the expiration information on the **Certificate** tab.
6. On the **General** tab, click **Export** to save the certificate for transfer to the device that needs it.
7. Click **OK**.

Uploading VPN Certificates Manually

Prerequisites: Creating a VPN Certificate or Certificate Request for an Internal Gateway

In most cases, certificates are transferred to the engines automatically. However, if there are problems with missing VPN certificates on the Gateways, you can transfer them to the engines manually.

▼ To transfer VPN certificates to the engines

1. Select **Configuration→Configuration→VPN** from the menu to switch to the VPN Configuration view. A VPN-specific toolbar is shown at the top edge of the element panel.

Illustration 55.3 VPN Tools Menu

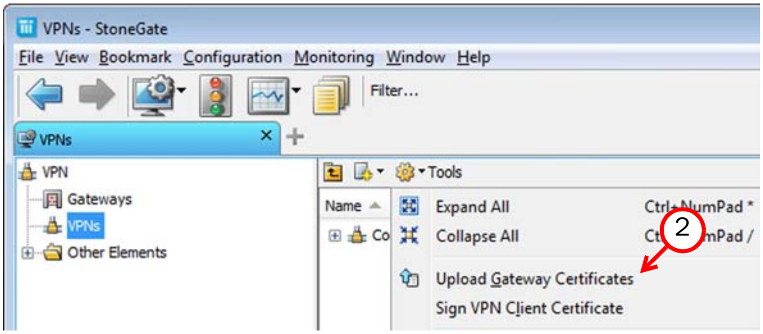
VPNs - StoneGate
File View Bookmark Configuration Monitoring Window Help
Filter...
VPNs
VPN
Gateways
VPNs
Other Elements
Tools
Name
Co
Expand All Ctrl+NumPad *
Collapse All Ctrl+NumPad /
Upload Gateway Certificates
Sign VPN Client Certificate
2

2. In the VPN-specific toolbar, click the Tools icon and select **Upload Gateway Certificates**. The certificates are either uploaded, or you will see a message that all certificates have already been automatically uploaded.

Related Tasks

- Troubleshooting Certificates (page 1091)

Renewing VPN Certificates

Prerequisites: None

For security reasons, certificates have an expiry date, after which the certificate must be replaced with a new one. The process is partially automatic when internally-signed certificates are used, and the necessary steps are explained in this section. To create a new externally-signed certificate, see Creating a VPN Certificate or Certificate Request for an Internal Gateway (page 975).

The certificates issued by the internal VPN Certificate Authority are valid for three years. If automatic RSA certificate management is activated for an internal security gateway, RSA certificates issued by the internal VPN Certificate Authority are renewed automatically without your intervention as long as the certificate-related files are intact (including the private key stored on the engines).

The internal VPN Certificate Authority itself is valid for ten years. A new internal VPN Certificate Authority is automatically created six months before the internal VPN Certificate Authority's expiration date. New certificates signed by the new internal VPN Certificate Authority are automatically created for internal gateways. If certificates are used to authenticate VPN client users and the certificates have been signed by the expiring VPN Certificate Authority, you must manually create new certificates for the VPN clients. You must also create new certificates manually for any other external components that have certificates signed by the internal VPN Certificate Authority.

Note – When you renew the VPN certificate, Stonesoft IPsec VPN Client users receive a notification about the certificate fingerprint change. Notify your users before you renew the certificate if possible.

▼ To renew an externally signed certificate of an internal Gateway

1. Create a new certificate request as explained in Creating a VPN Certificate or Certificate Request for an Internal Gateway (page 975).
2. Import the signed certificate as explained in Importing a VPN Gateway Certificate (page 981).

▼ To renew an internally signed certificate of an external component

1. Create a new certificate request in the external component. For Stonesoft IPsec VPN Clients, this is explained in the *IPsec VPN Client User's Guide*.
2. Sign the certificate as explained in Signing External Certificate Requests Internally (page 977).

▼ To manually renew internally signed Internal Gateway certificates

1. Select **Configuration→Configuration→VPN** from the menu to switch to the VPN Configuration view.
2. Browse to **Virtual Private Networks→Certificates→Gateway Certificates**. The certificates are shown with their expiration dates and signer information.
3. Right-click the certificate you want to renew and select **Renew Certificate**.
4. Click **Yes**. There is a delay while the certificate is renewed, after which you are notified that the certificate was renewed. The certificate is transferred to the engine automatically.
5. Refresh the policy of the Firewall/VPN engine to activate the new certificate.

The procedure explained above renews the certificate when the certificate-related information is intact on the engine and on the Management Server. If the certificate has not expired but is affected by other problems, delete the existing certificate element in the Management Client and create a new one (see Creating a VPN Certificate or Certificate Request for an Internal Gateway (page 975)).

Exporting the Certificate of VPN Gateway or VPN CA

Prerequisites: A signed certificate is present in the system or a new internal VPN CA has been created

In addition to certificate requests, you can also export signed gateway certificates and the certificates of VPN Certificate Authorities. This is not usually necessary, but can be done as needed.

If the system has created a new internal VPN Certificate Authority, you must export the certificate of the new internal VPN Certificate Authority and import the certificate on external gateways that use certificates signed by the internal VPN Certificate Authority or communicate with gateways that use certificates signed by the internal VPN Certificate Authority. If the external gateway itself uses a certificate signed by the internal VPN Certificate Authority, you must also create a new certificate for the external gateway (see Renewing VPN Certificates (page 979).

Certificates that are created when the internal VPN Certificate Authority signs an external certificate request must be exported at the time of signing and are not stored for exporting at a later time.

▼ To export a signed certificate

1. Select **Configuration→Configuration→VPN** from the menu to switch to the VPN Configuration view.
2. Browse to **Other Elements→Certificates→Gateway Certificates**. The certificates are shown with their expiration dates and signer information.
3. Right-click a certificate and select **Export Certificate**. A file save dialog opens.
4. Browse to the location where you want to save the file on your local workstation and click **Save**.

▼ To export the certificate of the Internal VPN Certificate Authority

1. Select **Configuration→Configuration→VPN** from the menu to switch to the VPN Configuration view.
2. Browse to **Other Elements→Certificates→VPN Certificate Authorities**. The VPN Certificate Authorities are shown with their expiration dates.
3. Right-click a VPN Certificate Authority and select **Tools→Export Elements**. The Export dialog opens.
4. Browse to the location where you want to save the file on your local workstation and click **Export**.

If the external gateway uses a certificate signed by the internal VPN Certificate Authority and the internal VPN Certificate Authority has been renewed, you must create a new certificate for the external gateway. See Renewing VPN Certificates (page 979).

Importing a VPN Gateway Certificate

Prerequisites: Creating a VPN Certificate or Certificate Request for an Internal Gateway

This procedure allows you to import a certificate signed by an external certificate issuer for an internal VPN Gateway when the certificate request has been created internally. For security reasons, it is not possible to import externally generated private keys.

Note – All CAs that issue certificates for your VPNs must be configured in the system (see Defining a VPN Certificate Authority (page 973)) and be included as trusted both at the Gateway and VPN Profile levels.

▼ To import a VPN gateway certificate

1. Select **Configuration→Configuration→VPN** from the menu to switch to the VPN Configuration view.
2. Open the **Gateways** branch and expand the tree under the Gateway.
3. Right-click the certificate request and select **Import Certificate**.
4. Select the certificate authority that signed the certificate.
5. Browse to the certificate file or copy-paste the file's content as text.
6. After the signed certificate is imported, delete the certificate request, which is still displayed under the Gateway along with the signed certificate.

The certificate is transferred to the engine automatically.

Checking When Gateway Certificates Expire

Prerequisites: A signed certificate is present in the system

By default, RSA certificates issued by the internal VPN Certificate Authority to internal VPN Gateways are renewed automatically. In other cases, the certificates expire according to the information written in the certificate when it was generated. Internal Gateways never accept expired certificates.

▼ To check when VPN certificates expire

1. Select **Configuration→Configuration→VPN** from the menu to switch to the VPN Configuration view.
2. Expand the **Other Elements→Certificates** branch in the element tree and select **Gateway Certificates**. The existing certificates are displayed in the other panel.
3. See the **Expiration Date** column for information on the certificate's expiry.
   - You can renew internally signed certificates through their right-click menu (necessary only if automatic certificate management is not active). VPN client users may be prompted to accept the change of certificate.
   - Elements with no expiration date are certificate requests (Status: To be signed).

Checking When an Internal VPN CA Expires

Prerequisites: None

The Management Server includes a dedicated internal VPN CA for signing VPN certificates, the “Internal IPsec CA”. The internal VPN CA is valid for ten years. A new internal VPN CA is automatically created six months before the internal VPN CA’s expiration date. You can check the expiration date and status of the internal VPN CA.

▼ To check the status of an internal VPN CA

1. Select **Configuration→Configuration→VPN** from the menu to switch to the VPN Configuration view.
2. Expand the **Other Elements→Certificates** branch in the element tree and select **VPN Certificate Authorities**. The existing VPN Certificate Authorities are displayed in the other panel.
3. See the **Expiration Date** column for information on the certificate’s expiration date.
4. To view detailed information, right-click a VPN Certificate Authority and select **Properties**. The Properties dialog for the internal VPN CA opens.
5. See the **Validity** information in the Valid from and Valid to fields.
6. See also the **Status** information:
   - Active: You can use this Internal VPN CA to sign certificates.
   - Renewal Started: This is a new Internal VPN CA that the system has created automatically. The process of renewing VPN certificates for internal gateways has begun.
   - Expires Soon: A new internal VPN CA has been created but some components may still use certificates signed by this internal VPN CA.
   - Inactive: This internal VPN CA has expired or no system component uses a certificate signed by this internal VPN CA.

Related Tasks

- ▶ Understanding Certificate-Related Problems (page 1092)
- ▶ Dealing with Expiring Certificate Authorities (page 1097)
- ▶ Checking When Gateway Certificates Expire (page 981)

Chapter 56

Reconfiguring Existing VPNs

The topics below provide additional instructions for some common configuration and tuning tasks for existing VPNs.

The following sections are included:

Adding or Removing Tunnels in a VPN

Prerequisites: You must have a VPN configured

Before changing the tunnels that are used in active VPNs, we recommend making a backup of the Management Server as instructed in Creating Backups (page 1009).

In the Route-Based VPN, you must add or remove tunnels manually. See Defining Route-Based VPN Tunnels (page 965).

In a policy-based VPN, the number of tunnels generated for a VPN is determined by the Gateway elements on the **Overall Topology** tab of the VPN element and the number of End-Points there are active in each Gateway element. After changing the topology of a policy-based VPN, always check that all new or changed Tunnels are valid on the **Tunnels** tab.

- Each Central Gateway forms a tunnel with each central and satellite Gateway in the VPN. No other Gateway<->Gateway tunnels are created. See Defining VPN Topology for Policy-Based VPNs (page 951) for detailed instructions. Tunnels are not generated between End-Points that cannot connect to each other. For example, tunnels are not generated between two End-Points if they both have a dynamic IP address.
- Adding a Gateway under another Gateway instead of directly at the main level in the central Gateways list may prevent tunnel generation. This configuration implies that the Gateway at the main level forwards connections to the Gateway(s) below it in the hierarchy. For the forwarding to work, it must be explicitly configured in the Central Gateway's Access rules with the **Use VPN→Forward** action.
- End-Point<->End-Point tunnels are created between all End-Points defined in the properties of any two Gateways that form a Gateway<->Gateway tunnel. See Defining End-Points for Internal Security Gateways (page 924) or Defining End-Points for External Security Gateways (page 927).

Related Tasks

- ▶ Defining Security Gateways (page 923)
- ▶ Adding New Gateways to an Existing VPN (page 986)

Configuring NAT Settings for an Existing VPN

Prerequisites: You must have a VPN configured

Activating NAT Traversal

NAT traversal prevents intermediary devices from applying NAT to VPN communications if NAT is found to prevent the communications from working.

NAT traversal encapsulates the IKE and IPsec communications inside UDP packets (NAT-T) or tunnels them inside TCP packets (proprietary method). The TCP tunneling option is available for client-to-gateway VPNs only. There is also a legacy UDP encapsulation method that is not supported on firewall engine versions 5.2 and higher.

The UDP and NAT-T encapsulation options do not affect client-to-gateway VPNs. NAT-T is always active in client-to-gateway VPNs.

Encapsulation is not always necessary. Usually, the VPN can be configured to work even when NAT is applied simply through the definition of Contact Addresses. For more information, see Getting Started with System Communications (page 64).

The encapsulation options are activated in the End-Point properties in the Gateway element. See Defining End-Points for Internal Security Gateways (page 924) and Defining End-Points for External Security Gateways (page 927).

Translating Addresses of VPN Communications Between Gateways

The communications needed to establish and maintain VPN tunnels between VPN Gateways are always translated according to the NAT rules (as opposed to the communications that use the VPN tunnel. See Translating Addresses in Traffic Inside a VPN Tunnel (page 986). Create a matching NAT rule that defines some NAT operation to translate addresses in these communications and make sure that the communications do not match the wrong NAT rule unintentionally.

When you add NAT, you may need to change VPN settings. Add contact addresses for the firewalls: see Getting Started with System Communications (page 64). Contact Addresses can be used with both internal and external gateways.

There is nothing VPN-specific about creating the actual NAT rules. See Editing Firewall NAT Rules (page 687).

Translating Addresses in Traffic Inside a VPN Tunnel

By default, IP addresses in traffic that enters or leaves a VPN tunnel are not translated. This behavior is controlled by an option in the properties of the VPN element, accessible through the right-click menu for the VPN in question.

If the option to translate the IP addresses is enabled, the IP addresses in traffic that uses gateway-to-gateway VPN tunnels are translated according to the NAT rules. There is nothing VPN-specific in creating these NAT rules. However, the VPN configuration is affected if local protected addresses are translated using NAT:

- Set the Site that contains the private local addresses (before translation) in the Private mode in VPNs in which those addresses are translated using NAT.
- Add the translated addresses as a new Site for the Gateway (disable the Site in other VPNs). This Site is in the default Normal mode.

If the option to translate the IP addresses is enabled, VPN client traffic is translated according to the NAT Pool settings in the Firewall element's properties (see Managing VPN Client IP Addresses (page 1000)) or as defined in the NAT rules.

For more information on configuring NAT rules. See Editing Firewall NAT Rules (page 687).

Adding New Gateways to an Existing VPN

Prerequisites: None

In the Route-Based VPN, you must define new tunnels when you add new Gateways. See Defining Route-Based VPN Tunnels (page 965).

If you have already configured a policy-based VPN, you can add new Gateways as needed following the general configuration workflow outlined below:

1. Create a Gateway element to represent the physical gateway device in VPNs if the element does not exist already. See Defining Security Gateways (page 923). Once created, the same element can be used in any number of VPNs.
2. If the VPN uses Certificates for authentication, you may need to create a new VPN certificate for the Gateway. See Getting Started With VPN Certificates (page 972). The same certificate can be used in any number of VPNs, providing it fulfills the following criteria:
   - The certificate must match the type of certificate selected for the VPN in the VPN Profile.
   - The certificate must be issued by a certificate authority that the other Gateways trust.
3. Edit the VPN element to add the Gateway on the **Overall Topology** tab. See Defining VPN Topology for Policy-Based VPNs (page 951).
4. Check and adjust the Tunnels between the new Gateway and the existing Gateways. See Defining VPN Tunnel Settings for Policy-Based VPNs (page 953).
5. Refresh the policies of all Firewall/VPN engines that are involved in the tunnel.

Changing Gateway IP Addressing in an Existing VPN

Prerequisites: None

For internal Gateways, the VPN End-Point addresses are determined by the IP addresses you have defined for the interfaces in the Firewall element. On clusters, only CVI addresses are used as VPN End-Points. VPNs are not supported on Layer 2 Firewalls.

Note – If the Gateway's identity in the VPN is tied to its IP address, you must update the configurations of all gateways in the VPN even if the IP address is NATed and not directly used for contact. For internal Gateways, this is done by refreshing the engine's policy after making changes. For external Gateways, change the information in the configuration of that gateway device.

- If you change the IP address of some interface of a firewall, the corresponding VPN End-Point IP address also changes automatically and the existing tunnels in the VPN element are preserved.
- If continuous connectivity is required, define the old and new IP address as two separate interfaces, select both as End-Points in the VPN, and refresh the policies of all affected Firewall/VPN engines. In this configuration, either of the IP addresses can be used. The Multi-Link VPN automatically selects the End-Point that works.
- If you add and remove interfaces, you may need to select and deselect End-Points manually and then check the tunnel configuration in the VPN element. See Defining End-Points for Internal Security Gateways (page 924) and Defining VPN Tunnel Settings for Policy-Based VPNs (page 953).

Note – You cannot use the same End-Point in a policy-based VPN and the Route-Based VPN.

For External Gateways, the VPN End-Point addresses are always entered manually. Change the IP address configured in the Management Client and refresh the policies of all affected Firewall/ VPN engines.

- If continuous connectivity is required, define the new address as a second End-Point before the address is actually changed. The MultiLink VPN automatically selects the IP address that works before and after the change.
- See Defining End-Points for External Security Gateways (page 927).

Giving VPN Access to Additional Hosts

Prerequisites: None

In the Route-Based VPN, it is not necessary to modify the VPN configuration to allow access through the VPN for additional hosts. Any traffic that is routed to a Tunnel Interface and allowed by the Access rules automatically uses the Route-Based VPN tunnel.

If you want to give access through a policy-based VPN to hosts with IP addresses that are not already configured for your VPN, proceed according to this general workflow:

1. Make sure the IP addresses (possibly the translated IP addresses if NAT is enabled in this VPN) are included in one of the Sites of the correct Gateway. If the IP addresses must not be included in other VPNs where the same Gateway element is used, add them to a separate Site and disable the Site in other VPNs. See Adding a New VPN Site (page 935).
2. If there are external Gateways involved, make sure any IP addresses you add to the Site definition in the Management Client are also added to the configuration of the external gateway device, so that it routes the traffic through the VPN.
3. Check that the Access rules of all Gateways involved specify that this traffic is sent/ allowed through the VPN. If NAT is enabled in the VPN, check also the NAT rules. See Creating Rules for Gateway Connections in Policy-Based VPNs (page 959) and Creating NAT Rules for Policy-Based VPN Traffic (page 963).

Routing Internet Traffic Through Policy-Based VPNs

Prerequisites: A working VPN between all Gateways involved

It is possible to force all traffic from VPN clients or clients in protected networks to be routed through a policy-based VPN so that the traffic can be inspected centrally. The configuration requires the following settings:

1. Enable NAT for tunneled traffic in the VPN element's properties. See Modifying an Existing VPN Element (page 951).
2. Change the mode of all of the central Gateway's Sites in this VPN to **Private** and replace them with a Site that contains the Any Network element. Disable the Any Network Site in other VPNs. See Defining Sites for VPN Gateways (page 933).
3. Reconfigure the policy:
   - Create Access rules. See Creating Forwarding Rules on Hub Gateways for Policy-Based VPNs (page 961).
   - Create NAT rules that translate any private IP addresses to public addresses for the Internet. See Creating NAT Rules for Policy-Based VPN Traffic (page 963).
4. Redirect the traffic from external components to the central gateway as necessary:
   - For Firewall/VPN Gateways, add an Access rule that sends the desired traffic to the VPN. See Creating Rules for Gateway Connections in Policy-Based VPNs (page 959).
5. (*VPN Clients only*) Configure the Virtual Adapter as explained in Configuring Virtual IP Addressing for VPN Clients (page 1002).

Redirecting Traffic Between VPN Tunnels

Prerequisites: (*If VPN client traffic is forwarded*) Configuring Virtual IP Addressing for VPN Clients

In policy-based VPNs, it is possible to redirect traffic from one VPN tunnel to some other VPN tunnel through a hub gateway. This is especially useful for VPN client users who need access to resources at different locations, because the users do not have to separately connect to different gateways according to the services they want to use.

See Getting Started With IPsec VPNs (page 918) if you need more instructions for creating and editing VPNs and the elements involved.

See Configuration 4: Basic VPN Hub (page 910) for a step-by-step configuration example.

▼ To redirect traffic between VPN tunnels

1. Place the forwarding hub Gateway at the top level of the Topology tree.
   - The **IPsec Client** Gateway is always inserted at the top level.

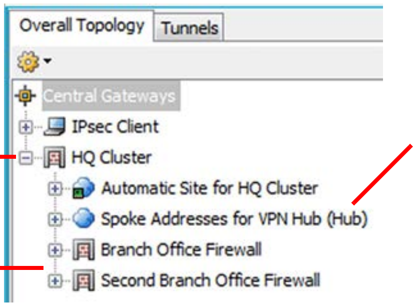
Overall Topology
Tunnels
Central Gateways
IPsec Client
HQ Cluster
Automatic Site for HQ Cluster
Spoke Addresses for VPN Hub (Hub)
Branch Office Firewall
Second Branch Office Firewall

Hub Gateway.

These Gateways are contacted through the hub.

Protected IP addresses behind the spoke Gateways.

2. Place the Gateways that are contacted through the hub under the hub Gateway.

Note – Duplicate tunnels are not allowed. There must not be gateway-to-gateway connections between the hub and the other gateways in other active VPNs.

3. Add a Site that contains the IP addresses behind the spoke Gateways under the hub Gateway.
   - Set the Mode of the Site to **Hub**.
   - Disable the Site in any other VPNs where it is used.
4. If you want to forward VPN client traffic, add a Site that contains the virtual IP address space used for the VPN clients under the hub Gateway (if the IP address space is not already included).
5. Add Access rules that forward the traffic between tunnels as instructed in Creating Forwarding Rules on Hub Gateways for Policy-Based VPNs (page 961).
6. Refresh the policies of all Gateways, starting from the hub gateway.

Optionally, all traffic (including Internet traffic) can be routed through the hub Gateway (see Routing Internet Traffic Through Policy-Based VPNs (page 988)).

Renewing or Generating Pre-Shared Keys

Prerequisites: Pre-shared key authentication is selected in the VPN Profile and allowed in the Gateway Profiles

Renew or generate pre-shared keys according to one of the two sets of instructions below. You can automatically generate pre-shared keys only for policy-based VPNs. You can manually configure pre-shared keys for policy-based VPNs and for the Route-Based VPN.

Generating a New Pre-Shared Key Automatically

As a security precaution, we recommend that you periodically change the pre-shared key (for example, monthly).

▼ To regenerate new pre-shared keys automatically

1. Select **Configuration→Configuration→VPN**. The VPN Configuration view opens.
2. Right-click a VPN element and select **Edit VPN**. The VPN opens for editing.
3. Switch to the **Tunnels** tab.
4. Select all Tunnels for which you want to renew the pre-shared keys.
5. Right-click one of the selected Tunnels and select **Delete Pre-Shared Key**.
6. Right-click one of the selected Tunnels again and select **Generate Missing Pre-Shared Key**.
7. If you need to transfer keys to external components, double-click the **Key** column for those tunnels and copy-paste or export the key.

Caution – Make sure that outsiders cannot obtain the key while you transfer it to other devices. The key must remain secret to be an effective security measure.

Renewing Pre-Shared Keys Manually

As a security precaution, we recommend that you periodically change the pre-shared key (for example, monthly).

▼ To renew pre-shared keys manually

1. Select **Configuration→Configuration→VPN**. The VPN Configuration view opens.
2. Right-click a VPN element or the Route-Based VPN element and select **Edit VPN**. The VPN editing view or Route-Based VPN editing view opens.
3. (*VPN elements only*) Switch to the **Tunnels** tab.
4. Double-click the **Key** column for the tunnel for which you want to change the key.
5. Copy-paste or type in a new key into the dialog and click **OK**.

Caution – The pre-shared key must be long and random to provide a secure VPN.

Advanced VPN Tuning

Prerequisites: None

The Gateway Settings element defines performance-related VPN options for the Firewall/VPN engines. These settings are used internally and there is no need to match them exactly with settings of other gateways in VPNs.

The Gateway Settings elements are stored under the **Other Elements→Profiles→Gateway Settings** branch in the VPN Configuration view.

Illustration 56.1 Gateway Settings in the Element Tree

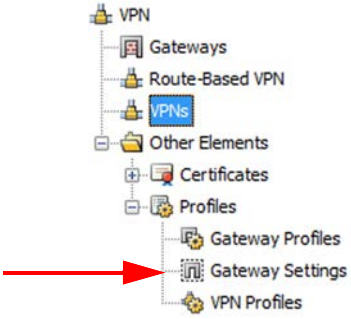

What's Next?

- If you want to change any of the Gateway settings, proceed to Defining a Custom Gateway Settings Element (page 991).

Defining a Custom Gateway Settings Element

▼ To create new Gateway Settings

1. Select **Configuration→Configuration→VPN**. The VPN Configuration view opens.
2. Expand the **Other Elements→Profiles** branch in the element tree.
3. Right-click the **Gateway Settings** branch in the element tree and select **New Gateway Settings**. The Gateway Settings Properties dialog opens.

Proceed according to the settings you want to change.

What's Next?

- Adjusting MOBIKE Settings (page 992)
- Adjusting Negotiation Retry Settings (page 992)
- Adjusting Certificate Cache Settings (page 993)

Adjusting MOBIKE Settings

MOBIKE (mobile IKE) return routablility checks (RRC) can be used together with IKEv2 to verify the validity of a VPN client's IP address or the gateway's IP address if the IP address changes in the middle of an open VPN connection. The IP address is updated in the negotiated SAs when the new IP address has been verified. If the new IP address cannot be verified, the VPN connection is closed. By default, no return routablility checks are done.

Table 56.1 MOBIKE RRC Settings

Parameters	Description
Before SA Update	The remote peer's IP address is first checked and then, if the IP address is valid, the IP address is updated in the SA.
After SA Update	The IP address is updated in the SA and then verified.

What's Next?

- Adjusting Negotiation Retry Settings (page 992)
- Adjusting Certificate Cache Settings (page 993)
- If you are done with Gateway Settings, click **OK** and continue by assigning the new Gateway settings as explained in Assigning the Gateway Settings for a Firewall/VPN Engine (page 994).

Adjusting Negotiation Retry Settings

See the table below for an explanation of the effect of the settings in the **Negotiation Retry** section in Gateway Settings properties. If a negotiation for a VPN does not complete successfully, the VPN establishment is retried according to the settings available here. The default settings are the recommended values.

If the VPN establishment often fails because you have frequent intermittent problems with network connectivity or because your network connection is very slow, increasing these values may be a work-around solution for getting the VPN to establish.

Table 56.2 Negotiation Retry

Parameters	Description
Retry Limit	Number of possible retries when sending a packet to a remote gateway. The negotiation for opening a tunnel is cancelled when negotiation has been attempted the number of times specified.
Retry Timer	Initial time between retry attempts. This value doubles at each attempt until the value entered for the Retry Timer Maximum is reached. Enter the value in milliseconds.
Retry Timer Maximum	Maximum delay between two retry attempts. When the Retry Timer reaches this value, all further retries are done at the interval defined here until the Retry Limit is reached. Enter the value in milliseconds.
Negotiation Expiration	Tunnel negotiation is cancelled if it is not completed before the expiration time is reached. Enter the value in milliseconds.

What's Next?

- Adjusting Certificate Cache Settings (page 993)
- If you are done with Gateway Settings, click **OK** and continue by assigning the new Gateway settings as explained in Assigning the Gateway Settings for a Firewall/VPN Engine (page 994).

Adjusting Certificate Cache Settings

See the table below for an explanation of the effect of the settings in the **Certificate Cache** section in Gateway Settings properties. These settings have effect only if you use certificates to authenticate the security gateways in IKE negotiations. The default settings are the recommended values. Adjust the settings if you have a specific need to do so.

Table 56.3 Certificate Cache

Parameters	Description
Maximum Number	Maximum number of certificates stored in the cache.
Maximum Size	Maximum size of the certificate cache. Define the value in bytes.
CRL Validity	Defines how long any Certificate Revocation Lists (CRL) are valid. Overrides the validity stated on the received CRLs. Typically, individual CRLs should not be trusted over long periods of time. When the CRL validity time is reached, the Firewall/VPN engine fetches new CRLs from their issuing authorities. Define the value in seconds.

What's Next?

- Click **OK** and continue by assigning the new Gateway settings as explained in Assigning the Gateway Settings for a Firewall/VPN Engine (page 994).

Assigning the Gateway Settings for a Firewall/VPN Engine

By default, all Firewall/VPN engines refer to the **Gateway Default Settings** Gateway Settings element. If you want to change the Gateway Settings to a set you have created yourself, follow the directions below. VPNs are not supported on Layer 2 Firewalls.

▼ To define the Gateway Settings for a Firewall/VPN engine

1. Double-click the Firewall element. The Firewall Properties dialog opens.
2. Switch to the **Advanced** tab.
3. Click the **VPN Settings** button. The VPN Settings dialog opens.

Illustration 56.2 VPN Settings

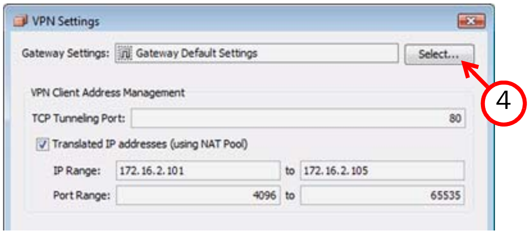

VPN Settings

Gateway Settings: Gateway Default Settings | Select...

VPN Client Address Management

TCP Tunneling Port: 80

☑ Translated IP addresses (using NAT Pool)

IP Range: 172.16.2.101 to 172.16.2.105

Port Range: 4096 to 65535

4. Click **Select** and browse to the correct profile.
5. Click **OK** to close the VPN Settings dialog.
6. Click **OK** to close the Firewall Properties dialog.

What's Next?

▶ Refresh the policy of the firewall to transfer the changes.

CHAPTER 57

VPN CLIENT SETTINGS

Stonesoft IPsec VPN Clients do not have controls for many settings that are needed for establishing a VPN. These settings are defined in the SMC, and the IPsec VPN Client downloads it from the gateways it connects to. VPN clients are only supported in policy-based VPNs.

The following sections are included:

Getting Started With VPN Client Settings

Prerequisites: See Getting Started With IPsec VPNs

For instructions for the installation and daily use of Stonesoft IPsec VPN Clients, see the *Stonesoft IPsec VPN Client Administrator's Guide* and the *Stonesoft IPsec VPN Client User's Guide*.

What VPN Client Settings in the SMC Do

The Stonesoft IPsec VPN Client settings are configured centrally in the Management Center and then automatically updated to the IPsec VPN Clients from the engines when the clients connect. The following settings are transferred from the gateway to the client:

- Routing information (VPN Site definitions): generally, if an IP address that the client wants to contact is included in the Site definition, the traffic is routed into the VPN.
- Authentication settings.
- Encryption settings.
- Information about the gateway's end-points.
- Setting for NAT traversal method allowed.
- Settings for local security checks on the client computer (for Stonesoft IPsec VPN Clients version 5.0 or higher).
- Backup gateways to contact in case there is a disruption at the gateway end (for Stonesoft IPsec VPN Client version 5.1 or higher).

Limitations

- When the Stonesoft IPsec VPN Client is first installed, it has no configuration. The basic information on Gateways (such as the IP address to use for connecting) must be introduced by either the user after installation or the administrator when preparing for the deployment.
- There are version-specific dependencies between Stonesoft IPsec VPN Client and Firewall/VPN software. See the *Release Notes* of the Stonesoft IPsec VPN Client version you intend to use for information on compatibility with your Firewall/VPN gateway's software version.
- The SMC does not create configurations for third-party VPN clients. You must create the configuration through the controls and tools of the third-party VPN client product.

Related Tasks

- ▶ Getting Started With IPsec VPNs (page 918)
- ▶ Configuration 3: Basic VPN for Remote Clients (page 902)
- ▶ Creating Rules for VPN Client Connections in Policy-Based VPNs (page 960)
- ▶ Routing Internet Traffic Through Policy-Based VPNs (page 988)

List of VPN Client Settings in the Management Client

Prerequisites: See Getting Started With IPsec VPNs

The table below outlines the settings for Stonesoft IPsec VPN Clients available in the Management Client.

Table 57.1 IPsec VPN Client Settings in the Management Client

Location	Setting	Explanation
Firewall element properties →Advanced Settings tab →VPN Settings	TCP Tunneling Port	Port used for tunneling IPsec VPN Client connections inside TCP connections to bypass intermediary traffic filters and NAT devices.
	Translated IP address (Using NAT Pool)	Address range for translating IP addresses of incoming IPsec VPN Client connections for internal networks. Alternative to using the Virtual Adapter (next setting below in this table). (See Managing VPN Client IP Addresses (page 1000).)
Internal VPN gateway element properties →IPsec Client Tab (See Managing VPN Client IP Addresses (page 1000).)	Virtual IP address (Using Virtual Adapter)	Options for configuring the IPsec VPN Client with a second, virtual network adapter with a DHCP-assigned IP address for connections inside the VPN. Alternative to using the NAT Pool (previous setting above in this table).
	Firewall Advanced Settings	Shows the VPN Client settings from the properties of the Firewall element for your convenience, since the address space for the NAT Pool must not overlap with the Virtual Adapter address space if both are used.
	Backup Gateways	Backup gateways to contact in case there is a disruption at the gateway end (in the order of contact). This feature requires Stonesoft IPsec VPN Client version 5.1 or higher.

Table 57.1 IPsec VPN Client Settings in the Management Client (Continued)

Location	Setting	Explanation
VPN Profile element properties →IKE SA tab (See Defining IKE SA Settings for a VPN (page 941)	Versions	IKE version(s) used in IKE SA negotiations.
	Cipher Algorithms →AES-128 →AES-256 →AES-CGM →DES →3DES Message Digest Algorithms →SHA-1 →SHA-2 →MD5	The supported algorithms for the current version of Stonesoft IPsec VPN Clients.
	Diffie-Hellman Group	Diffie-Hellman group used in IKE SA negotiations.
	Authentication Method SA Lifetime in Minutes IKEv1 Negotiation Mode	These settings have no effect on IPsec VPN Client connections. See IPsec Client Tab instead.
VPN Profile element properties →IPsec SA tab (See Defining IPsec SA Settings for a VPN (page 943).)	IPsec Type	Only ESP is supported.
	Cipher Algorithms →AES-128 →AES-256 →AES-CGM →DES →3DES Message Digest Algorithms →SHA-1 →SHA-2 →AES-XCBC-MAC →MD5	The supported algorithms for the current version of Stonesoft IPsec VPN Clients. AES-XCBC-MAC requires Stonesoft IPsec VPN Client version 5.0 or higher.*
	Compression Algorithm →Deflate	Not supported in Stonesoft IPsec VPN Client (and Firewall/VPN engine) versions 4.2 and 4.3. Other versions support this option.
	PFS Diffie-Hellman Group	Diffie-Hellman group used in IKE SA negotiations.

Table 57.1 IPsec VPN Client Settings in the Management Client (Continued)

Location	Setting	Explanation
VPN Profile element properties →IPsec Client tab (See Defining VPN Client Settings (page 946).)	Authentication Method	The selected authentication method used with Stonesoft IPsec VPN Clients.
	Allow Hybrid/EAP Authentication	The IPsec VPN Client users authenticate by username and password (or other type of passcode) and the gateway authenticates itself to the client with a certificate.
	Allow CN authentication	Allows authentication using the common name in the certificate as the user name. The CN is checked against a value entered in the User elements.
	Allow Pre-Shared Key Authentication with IKEv1	This setting has no effect on Stonesoft IPsec VPN Client connections, as pre-shared key authentication is not supported.
	IPsec Security Association Granularity	Defines whether SAs are negotiated per network or per each connecting IP address. Only **SA Per Net** is supported by Stonesoft IPsec VPN Clients.
	Local Security Checks	Defines whether the IPsec VPN Client checks for the presence of basic security software to stop connections from risky computers. This feature requires Stonesoft IPsec VPN Client version 5.0 or higher.
VPN element →Tunnels tab	Pre-shared key fields	This setting has no effect on IPsec VPN Client connections. Pre-shared keys for VPN Client connections are defined per-user account in the User elements.

*)The Russian product version has no strong encryption algorithms.

Managing VPN Client IP Addresses

Prerequisites: See Getting Started With IPsec VPNs

There are two different methods to define the IP addresses VPN clients use in the internal network. You must always configure one or the other when you want to create a client-to-gateway VPN for the VPN to be valid. The methods are as follows:

1. You can use NAT to translate the IP addresses in communications, which gives the VPN Clients an ‘internal’ IP address in the internal network without the need for a DHCP server. This is called a *NAT Pool*.
   - This method is not recommended for Stonesoft IPsec VPN Clients, because it does not allow the clients to make queries to internal DNS servers without manual configuration.
   - NAT rules are not applied to communications from clients that receive their address through the NAT Pool feature. The NAT Pool translation is applied before the NAT rules.
   - The NAT Pool method does not require any additional client-side features.
2. (*Recommended for Stonesoft IPsec VPN Clients)* You can use a DHCP server to assign the VPN clients a second, virtual IP address that is used in communications through the VPN tunnel. The IP address is attached to a *Virtual Adapter*. Using this method provides the following benefits over the NAT Pool:
   - Centrally configure the DNS settings for any version of Stonesoft IPsec VPN Clients when connected (using the DHCP server).
   - Control the IP address each VPN Client is assigned (depending on the DHCP server).
   - Forward client-to-gateway VPN traffic to a site-to-site VPN and/or route the Internet traffic from the client computer through the gateway for inspection.
   - Open new connections also from the internal network to the VPN Client computers through the VPN.
   - To use the Virtual Adapter, the VPN client software must support this feature. Not all third-party VPN clients have a Virtual Adapter feature.

The Virtual Adapter is required when there is a need to open connections from the internal network to the VPN client. Activating both the NAT Pool and the Virtual Adapter is technically possible, but the NAT Pool address translation is applied to all VPN client traffic when activated, including connections from hosts that use a Virtual Adapter.

What’s Next?

- ▶ Configuring NAT Pool for VPN Clients (page 1001)
- ▶ Configuring Virtual IP Addressing for VPN Clients (page 1002)

Related Tasks

- ▶ For a detailed technical discussion on using a virtual IP address, see RFC 3456.

Configuring NAT Pool for VPN Clients

The NAT pool defines a range of IP addresses that the firewall can use to translate the source address of connections from VPN clients. The NAT pool translates the addresses in the same way as NAT rules do. Connections that use the NAT Pool must not match any NAT rules.

Note – Make sure NAT is enabled for this VPN. The Enable NAT with this VPN option in the properties of the VPN element must be selected. Otherwise, the NAT pool options have no effect.

▼ To define the NAT Pool

1. Open the properties of the Firewall element.
2. Switch to the **Advanced Settings** tab.
3. Click the **VPN Settings** button. The VPN Settings dialog opens.

Illustration 57.1 VPN Settings

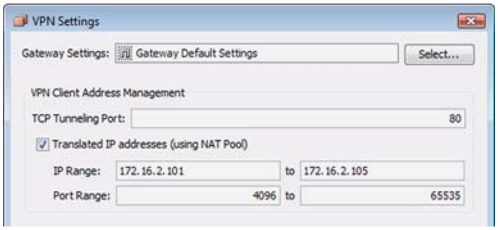

Gateway Settings: Gateway Default Settings

VPN Client Address Management

TCP Tunneling Port: 80

☑ Translated IP addresses (using NAT Pool)

IP Range: 172.16.2.101 to 172.16.2.105

Port Range: 4096 to 65535

4. Select the **Translated IP address (using NAT Pool)** option.

Note – If the NAT Pool is active, it is used for translating connections from VPN clients that have a virtual IP address. It is not possible to exclude hosts with a virtual IP address from being subject to the NAT Pool address translation.

5. Enter the IP addresses and ports you wish to use for translating VPN client traffic in **IP Range** and **Port Range**.

Caution – Make sure the addresses you define here do not overlap with addresses that are in use in your networks. Also, the addresses must not overlap with any translated address space in your NAT rules.

Configuring Virtual IP Addressing for VPN Clients

This feature requires the following:

- You use an external DHCP server to assign the IP addresses.
- The users use a VPN client that has a Virtual Adapter feature. All Stonesoft IPsec VPN Clients always have this feature installed and active.

The Virtual Adapter assigns the VPN client a second IP address for communications through the VPN, independent of the address the VPN client computer uses in its local network.

Most DHCP servers allow a configuration in which a particular client computer is always assigned a particular IP address, for example, based on the MAC address (if VPN clients are set fixed MAC addresses for their Virtual Adapters).

What's Next?

▶ To configure virtual addressing, proceed to Configuring the Gateway for Virtual IP Address Clients.

Configuring the Gateway for Virtual IP Address Clients

▼ To define Virtual IP Addressing for VPN Clients

1. Select **Configuration→Configuration→VPN** from the menu to switch to the VPN Configuration view.
2. Select the **Gateways** branch in the element tree and double-click the Gateway to which VPN clients connect.
3. Switch to the **IPsec Client** tab.

Illustration 57.2 Internal Security Gateway Properties - VPN Client

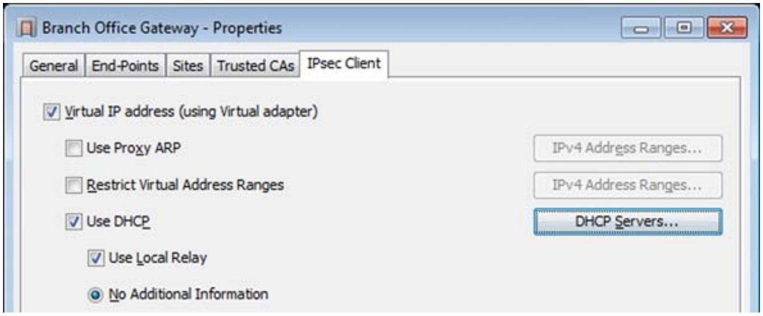

Branch Office Gateway - Properties

General | End-Points | Sites | Trusted CAs | IPsec Client

☑ Virtual IP address (using Virtual adapter)

□ Use Proxy ARP — IPv4 Address Ranges...

□ Restrict Virtual Address Ranges — IPv4 Address Ranges...

☑ Use DHCP — DHCP Servers...

☑ Use Local Relay

● No Additional Information

4. Select the **Virtual IP Address (using Virtual adapter)** option. The other options beneath are enabled.
5. (*Optional*) Select **Use Proxy ARP** to make the Firewall/VPN engine act as a proxy for the VPN client's ARP requests and click the corresponding **Edit** button to select the address range to define the scope for this option.

Note – The Proxy ARP option may be required for a working VPN depending on your network configuration.

6. (*Optional*) Select **Restrict Virtual Address Ranges** and click **Edit** to select IP addresses that the DHCP is allowed to assign.
   - With this option, you can enforce that the VPN clients' addresses stay within a set range, even if the DHCP server tries to assign some other IP address. If an incorrect address is assigned, the user may not be able to access resources.
   - These address ranges must not overlap with the NAT Pool. The NAT Pool configuration is shown for your reference at the bottom of the IPsec Client Tab.

Note – If the NAT Pool is active, it is used for translating connections from VPN clients that have a virtual IP address. It is not possible to exclude hosts with a virtual IP address from being subject to the NAT Pool address translation.

7. Select **Use DHCP** to define DHCP settings for the VPN clients.
8. Click the **DHCP Servers** button and select the correct DHCP Server element.
9. (*Optional*) Select the **Use Local Relay** option to force the use of unicast DHCP relay messages for VPN clients' DHCP requests even if the DHCP server is in a directly connected network in relation to the firewall engine.
   - By default, the Firewall/VPN engine sends a normal DHCP client broadcast message to a DHCP server located in a directly connected network.
   - Unicast DHCP relay messages are always used when sending the DHCP requests to DHCP server that is behind an external gateway device.
10. (*Optional*) Select **Add User information** or **Add Group information** to add the VPN Client user or user group information (in the form: user@domain or group@domain) to the DHCP Request packets' Remote ID option field.
    - Your DHCP server must support the DHCP Relay Agent Information option to use this information.
    - Depending on your DHCP server configuration, this information can be used as a basis for IP address selection.
11. Choose one NDI address in **NDI for DHCP Relay** to be used as the source address for the DHCP packets when querying the DHCP server (the interface towards the DHCP server).
12. If you have a specific need to do so, you can change the largest allowed packet size for the relayed DHCP requests in the **Max Packet Size** drop-down list.

What's Next?

▶ DHCP relay is not allowed in the Firewall Template policy. To allow this traffic, continue the configuration in Allowing DHCP Relay in the Policy (page 1004).

Allowing DHCP Relay in the Policy

DHCP relay is not allowed in the Firewall Template policy. However, there is a ready-made DHCP relay sub-policy that you can use in your own policy.

Note – The sub-policy also contains rules for local DHCP relay between internal networks. If just one of these DHCP relay features is active, the rules for the other feature are invalid and ignored when the policy is installed.

▼ To activate the DHCP relay sub-policy

1. Open the firewall's policy (or a template policy) for editing.
2. Insert a rule that redirects the inspection to the DHCP relay Sub-Policy:
   - Source, Destination, and Service: right-click each cell and select **Set to ANY**.
   - Action: Click the cell and select **Jump** from the list that opens. Select **DHCP Relay** and click **Select**.

What's Next?

▶ Refresh the firewall's policy before using this configuration.

Exporting VPN Client Configuration to a File

Prerequisites: The client-to-gateway VPN must be fully configured, see Getting Started With IPsec VPNs

You can export the settings for Stonesoft IPsec VPN Clients from the Management Center. This allows deploying new clients without requiring the users to add the VPN gateways manually.

▼ To export the Stonesoft IPsec VPN Client configuration

1. Select **Configuration→Configuration→VPN** from the menu to switch to the VPN Configuration view.
2. Select the **Gateways** branch in the element tree and right-click the Gateway to which VPN clients connect.
3. Select **Tools→Save Gateway Contact Information** and save the resulting file.

Generate a separate file for each gateway to which clients directly connect.

What's Next?

▶ Make the configuration available to VPN Clients as explained in the *Stonesoft IPsec VPN Client Administrator's Guide*.

Maintenance and Upgrades

In this section:

CHAPTER 58

BACKING UP AND RESTORING SYSTEM CONFIGURATIONS

Backups contain the necessary configuration information to restore the SMC to the state it was in when the backup was taken, including the configuration information for the Firewall, IPS, and Layer 2 Firewall engines that the Management Server stores.

The following sections are included:

Getting Started with Backups

Prerequisites: None

What Backups Do

Backups allow you to save and restore Management Server and Log Server configurations on the same system or on a different physical host.

- The Management Server backup contains the policies, elements, and other essential configuration details for all Firewall, IPS, and Layer 2 Firewall components that they manage, as well as the configuration information of the Web Portal Server and of the Management Server itself.
- The Log Server backup contains the Log Server's local configuration and optionally the logs.
- The Authentication Server backup only contains information on the user accounts in the user database. Authentication Server backups must be taken separately for each Authentication Server node.

Restoring the backups allows you to restore the SMC configurations to the state they were in when taking the backup, even if you restore it on a completely new installation.

Backups are needed to recover from the loss of the system configurations, for example, due to hardware failure. A backup also allows you to relocate the SMC servers onto different hardware.

Limitations

The private keys of engine certificates are stored locally on the engines and are not backed up.

What Do I Need to Know Before I Begin?

The Management Server is the only place that contains usable, complete configuration information for any individual engine component in the system. The engines contain a working copy of the configuration details that allows them to carry out traffic inspection independently, but it is not possible to extract this information from the engines in the event that the Management Server is lost. Regular Management Server backups are therefore essential and must be stored in a safe storage location outside the Management Server host machine.

Always take the backups using the Management Server's internal backup tool. External backup applications that back up the host server may not produce usable backups of your SMC servers, especially if the SMC servers are running at the time the backup is taken.

Backups from the previous major version of the SMC can always be restored in the current major version of the SMC. Backups taken from older versions may not always be restorable. Generally, backups can be restored between versions that support direct upgrades between the versions. See the Release Notes for version-specific details.

Note – If your configuration contains elements for TLS Inspection, the private keys and certificates of the Server Protection Credentials and Client Protection Certificate Authorities are included as plain text in the Management Server backup. Use the encryption option for the backups when the configuration contains elements for TLS Inspection. For more information, see Setting up TLS Inspection (page 793).

Configuration Overview

1. Back up the Management Server(s), Log Server(s) and Authentication Server regularly as instructed in Creating Backups or schedule backup tasks to run at regular intervals as instructed in Creating Backup Tasks (page 1035) and Scheduling Tasks (page 1040).
2. Store the backup files in a safe location as instructed in Storing Backup Files (page 1010).
3. When necessary, restore a backup as instructed in Restoring Backups (page 1011).

Creating Backups

Prerequisites: None

Management Server backups include all configuration information, including licenses, server components' certificates needed for system communications, the root CA, and locally stored user accounts. The configurations you create in the SMC for the engine components are included in the Management Server backup and do not need separate backups.

Log Server backups contain the Log Server configuration information and optionally the log data stored on the server. There is a configurable limit to how large the Log Server backup can be.

Authentication Server backups must be created separately for each Authentication Server node. Authentication Server backups only contain information about user accounts in the Authentication Server's user database. Configuration information about the Authentication Server is included in the Management Server backup.

The steps below explain how to use the Management Client to take and manage the backups. It is also possible to create backups on the command line. See Command Line Tools (page 1145). The contents of the backup file are the same regardless of the method used.

Note – To back up a Management Server, there must be enough free disk space on the server. Twice the size of the management database is required. If there is not enough available disk space, the backup process does not start.

▼ To create backups

1. Right-click the Management Server or Log Server, or the Authentication Server node you want to back up and select **Backup**. The Backup Task Properties dialog opens.
2. (*Optional*) If you want to back up additional servers, select the server(s) from the list on the left and click **Add**.
3. (*Optional*) If you want to create an encrypted backup, select **Encrypted** and enter and confirm a password.
   - We recommend this option if the configuration contains elements for TLS Inspection.
4. (*Optional*) If you are creating a backup of Log Server(s) and you want to back up the log files, select **Back up Log Files**.
5. Click **OK**. The progress is shown on a new tab.

What's Next?

▶ Copy the backup files from the backup directory to a separate, safe location for storage. See Storing Backup Files (page 1010).

Related Tasks

- If you want to create backup tasks and schedule them to run at regular intervals, see Creating Backup Tasks (page 1035) and Scheduling Tasks (page 1040).
- To back up and delete log data with the log management tools, see Getting Started with Log Data Management (page 1016).

Storing Backup Files

Prerequisites: Creating Backups

The backup files are saved in the *<installation directory>/backups/* directory of the server on which they were created. We recommend copying the backup file to a safe location, for example, to removable media or another host. Otherwise, you will have to manually recreate all configurations if the data on the host computer is irrecoverably lost.

Note – If you installed the Management Server in the `C:\Program Files\Stonesoft\Management Center` **directory in Windows, some program data may be stored in the** `C:\ProgramData\Stonesoft\Management Center` **directory.**

The backups files are compressed to `.zip` files or `.enc` files and they can also be decompressed manually if needed. If necessary, the backups are split into several files to fit the maximum file size. Each backup has its own subdirectory.

Note – Remember to handle the backup files securely, as they contain all the configuration information for the system.

To store backup files

1. Browse to the backup directory on the server on which the backup was created.
   - Backup files are stored in the *<installation directory>/backups* directory or a subdirectory under it.
   - Unencrypted backups are `.zip` files.
   - Encrypted backups are `.enc` files.
2. Copy the backup files to a safe storage location.

Restoring Backups

Prerequisites: Creating Backups

Backups created in one operating system can be restored to an installation running on another operating system without any special measures. This is useful when changing the operating system or hardware platform.

See the upgrade instructions in the Release Notes. If an intermediate upgrade is required between your current version and the newest version, upgrade the existing installation to (at least) the intermediate version to create a working backup.

When you restore a backup, the backup restoration process checks that there is enough disk space on the destination drive. Twice the size of the backup file is required. If there is not enough available disk space, the restoration fails.

It is also possible to restore backups on the command line (see Command Line Tools (page 1145)).

What's Next?

- ▶ To restore a Management Server backup, see Restoring a Management Server Backup.
- ▶ To restore a Log Server backup, see Restoring a Log Server Backup (page 1012).
- ▶ To restore an Authentication Server backup, see Restoring an Authentication Server Backup (page 1013).

Restoring a Management Server Backup

▼ To restore a Management Server backup

1. Check that the backup file is in the `<`*`installation directory`*`>/backups/` directory of the server in question.
   - If you have moved the backup file to a different location, you must first copy it back to the `<`*`installation directory`*`>/backups/` directory.

Note – If you installed the Management Server in the `C:\Program Files\Stonesoft\Management Center` **directory in Windows, some program data may be stored in the** `C:\ProgramData\Stonesoft\Management Center` **directory.**

2. Stop the Management Server service through the operating system's service management feature or using the command line script.
   - If you have trouble stopping the services, disable the automatic startup of SMC services in the operating system and restart the computer.
3. Start the backup restoration script:
   - In Windows, run `<`*`installation directory`*`>/bin/sgRestoreMgtBackup.bat`
   - In Linux, run `<`*`installation directory`*`>/bin/sgRestoreMgtBackup.sh`
4. Select the backup file to be restored. The default Management Server backup file names have the following structure: `sgm_vVERSION.[BUILD]_YYYYMMDD_HHMMSS[comment]`.
5. Type **y** and press enter to confirm the restoration. Encrypted backups require you to enter the password that was used to encrypt the backup when it was created.

If the restore operation fails, the original configuration remains unchanged.

What’s Next?

- If the backup is restored on a system that uses a different IP address than the Management Server that the backup is from, you must complete the relevant steps to change the IP address. See Changing the Management Server IP Address (page 345).
- The backup contains the internal CAs (certificate authorities). If components in the system have certificates from a different CA than the one contained in the backup, the certificates are not accepted as valid after restoring the backup and have to be regenerated as explained in Troubleshooting Certificates (page 1091).
- Otherwise, start the Management Server. If you disabled automatic startup for any SMC services and want the services to start automatically after reboot, start the services and enable automatic startup for them.

Restoring a Log Server Backup

To restore a Log Server backup

1. Check that the backup files are in the `<`*`installation directory`*`>/backups/` directory of the server in question.
   - If you have moved the backup files to a different location, you must first copy them back to the `<`*`installation directory`*`>/backups/` directory.
2. Stop the Log Server service through the operating system’s service management feature or using the command line script.
   - If you have trouble stop the services, disable the automatic startup of SMC services in the operating system and restart the computer.
3. Start the backup restoration script:
   - In Windows, run `<`*`installation directory`*`>/bin/sgRestoreLogBackup.bat`
   - In Linux, run `<`*`installation directory`*`>/bin/sgRestoreLogBackup.sh`
4. Select the backup file to be restored. The default Log Server backup file names have the following structure: `sgl_vVERSION.[BUILD]_YYYYMMDD_HHMMSS[comment]`.
5. Type **y** and press enter to confirm the restoration. Encrypted backups require you to enter the password that was used to encrypt the backup when it was created.

If the restore operation fails, the original configuration remains unchanged. If it is not possible to transfer the logs through a backup, log files can be copied to the Log Server through the operating system like any other files.

What’s Next?

- If you restore the Log Server backup on a computer that has a different IP address than the Log Server that the backup was created with, complete the relevant steps in Changing the Log Server IP Address (page 346).
- Otherwise, restart the Log Server. If you disabled automatic startup for any SMC services and want the services to start automatically after reboot, start the services and enable automatic startup for them.

Restoring an Authentication Server Backup

▼ To restore an Authentication Server backup

1. Check that the backup files are in the `<installation directory>/backups/` directory of the server in question.
   - If you have moved the backup files to a different location, you must first copy them back to the `<installation directory>/backups/` directory.
2. Start the backup restoration script:
   - In Windows, run `<installation directory>/bin/sgRestoreAuthBackup.bat`
   - In Linux, run `<installation directory>/bin/sgRestoreAuthBackup.sh`
3. Select the backup file to be restored.
4. Type **y** and press enter to confirm the restoration. Encrypted backups require you to enter the password that was used to encrypt the backup when it was created.
5. In the Management Client, right-click the Authentication Server and select **Apply Configuration** when the backup has been successfully restored.

Note – The original configuration remains unchanged until you apply the Authentication Server's configuration.

Recovering from a Hardware Failure

Prerequisites: Creating Backups

▼ To restore Management Server configurations on replacement hardware

1. Install the Management Server software (see the *Stonesoft Management Center Installation Guide*). The same exact version is not required for recovery, but all SMC components must run the same version to work together.
2. Restore the Management Server backup as explained in Restoring a Management Server Backup (page 1011).

▼ To restore Log Server configurations on replacement hardware

1. Install the Log Server software, if not installed together with the Management Server software (see the *SMC Installation Guide*). The same exact version is not required for recovery, but all SMC components must run the same version to work together.
2. Restore the Log Server backup as explained in Restoring a Log Server Backup (page 1012).

▼ To restore engine configurations on replacement hardware

1. Generate an initial configuration for the engine in the SMC as explained in Saving an Initial Configuration for Security Engines (page 515).
2. Add the hardware to the network and configure it in the same way as a new installation (see the *SMC Installation Guide* or *Appliance Installation Guide*).

3. When contact with the Management Server is established, install the policy. The full working configuration is transferred to the engine.

Note – In some cases, the IPsec VPN certificate information may be lost and the policy installation fails. If this happens, delete the old IPsec VPN certificates in the Management Client and create new VPN certificates for the engine. When you use the same CA and certificate details, the new certificates are accepted by other components. Policy installation is also possible if you disable the invalid configurations (for example, by disabling all VPN-specific Access rules in the policy). See Creating and Signing VPN Certificates (page 975).

CHAPTER 59

MANAGING LOG DATA

Logs must be actively managed to prevent the Log Server from filling up the hard disk with logs. Management Center tools help you manage the generated log entries automatically according to settings you define.

The following sections are included:

Getting Started with Log Data Management

Prerequisites: None

Any Log Server will gradually fill up completely if log data is never removed. You must actively manage the log data to prevent this from happening.

What Log Management Tools Do

You can manage the log data in the following ways:

- Export log data so that it can be used elsewhere.
- Copy log data to an archive location.
- Delete old or unnecessary log data.
- Set up automatic log management tasks for exporting, copying, and removing selected data.
- Reduce the amount of logs by pruning some of the log entries before they are stored on the Log Server. However, preventing the unwanted log entries from being created should always be preferred over pruning to avoid unnecessary use of resources.

Limitations

Only the logs in the active storage are used in reporting. If you archive logs, you can still view them in the Logs view, but they are no longer available when you generate reports (see Reports (page 167)).

Alert and audit logs cannot be pruned.

Configuration Overview

1. Tune your system to generate the log entries you need. See Defining When Logs Are Generated (page 1017).
2. (*Optional*) Set up log archiving to store older important logs for possible later use and free up the space on the Log Server, see Archiving Log Data (page 1018).
3. (*Recommended*) Set up scheduled log data tasks for deleting logs that are not needed in the long term. See Deleting Log Data (page 1021).
4. (*Optional*) Configure log pruning to prune out any unnecessary logs if any are generated. See Pruning Log Data (page 1024).

Related Tasks

- Changing Log Server Configuration Parameters (page 320)
- Exporting Log Data (page 1026)
- Forwarding Log Data to Syslog (page 323)
- To export and archive log data directly from the Logs view, see Exporting Extracts of Log Data (page 163).

Defining When Logs Are Generated

Prerequisites: None

Normal and Alert logs are generated both based on internal conditions in the operation of a component and based on traffic that the engines handle.

Internal conditions that trigger logs or alerts:

- There is a system error or warning.
- An engine test fails. You can configure the engine tester in detail and select whether test failures trigger an alert as explained in Configuring the Engine Tester (page 519).
- The status of an engine changes (not active by default). See Enabling/Disabling Engine Status Monitoring (page 221).
- When the values of a monitored item exceed a threshold limit in an Overview (not active by default). See Setting Thresholds for Monitored Items (page 109).
- Diagnostics are active on a Firewall engine (not active by default). See Enabling/Disabling Firewall/VPN Diagnostics (page 222).

Traffic conditions that trigger logs and alerts:

- An IPS engine's or a Layer 2 Firewall's limit for the number of times tunneled traffic is rematched has been reached (not active by default). See Configuring Inspection of Tunneled Traffic (page 577).
- Traffic matches a rule in your policy. See Getting Started with Editing the Rules in Policies (page 660).
- Diagnostics are active on an engine (not active by default). See Enabling/Disabling Firewall/VPN Diagnostics (page 222).

Additionally, you can set up Log Servers to receive logs from any devices that can be set up to send syslog, see Monitoring Third-Party Devices (page 131).

In addition to activating and deactivating logging and the features listed above, you can optimize the number of generated logs on the engines in the following ways:

- You can configure log compression for Discard logs for Firewalls, IPS engines, and Layer 2 Firewalls.
- On Firewalls, you can configure log compression also for antispoofing logs.

For more information on compressing logs on engines, see Configuring Default Log Handling Settings (page 580).

What's Next?

- ▶ To archive (and delete) logs, see Archiving Log Data (page 1018).
- ▶ To remove unnecessary logs, see Deleting Log Data (page 1021).

Archiving Log Data

Prerequisites: None

You can set up an Archive Log Task for copying log data from the active storage on the Log or Management Server to some other location. The same task can also delete the log data from the active storage, so that you do not have to set up a separate task for freeing up the space.

By default, the log archive location is on the same disk drive as the active storage. To change the archive folder, see Changing Log Server Configuration Parameters (page 320).

Note – The Archive Log Task copies the existing log files without compression. This enables you to view the archived logs in the Logs view but they are not used in the Reports view when reports are generated.

What's Next?

- Start by Creating an Archive Log Task.

Creating an Archive Log Task

▼ To create an Archive Log Task

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Right-click **Tasks** in the Administration tree and select **New→Archive Log Task** from the menu. The Archive Log Task Properties dialog opens.

Illustration 59.1 Archive Log Task Properties - Genera Tab

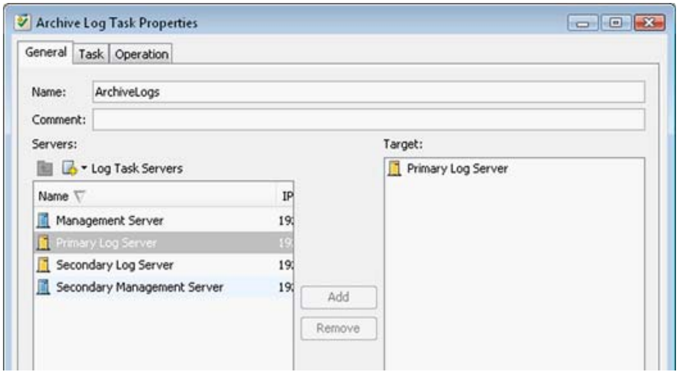

Archive Log Task Properties

Name: ArchiveLogs

Comment:

Servers: Log Task Servers

Name	IP
Management Server	19
Primary Log Server	19
Secondary Log Server	19
Secondary Management Server	19

Target: Primary Log Server

3. Enter a unique **Name** and an optional free-form **Comment** for your own reference.
4. Select the server from which the logs are archived and click **Add.**

What's Next?

- Continue by Selecting Log Data for Archiving (page 1019).

Selecting Log Data for Archiving

▼ To select log data for archiving

1. In the Archive Log Task properties, switch to the **Task** tab.

Illustration 59.2 Archive Log Task Properties - Task Tab

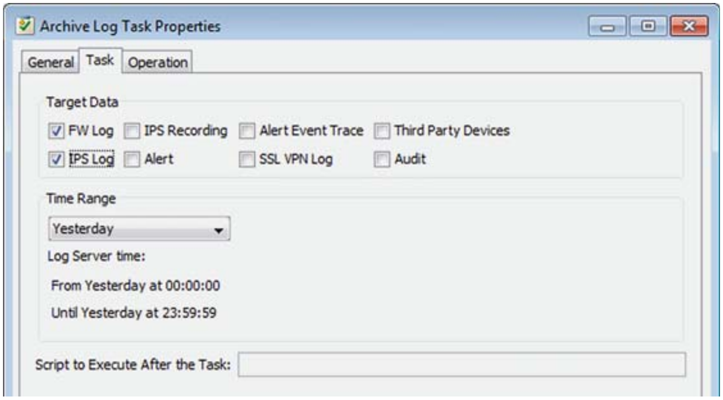

Target Data

☑ FW Log ☐ IPS Recording ☐ Alert Event Trace ☐ Third Party Devices

☑ IPS Log ☐ Alert ☐ SSL VPN Log ☐ Audit

Time Range

Yesterday

Log Server time:

From Yesterday at 00:00:00

Until Yesterday at 23:59:59

Script to Execute After the Task:

2. Select the type of log data for archiving by checking the **Target Data** options.
3. Select the **Time Range** of the log entries.
   - You have several options to limit the time range. For example, select **Absolute Time Range** in the Time Range list and define the Start and End Time.
4. (*Optional*) Specify a script in the **Script to Execute After the Task** field. The Log Server triggers this script after completing the task.

What's Next?

▶ Continue by Selecting Operation Settings for Archiving Log Data (page 1019).

Selecting Operation Settings for Archiving Log Data

▼ To select operation settings for archiving log data

1. In the Archive Log Task properties, switch to the **Operation** tab.

Illustration 59.3 Archive Log Task Properties - Operation Tab

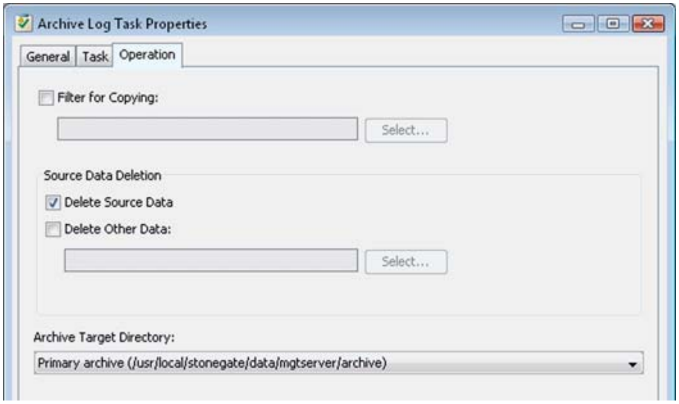

☐ Filter for Copying:

Source Data Deletion

☑ Delete Source Data

☐ Delete Other Data:

Archive Target Directory:

Primary archive (/usr/local/stonegate/data/mgtserver/archive)

2. (*Optional*) Select the **Filter for Copying**. For instructions on how to define filters, see Filtering Data (page 183).

3. In the **Source Data Deletion** section, select whether the data to be archived is kept in active storage after it has been copied to the target location. You can also select that the task removes some other data from the active storage (for example, the archive operation can copy important logs from within the time range and then clear the active storage of all logs within the time range).
   - If you want to delete the archived data from the active storage, select the **Delete Source Data** option.
   - If you want to delete some other data from active storage while you archive data, select the **Delete Other Data** option and optionally a log filter for deleting the other data.

4. Select the **Archive Target Directory** from the list. The directory is determined in the Log Server's configuration file (LogServerConfiguration.txt). See Changing Log Server Configuration Parameters (page 320) for details.

5. Click **OK**.

The task appears under Task Definitions in the Tasks branch of the Administration tree. You can run the task either manually or according to a fixed schedule.

What's Next?

▶ To run the task manually, see Starting Tasks Manually (page 1040).

▶ To make the task run automatically, see Scheduling Tasks (page 1040).

Related Tasks

▶ Viewing Logs From Specific Servers and Archive Folders (page 158)

Deleting Log Data

Prerequisites: None

To permanently remove generated log data, you can delete it from the active storage or delete it as it arrives to the Log Server using pruning filters.

What's Next?

- To delete stored log data, proceed to Creating a Delete Log Task.
- To delete log data as it arrives on the Log Server, see Pruning Log Data (page 1024)

Creating a Delete Log Task

The recommended way to delete logs is to set up a Delete Log Task. In an emergency, you can also delete some of the log files from the Log Server (the default location for logs is in the `<installation directory>/data/storage` folder).

To create a Delete Log Task

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Right-click **Tasks** in the Administration tree and select **New→Delete Log Task**. The Delete Log Task Properties dialog opens.

Illustration 59.4 Delete Log Task Properties - General Tab

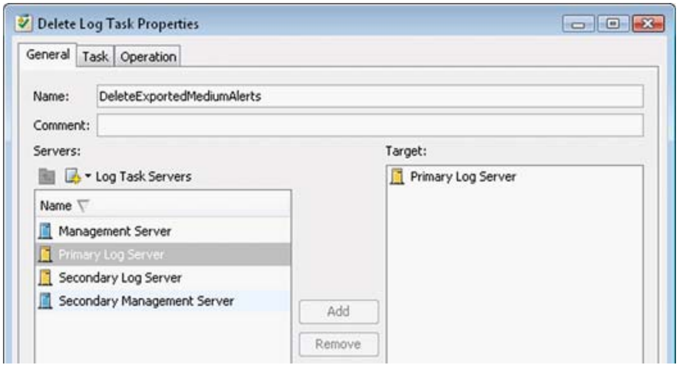

Delete Log Task Properties

General | Task | Operation

Name: DeleteExportedMediumAlerts

Comment:

Servers: Log Task Servers

Name
Management Server
Primary Log Server
Secondary Log Server
Secondary Management Server

Target: Primary Log Server

3. Enter a unique **Name** and an optional free-form **Comment** for your own reference.
4. Select the server from which the logs are deleted and click **Add**.

What's Next?

- Continue by Selecting Data for Deleting Logs.

Selecting Data for Deleting Logs

▼ To select data for deletion

1. In the Delete Log Task properties, switch to the **Task** tab.

Illustration 59.5 Delete Log Task Properties - Task Tab

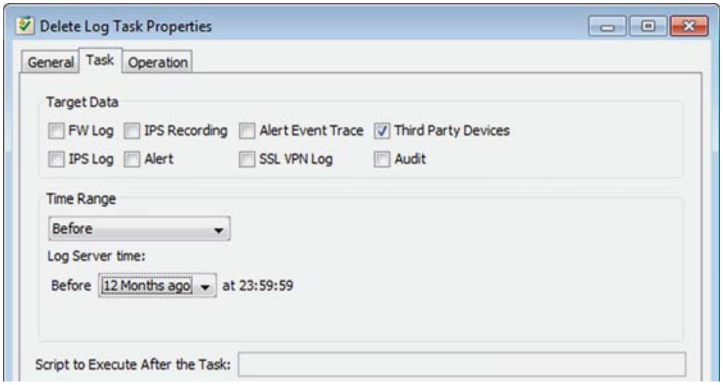

Delete Log Task Properties

Target Data

□ FW Log □ IPS Recording □ Alert Event Trace ☑ Third Party Devices

□ IPS Log □ Alert □ SSL VPN Log □ Audit

Time Range

Before

Log Server time:

Before 12 Months ago at 23:59:59

Script to Execute After the Task:

2. Select the type of log data for deletion by checking the **Target Data** options.
3. Select the **Time Range** of the log entries.
   - You have several options to limit the time range. For example, select **Absolute Time Range** in the Time Range list and define the Start and End Time.
4. (*Optional*) Specify a script in the **Script to Execute After the Task** field. The Log Server triggers this script after completing the task.

What's Next?

▶ Continue by Selecting Operation Settings for Deleting Logs (page 1023).

Selecting Operation Settings for Deleting Logs

▼ To select operation settings for deleting logs

1. In the Delete Log Task properties, switch to the **Operation** tab.

Illustration 59.6 Delete Log Task Properties - Operation Tab

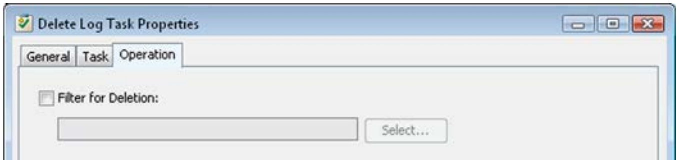

Delete Log Task Properties

General | Task | Operation

□ Filter for Deletion:

2. (*Optional*) Select a log filter that defines which log entries are deleted. For instructions on how to define filters, see Filtering Data (page 183).
3. Click **OK**.

The task appears under Task Definitions in the Tasks branch of the Administration tree. You can run the task either manually or according to a fixed schedule.

Caution – When this task is started (either manually or as you schedule it), all logs matching the selected filter and time range are permanently deleted from the active storage. Make sure that the data you want to keep is exported and/or copied to a safe location before the operation is started.

What's Next?

- To run the task manually, see Starting Tasks Manually (page 1040).
- To make the task run automatically, see Scheduling Tasks (page 1040).

Related Tasks

- Exporting Log Data (page 1026)
- Archiving Log Data (page 1018)

Pruning Log Data

Log data pruning allows you to discard some of the generated logs according to detailed filtering criteria you set. You can prune Log Entries as soon as they arrive on the Log Server or before they are stored (allowing you to still view them in the Current logs view).

Use pruning with caution. The data is deleted without leaving any traces and there is no way to recover incorrectly pruned entries. Pruning also wastes resources compared to preventing the entries from being generated, since pruned entries still have to be created and transferred to the Log Server, and the Log Server still has to process them. To prevent log entries from being generated, see Defining When Logs Are Generated (page 1017).

You can prune normal log entries; alert and audit entries cannot be pruned. The logs are pruned using log filters. If you need more information on Log Filters, see Filtering Data (page 183).

Note – Pruning has no effect on logs that have already been stored. See Deleting Log Data (page 1021) on how to delete stored data.

▼ To prune log data

1. Select **Configuration→Log Data Pruning**. The Log Pruning view opens.

Illustration 59.7 Log Data Pruning

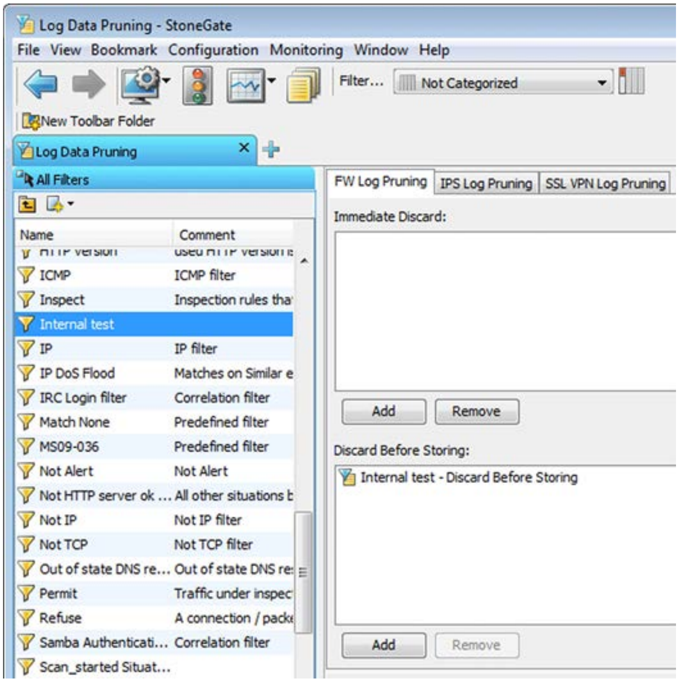

Name	Comment
HTTP version	used HTTP version i
ICMP	ICMP filter
Inspect	Inspection rules tha
Internal test	
IP	IP filter
IP DoS Flood	Matches on Similar e
IRC Login filter	Correlation filter
Match None	Predefined filter
MS09-036	Predefined filter
Not Alert	Not Alert
Not HTTP server ok ...	All other situations b
Not IP	Not IP filter
Not TCP	Not TCP filter
Out of state DNS re...	Out of state DNS re
Permit	Traffic under inspec
Refuse	A connection / pack
Samba Authenticati...	Correlation filter
Scan_started Situat...	

Internal test - Discard Before Storing

2. Select the tab of the log data type you want to prune.

3. Select a log filter for pruning log data from the tree in the Resources panel. If you add several Log Filters, they are combined with logical OR (each filter is matched individually, so all logs that match any of the selected filters are deleted).
   - You can also create a new Filter through the New icon above the list of Filters. For instructions on how to create Filters, see Filtering Data (page 183).

Caution – Never select the Match All filter for log pruning. Pruning with the Match All filter irreversibly destroys all new logs that are created.

4. Activate the pruning for the correct stage:
   - Click **Add** below the **Immediate Discard** field to prune logs before they are even shown in the Current Logs view.
   - Click **Add** below the **Discard Before Storing** field to show the log entries in the Current Logs view and then delete them before they are permanently stored.

Caution – Any log entry that matches the filter you have selected is irrevocably deleted. The changes you make to pruning filters are applied immediately.

5. A warning message is displayed. Click **Yes** to prune the selected log entries. This change is applied immediately without any further action.

Related Tasks

▶ Disabling Pruning Filters (page 1025)

▶ Deleting Log Data (page 1021)

Disabling Pruning Filters

Pruning filters are disabled by removing them from the Log Data Pruning panels. An empty panel means that no logs are pruned.

▼ To remove filters from log pruning

1. Select **Configuration→Log Data Pruning**. The Log Data Pruning view opens.
2. Select the correct tab according to log data type.
3. Select the filter you want to remove from pruning in the **Immediate Discard** or **Discard Before Storing** field.
4. Click the **Remove** button below the field that contains the filter.

Note – Log Filters that are removed from pruning remain available for other use until you delete them separately.

5. Click **OK** in the dialog that is displayed to affirm this action. This change is applied immediately without any further action.

Exporting Log Data

Prerequisites: None

Following the steps here, you can set up a Log Export task, which can be used in manual and scheduled export operations to export a specific log data set from the Log Server and the Management Server log storages.

You can also export extracts of log data while browsing logs, see Exporting Data from the Logs View (page 163).

To send log data to external monitoring products, see Forwarding Log Data to Syslog (page 323).

What's Next?

▶ Start by Creating an Export Log Task (page 1026).

Creating an Export Log Task

▼ To create an Export Log Task

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Right-click **Tasks** in the Administration tree and select **New→Export Log Task**. The Export Log Task Properties dialog opens.

Illustration 59.8 Export Log Task Properties - General Tab

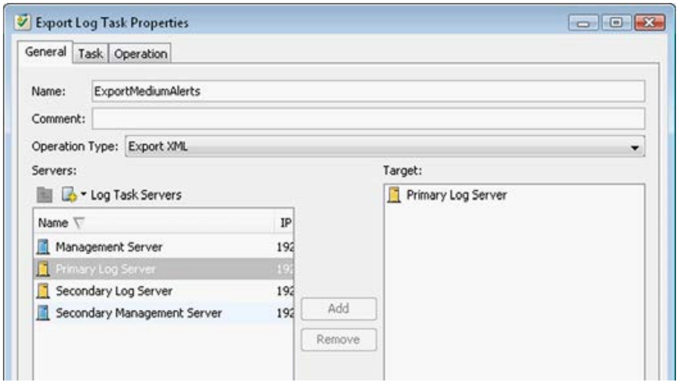

Name: ExportMediumAlerts

Comment:

Operation Type: Export XML

Servers: Log Task Servers

Name	IP
Management Server	192
Primary Log Server	192
Secondary Log Server	192
Secondary Management Server	192

Target: Primary Log Server

3. Enter a unique **Name** and an optional free-form **Comment** for your own reference.
4. Set the **Operation Type** to select the format for the exported data:
   - **Export XML**: logs are exported in XML format
   - **Export CSV**: logs are exported in CSV (comma-separated value) format
   - **Export CEF**: logs are exported in CEF format.
   - (*For exporting IPS traffic recordings*) **Export IPS Recordings as PCAP**: IPS recording file is converted to PCAP format that is compatible with “sniffer” tools such as tcpdump, WinDump, or Wireshark.
   - (*For exporting IPS traffic recordings*) **Export IPS Recordings as Snoop**: IPS recording file is converted to Snoop format that is compatible with “sniffer” tools such as Wireshark.

5. Select the server from which the log data is exported and click **Add**.

What's Next?

- Continue by Selecting Data for Log Export.

Selecting Data for Log Export

To select data for log export

1. In the Export Log Task Properties, switch to the **Task** tab.

Illustration 59.9 Export Log Task Properties - Task Tab

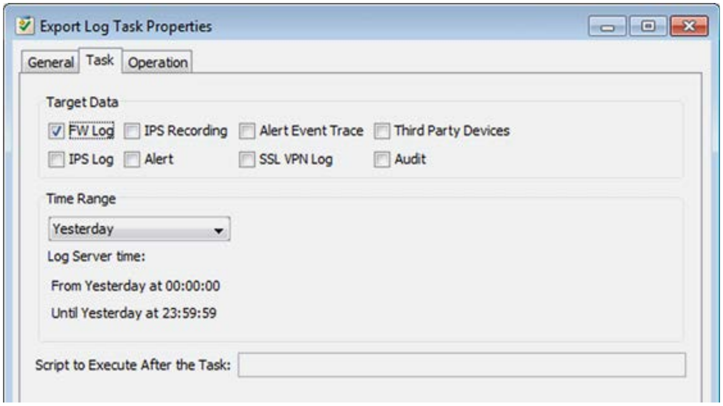

Export Log Task Properties

General | Task | Operation

Target Data

☑ FW Log ☐ IPS Recording ☐ Alert Event Trace ☐ Third Party Devices
☐ IPS Log ☐ Alert ☐ SSL VPN Log ☐ Audit

Time Range

Yesterday

Log Server time:

From Yesterday at 00:00:00

Until Yesterday at 23:59:59

Script to Execute After the Task:

2. Select the type of log data for export by checking the **Target Data** options.
   - If you selected **Export IPS Recordings as PCAP** or **Export IPS Recordings as Snoop** as the Operation Type on the General tab, **IPS Recording** is automatically selected.
3. Select the **Time Range** of the log entries.
   - You have several options to limit the time range. For example, select **Absolute Time Range** in the Time Range list and define the Start and End Time.
4. (*Optional*) Specify a script in the **Script to Execute After the Task** field. The Log Server triggers this script after completing the task.

What's Next?

- Continue by Selecting Operation Settings for Log Export (page 1028).

Selecting Operation Settings for Log Export

▼ To select operation settings for log export

1. In the Export Log Task Properties, switch to the **Operation** tab.

Illustration 59.10 Export XML Task - Operation Tab

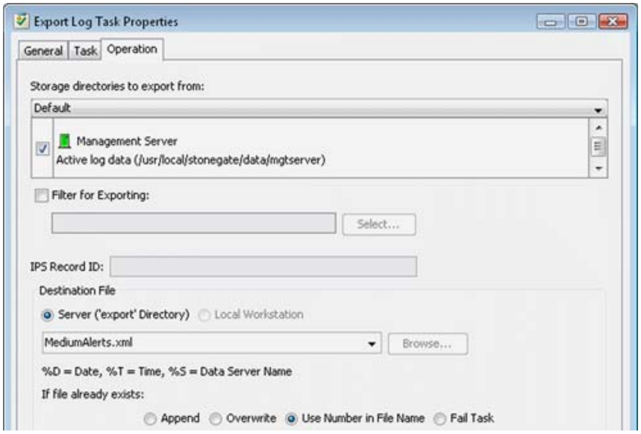

Storage directories to export from: Default

☑ Management Server
Active log data (/usr/local/stonegate/data/mgtserver)

□ Filter for Exporting:

IPS Record ID:

Destination File

● Server ('export' Directory) ○ Local Workstation

MediumAlerts.xml

%D = Date, %T = Time, %S = Data Server Name

If file already exists:

○ Append ○ Overwrite ● Use Number in File Name ○ Fail Task

2. Select the **Storage Directories to Export From**:
   - **Default:** the directory where active log data is stored on the selected server.
   - **Primary Archive:** the directory where archived log data is stored.
   - **Custom**: allows you to select one or more directories from the list.
3. (*Optional*) Select the **Filter for Exporting**. For instructions on how to define Filters, see Filtering Data (page 183).
4. (*Optional; IPS traffic recordings only*) Enter the **IPS Record ID** of the traffic capture.
5. Define the destination file location:
   - **Server**: file is exported to the `export` directory on the server.
   - **Local Workstation:** file is exported to the workstation you are using.
6. Specify a name for the destination file (for example, "*Mediumalerts.xml*") or click **Browse** to browse for an existing export file.
7. In **If file already exists**, specify what happens when a previous file with the same name exists in the same folder:
   - **Append**: the new data is inserted at the end of the existing file.
   - **Overwrite**: the previous file is replaced with the new export file.
   - **Use Number in File Name**: a number is added to the end of the new file's name.
   - **Fail Task**: the operation is cancelled.
8. Click **OK**.

The task appears under Task Definitions in the Tasks branch of the Administration view. If you want to delete the data after exporting, set up a task for that as explained in Deleting Log Data (page 1021).

What's Next?

- ▶ To run the task manually, see Starting Tasks Manually (page 1040).
- ▶ To make the task run automatically, see Scheduling Tasks (page 1040).

Related Tasks

- Archiving Log Data (page 1018)
- Deleting Log Data (page 1021)

Viewing a History of Executed Log Tasks

Prerequisites: Exporting Log Data / Archiving Log Data / Deleting Log Data

There is a history file from which it is possible to view all previously executed tasks related to logs (Export Logs Tasks, Archive Log Tasks, and Delete Log Tasks). The system never erases this file.

▼ To view previously executed Tasks

➥ Open the `<`*`installation directory`*`>/data/logfile.txt` file in a text editor to view the previously executed tasks.

Related Tasks

- Getting Started with Log Data Management (page 1016)

Chapter 60

Managing and Scheduling Tasks

Tasks define parameters of system maintenance operations. You can run maintenance operations manually or automatically according to a schedule you set.

The following sections are included:

Getting Started with Tasks

Prerequisites: None

What Tasks Do

With Task elements, you can start maintenance operations either manually or according to a schedule. You can do the following with Tasks:

- Back up the Management Server(s), Log Server(s), and the Authentication Server.
- Upload policies.
- Upgrade engine software remotely.
- Export, archive, and delete logs.

There are also certain predefined system tasks that run automatically in your system, and which you can see when you work with Task elements.

Scheduling the Tasks allows you to run regular or one-time maintenance operations automatically, for example, during a regular maintenance window.

What Do I Need to Know Before I Begin?

When scheduling automatic backups, you may want the data to be moved to a safe place automatically. This can be achieved through operating system scripts, which Tasks can launch automatically upon completion. With Log Servers, you can change the backup and log archive locations in the Log Server's local configuration file.

See Task Types (page 1033) for short descriptions of all Task types.

Configuration Overview

1. Define the Task parameters as explained in Creating New Task Definitions (page 1035).
2. (*Optional*) Set up automatic Task execution as explained in Scheduling Tasks (page 1040).
3. Run Tasks when necessary as described in Starting Tasks Manually (page 1040).

Related Tasks

- ▶ Pausing the Scheduled Execution of a Task (page 1041)
- ▶ Canceling a Task Schedule (page 1041)
- ▶ Stopping Task Execution (page 1041)

Task Types

Prerequisites: None

Tasks are based on Task Definitions. There are two kinds of Task Definitions: custom Task Definitions and predefined System Task Definitions. To view Task Definitions or to create new Task Definitions, browse to **Tasks→Definition** in the Administration Configuration view.

The table below explains the types of custom Tasks that you can create.

Table 60.1 Custom Task Definitions

Task Definition	Explanation
Backup Task	Creates backup files for the selected Management Server(s), Log Server(s) and/or the Authentication Server. See Creating Backup Tasks (page 1035).
sgInfo Task	Creates a .zip file that contains copies of configuration files and system trace files for the selected component(s) for Stonesoft's technical support. See Creating sgInfo Tasks (page 1039).
Refresh Policy Task	Refreshes the currently installed policy on the selected engine(s). See Creating Refresh Policy Tasks (page 1036).
Upload Policy Task	Uploads the selected policy to the selected engine(s). See Creating Upload Policy Tasks (page 1037).
Remote Upgrade Task	Remotely upgrades the software on the selected engine(s). See Creating Remote Upgrade Tasks (page 1038).
Export Log Task	Copies log data from the active storage or archive to the selected location. See Exporting Log Data (page 1026).
Archive Log Task	Copies log data from the active storage to the selected location. See Archiving Log Data (page 1018).
Delete Log Task	Deletes log data from the active storage. See Deleting Log Data (page 1021).

In addition to Task Definitions that you create and customize, there are predefined Task Definitions for several system tasks. You can run the System Tasks manually or reschedule them, but you cannot change the options in System Task Definitions.

Table 60.2 System Task Definitions

Task Definition	Explanation
Create Snapshot of All System Elements	Automatically creates a snapshot of all system elements after an update package has been activated. The snapshot information is used when administrators compare policy snapshots.
Delete Old Executed Tasks	Deletes old entries from the list of Executed Tasks.
Delete Old Snapshots	Deletes old Policy Snapshots. For more information about Policy Snapshots, see Checking and Comparing Policy Versions (page 654).

Table 60.2 System Task Definitions (Continued)

Task Definition	Explanation
Disable Unused Administrator Accounts	Disables the accounts of administrators who have not been active within the time period defined in the `SGConfiguration.txt` file. The accounts are disabled if the Enforce Password Settings option is enabled. An Administrator with unrestricted privileges can re-enable the disabled accounts. See Defining Password and Login Settings for Administrators (page 259).
Refresh Certificate Revocation Lists	Connects to the CRL servers defined in Trusted Certificate Authority elements (see Defining Trusted Certificate Authorities for TLS Inspection (page 800)) and refreshes the certificate revocation lists on engines.
Renew Gateway Certificates	Generates new certificates for internal gateways if automatic certificate renewal is enabled in the internal gateways' properties.
Renew Internal Certificate Authorities	Checks the status of Internal Certificate Authorities for automatic renewal. To ensure that the automatic certificate authority renewal works correctly, do not change the schedule of this Task.
Renew Internal Certificates	Checks the status of internal certificates for automatic renewal. To ensure that the automatic certificate renewal works correctly, do not change the schedule of this Task.
Upload Gateway and Certificate Authority Certificates	Uploads new IPsec VPN certificates to the Firewall/VPN engines.

What’s Next?

- ▶ Creating New Task Definitions (page 1035).

Creating New Task Definitions

Prerequisites: None

You can create new Tasks to help you maintain your system. The available Task Definitions are explained in Task Types (page 1033). Tasks for managing logs (for example, exporting, archiving, or deleting logs) are Domain-specific. It is unclear at this point if such log data Tasks will work across all Domains. For log data Tasks, see Getting Started with Log Data Management (page 1016).

What's Next?

- ▶ Creating Backup Tasks (page 1035)
- ▶ Creating Upload Policy Tasks (page 1037)
- ▶ Creating Refresh Policy Tasks (page 1036)
- ▶ Creating Remote Upgrade Tasks (page 1038)
- ▶ Creating sgInfo Tasks (page 1039)

Creating Backup Tasks

▼ To create a Backup Task

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Right-click **Tasks** and select **New→Backup Task**. The Backup Task Properties dialog opens.
3. Give the Task a descriptive **Name** and optionally a free-form **Comment**.
4. Select the server(s) you want to back up from the list on the left and click **Add**. The selected server(s) are added to the list on the right.
5. (*Optional*) Write a **Backup Comment**, which is shown in the Management Client and added to the backup file name.
6. (*Optional*) If you want to create an encrypted backup, select **Encrypted**, and enter and confirm a password.
7. (*Optional*) If you are backing up Log Server(s), select the **Back up Log Files** option to back up the logs in addition to the server's configuration.
8. Click **OK**. The new Backup Task is added to the list of Task Definitions.

What's Next?

- ▶ To make the Task run automatically, proceed to Scheduling Tasks (page 1040).
- ▶ To start the Task now, proceed to Starting Tasks Manually (page 1040).

Creating Refresh Policy Tasks

▼ To create a Refresh Policy Task

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Right-click **Tasks** and select **New→Refresh Policy Task**. The Refresh Policy Task Properties dialog opens.
3. Give the Task a descriptive **Name** and optionally a free-form **Comment**.
4. Select the engine(s) on which you want to refresh the policy from the list on the left and click **Add**. The selected engine(s) are added to the list on the right.

> **Note – If you want already established connections to continue using the same configuration information (such as NAT rules), make sure Keep Previous Configuration Definitions is selected.**

5. (*Optional*) Leave **Validate Policy before Upload** selected if you want to validate the rules when the Task is launched manually, and select the related settings. See Validating Rules Automatically (page 714) for more information.
   - Policy validation is possible if you select only one element for the Refresh Policy Task.
6. (*Optional*) Add an **Upload Comment** that is shown in Policy Snapshots created from the policy installations.
7. Click **OK**. The new Refresh Policy Task is added to the list of Task Definitions.

What's Next?

- ▶ To make the Task run automatically, proceed to Scheduling Tasks (page 1040).
- ▶ To start the Task now, proceed to Starting Tasks Manually (page 1040).

Related Tasks

- ▶ Creating Upload Policy Tasks

Creating Upload Policy Tasks

▼ To create a Upload Policy Task

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Right-click **Tasks** and select **New→Upload Policy Task**. The Upload Policy Task Properties dialog opens.
3. Give the Task a descriptive **Name** and optionally a free-form **Comment**.
4. Select the engine(s) to which you want to upload the policy from the list on the left and click **Add**. The selected engine(s) are added to the list on the right.
5. Click the **Select** button next to the **Policy** field. A list of policies opens.
6. Select the policy you want to upload and click **Select**.

> Note – If you want already established connections to continue using the same configuration information (such as NAT rules), make sure **Keep Previous Configuration Definitions** is selected.

7. (*Optional*) Leave **Validate Policy before Upload** selected if you want to validate the rules when the Task is launched manually, and select the related settings. See Validating Rules Automatically (page 714) for more information.
   - Policy validation is possible if you select only one element for the Policy Refresh Task.
8. (*Optional*) Add an **Upload Comment** that is shown in Policy Snapshots created from the policy installations.
9. Click **OK**. The new Policy Upload Task is added to the list of Task Definitions.

What's Next?

▶ To make the Task run automatically, proceed to Scheduling Tasks (page 1040).

▶ To start the Task now, proceed to Starting Tasks Manually (page 1040).

Related Tasks

▶ Creating Refresh Policy Tasks (page 1036)

Creating Remote Upgrade Tasks

▼ To create a Remote Upgrade Task

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Right-click **Tasks** and select **New→Remote Upgrade Task**. The Remote Upgrade Task Properties dialog opens.
3. Give the Task a descriptive **Name** and optionally a free-form **Comment**.
4. Select the upgrade **Operation**:
   - **Remote Upgrade (transfer + activate)** loads the new configuration and reboots the node.
   - **Remote Upgrade (transfer)** loads the new configuration without rebooting the node.
   - **Remote Upgrade (activate)** reboots the node to activate a previously loaded configuration.

Caution – Do not activate the new configuration simultaneously on all the nodes of a cluster. If you want to schedule a Remote Upgrade Task for several nodes, create two separate Remote Upgrade Tasks: one to transfer the new configuration and another to activate it. Schedule the Activate Task to run only after the Transfer Task is complete. See Scheduling Tasks (page 1040).

5. Select the engine(s) that you want to upgrade from the list on the left and click **Add**. The selected engine(s) are added to the list on the right.
6. Select the correct previously imported **Engine Upgrade** file for the upgrade.
7. Click **OK**. The new Remote Upgrade Task is added to the list of Task Definitions.

What's Next?

▶ To make the Task run automatically, proceed to Scheduling Tasks (page 1040).

▶ To start the Task now, proceed to Starting Tasks Manually (page 1040).

Creating sgInfo Tasks

The sgInfo Task collects information about your system for Stonesoft's technical support personnel.

▼ To create an sgInfo Task

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Right-click **Tasks** and select **New→sgInfo Task**. The sgInfo Task Properties dialog opens.
3. Give the Task a descriptive **Name** and optionally a free-form **Comment**.
4. Select the engine(s) whose configuration files and system trace files you want to get from the list on the left and click **Add**. The selected engine(s) are added to the list on the right.
5. Select **Include Core Files** to also include the core files for troubleshooting, unless instructed otherwise.
6. Click **OK**. The new sgInfo Task is added to the list of Task Definitions.

What's Next?

- To make the Task run automatically, proceed to Scheduling Tasks.
- To start the Task now, proceed to Starting Tasks Manually (page 1040).

Scheduling Tasks

Prerequisites: Creating New Task Definitions

After creating Task Definitions, you can schedule Tasks to run at a convenient time. If necessary, you can also schedule System Tasks. See Task Types (page 1033).

▼ To schedule a Task

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Browse to **Tasks→Task Definition**. A list of defined Tasks opens.
3. Right-click the Task and select **Schedule**. The Task Schedule Properties dialog opens.
4. Set the schedule properties:
   - **Repeat**: how often the Task is repeated. If you select **No**, the Task is executed only once (at the specified time).
   - **Start at**: the date and time when the Task starts its first scheduled run.

Tip – The date and time can be entered manually in the format *YYYY-MM-DD HH:MM:SS*. You can also right-click the Up or Down arrows next to the date field to select a date from the calendar.

5. Select the **Final Action**:
   - **Send Alert Only if Task Fails**: An alert is sent only if the Task fails to complete.
   - **Always Send Alert**: An alert is sent regardless of whether the Task fails or succeeds.
   - **Do Not Notify**: No notification is sent.
6. Click **OK**. The schedule is added under the Task Definition.

Related Tasks

- ▶ Task Types (page 1033)
- ▶ Creating New Task Definitions (page 1035)
- ▶ Starting Tasks Manually
- ▶ Pausing the Scheduled Execution of a Task (page 1041)
- ▶ Canceling a Task Schedule (page 1041)

Starting Tasks Manually

Prerequisites: Creating New Task Definitions

If you want a Task to run immediately, you can start it from the Task Definitions list.

▼ To start a Task

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Browse to **Tasks→Task Definitions**. A list of Task Definitions opens.
3. Right-click the Task you want to start and select **Start**. The Task starts.
   - To view the Task, make sure the **Show Executed Tasks** option is enabled (**Tools→Show Executed Tasks**).

Pausing the Scheduled Execution of a Task

Prerequisites: Scheduling Tasks

If you want to temporarily stop a Scheduled Task from running at the scheduled time, you can suspend the Scheduled Task. When a Scheduled Task is suspended, the schedule remains under the Task Definition, but the Task does not run at the scheduled time.

▼ To suspend a Scheduled Task

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Browse to **Tasks→Definition**. The Task Definitions list opens.
3. Expand the Task Definition you want to suspend. The schedule information for the Task is displayed below the Task Definition.
4. Right-click the schedule information and select **Suspend**. The Task is suspended.
5. To restart a suspended Task, right-click the schedule information and select **Continue**. The Task resumes and runs at the next scheduled time.

Canceling a Task Schedule

Prerequisites: Scheduling Tasks

You can remove Tasks from the schedule by marking the schedule information from the Task Definition as obsolete. Marking the schedule information as obsolete does not delete the Task Definition: the same Task can be scheduled again.

▼ To remove a Task from the schedule

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Browse to **Tasks→Definition**. The Task Definitions list opens.
3. Expand the Task you want to remove from the schedule. The schedule information for the Task is displayed below the Task Definition.
4. Right-click the schedule information and select **Mark as Obsolete**. A confirmation dialog opens.
   - The schedules of default System Tasks cannot be deleted.
5. Click **Yes** to confirm that you want to mark the selected Task as obsolete. The Task is removed from the Scheduled Tasks list.

Stopping Task Execution

Prerequisites: Creating New Task Definitions, Scheduling Tasks, Starting Tasks Manually

▼ To abort a running Task

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Browse to **Tasks→History**.
3. Right-click the Task you want to abort and select **Abort**. A confirmation dialog opens.
4. Click **Yes** to stop the running Task.

CHAPTER 61

MANAGING LICENSES

All system components must be licensed as a proof of purchase. In addition, some additional features can be activated by importing a feature license.

The following sections are included:

Getting Started with Licenses

Prerequisites: None

Licenses prove that your organization has legally purchased the Stonesoft components you configure. Licenses are issued when you purchase a product and you can upgrade them to new supported versions as part of each component's support and maintenance contract.

Licenses are generated by Stonesoft's online licensing servers and imported into the system as files. Licenses are shown as elements in the Administration Configuration view.

When Do I Have to Generate Licenses?

Generally, you must generate and import a license for each engine and each Management Center server to start using the component. However, there are the following exceptions:

- All Management Servers in the same Management Center share a single license. A high availability setup has multiple Management Servers.
- All currently available Security Engine appliance models can fetch a license automatically through the Management Server if automatic updates are enabled. If automatic licensing fails, the appliances have a 30-day temporary license to allow time for manual licensing.
- SSL VPN licenses can be managed in the SSL VPN Administrator interface or through the Management Client.
- Some optional features are activated by purchasing and importing a feature-specific license.

What Do the Different License Types Mean?

Licenses can be bound to a component in three different ways. The possible binding methods depend on the licensed component and the software version.

Table 61.1 License Binding Methods

License Binding	Description
IP address binding	The license is statically bound to the IP address of the licensed component. **Note!** Starting from SMC version 5.0, only licenses for Management Center servers can be bound to an IP address. Existing IP-address-bound licenses for other components continue to work and can be upgraded. Any new licenses for other components must be bound to the Management Server's proof-of-license (POL) code.
Management Server proof-of-license (POL) code binding	Licenses are dynamically bound to the Management Server's proof-of-license (POL) code. You must manually bind Management Server POL-bound licenses to the correct element. Licenses are valid only for components that are managed by the Management Server with the same POL code.
Appliance proof-of-serial (POS) code binding	The license is bound to the unique POS code of a preinstalled Stonesoft appliance. The appliance identifies itself when contacting the Management Server. The Management Server allows the use of the appliance if the license POS code matches the reported code. The Management Server automatically binds the correct license to the engine element based on the POS code. For Management Server and preinstalled appliances with engine version 5.0 and higher, the Management Server can use this licensing method automatically with new appliances.

The license types that are available depend on the system component:

Table 61.2 Possible License Types for System Components

Component	License Type
Management Servers	Always a static IP-address-bound license.
Log Servers	A static IP-address-bound license for servers prior to version 5.0 or a dynamic license bound to the Management Server's POL code.
Web Portal Servers	
Authentication Servers	
Preinstalled Stonesoft appliances	A license bound to the POS code of the appliance (all current models) or a dynamic license bound to the Management Server's POL code (older models).
Engines installed on your own hardware	Always a dynamic license bound to the Management Server's POL code.
Engines installed on a virtualization platform	Always a dynamic license bound to the Management Server's POL code.
SSL VPN Gateway	A special dynamic license bound to DNS domains and/or IP addresses. The license can be bound to several DNS domains, and all DNS domains that you use must be bound to the same license.
Feature-specific licenses	A dynamic license bound to the Management Server's POL code or a license bound to the POS code of the appliance depending on the feature.

When Do I Have to Update Licenses?

Components do not run without a valid license. Always make sure you have an updated license before you make any change that is not supported by the current license. Licenses must be updated for new software versions and if the binding detail in the license changes:

- Software upgrades: No action is required if automatic license upgrades have been enabled on the Management Server. In an environment with multiple Management Servers, no action is required if the automatic license upgrade feature has been enabled on the Management Server that controls the Shared Domain. See Configuring Automatic Software Updates (page 241). Otherwise, upgrade the licenses manually. License upgrades are available shortly before a new version is released.
- Changes in license binding: Change licenses manually if you change a control IP address used for license binding in a static IP-address-bound license, or if you move a dynamic Management Server POL-bound license to a different Management Server. If you have IP-address-bound licenses for engine components, you must switch the license to a Management Server POL-bound license if you change the engine's control IP address.

Each license indicates the *maximum version* for which the license is valid, but the license is also valid for all previous software versions. You must update the license if you upgrade a component to a new *major release* indicated by a change in the first two digits of the version number (for example, an upgrade from 1.2.3 to 1.3.0 or an upgrade from 1.2.3 to 2.0.0). If only the last number changes, the existing license is valid also for the higher software version.

Appliance licenses do not allow upgrading an appliance beyond the last supported software version on which the appliance can run. See http://www.stonesoft.com/en/customer_care/product_life_cycle/ to check which appliance models have a last supported software version. With third-party hardware, be careful not to upgrade the software to a version that exceeds the hardware's capabilities.

Management Server License Limits

The license limit concerns Management Server licenses that impose a restriction on the number of managed or monitored elements. If your Management Server license allows an unlimited number of managed components, this restriction does not apply to you. To check this information, see the properties of the Management Server's License.

Each engine is generally counted as one managed unit in the limitation, with some exceptions:

- Engines installed on virtualization platforms are not counted.
- Any number of clustered engines counts as a single managed unit, regardless of the number of engines in the cluster (one cluster equals one managed unit).
- L-model appliances (for example, FW-310L) count as half a managed unit (two components equal one managed unit). If all nodes in a cluster are L-model appliances, the whole cluster is counted as half a managed unit (mixed clusters are counted as one managed unit).
- Third-party components that are monitored through the Management Center count as one-fifth of a unit (five components equal one managed unit).

You cannot combine licenses in your system. For example, two Management Server licenses that each contain a restriction for five managed components only allow you to manage five components even if you bind both licenses to a single Management Server.

Authentication Server License Limits

The basic license that enables the use of the optional Authentication Server component limits the number of named users for user linking in the Authentication Server's user database and the number of RADIUS clients (excluding other Stonesoft components) that use the authentication services provided by the Authentication Server. The number of named users and RADIUS clients can be increased by purchasing a license upgrade.

Master Engine License Limits

Master Engines use the same Security Engine licenses as other engines. The Security Engine licenses enable the Virtual Security Engine features. Virtual Security Engines do not require their own licenses. However, the Security Engine license limits the number of Virtual Resources that can be created. The limit for the number of Virtual Resources limits how many Virtual Security Engines can be created. You can optionally increase the allowed number of Virtual Resources by purchasing and importing a feature-specific license.

What's Next?

▶ Generating New Licenses (page 1047)

▶ Installing Licenses (page 1050)

Generating New Licenses

Prerequisites: Appliance delivery or license code information from Stonesoft

Licenses always indicate the newest software version you are entitled to, but they are valid for licensing any older software versions as well. Generally, each Management Center component must have a separate license. Some additional features may also require a separate license.

Note – Your system may be able to automatically generate licenses for new Stonesoft appliances. For automatic licensing to work, install a license for the Management Center components and make sure that automatic updates are enabled on the Management Server. The factory-installed temporary license is automatically replaced with a permanent POS-bound license when the appliance is configured for use.

▼ To generate a new license

1. Go to the Stonesoft License Center at www.stonesoft.com/en/customer_care/licenses/.
2. Enter the POL or POS code in the **License Identification** field.
   - The proof-of-license (POL) code identifies a license. You can find it in the order delivery message sent by Stonesoft (usually by e-mail). For previously licensed components, the POL code is also shown in the Licenses tree in the Administration Configuration view.
   - Stonesoft appliances additionally have a proof-of-serial number (POS) that you can find on a label attached to the appliance hardware.
3. Click **Submit**. The License Center page opens.
4. Check which components are listed as included in this license and click **Register**. The License Generation page opens.
5. Read the directions on the page and fill in the required fields for all included components.
6. Enter the details that bind each license to a component in your system:
   - Management Server(s): enter the IP address you plan to use on the server. If your license allows several Management Servers in the same Management Center (for high availability), enter a comma-separated list of the IP addresses of all the Management Servers.
   - Other SMC servers: enter the Management Server's POL code or the IP address you plan to use on the server.
   - Master Engines: enter the POS code of a Stonesoft appliance (see the label attached to the appliance) or the Management Server's POL code. POS binding is always recommended when the option is available.
   - Security Engines: enter the POS code of a Stonesoft appliance (see the label attached to the appliance) or the Management Server's POL code. POS binding is always recommended when the option is available.
   - SSL VPN Gateways: enter the DNS domain name(s) and/or IP addresses to which the gateways are bound.

Note – If the binding information is incorrect, the license is unusable. If you accidentally generate a license with the wrong binding information, request a license change through the Stonesoft License Center.

7. Click **Submit Request**. The license file is sent to you and is also available for download at the License Center.

What's Next?

▶ Proceed to Installing a License for an Unlicensed Component (page 1050).

Upgrading Licenses Manually

Prerequisites: See Getting Started with Licenses for an overview

Licenses are valid for any older software versions in addition to the version indicated on the license, so you can upgrade the licenses at any time without affecting the system's operation.

> **Note – IP-address-bound licenses have been previously available for Firewalls and IPS engines. You can use and update a previously generated IP-address-bound engine license, but you must switch the license binding to the Management Server's POL code if the engine's control IP address changes.**

You can view, change, and download your current licenses at the Stonesoft License Center at http://www.stonesoft.com/en/customer_care/licenses/ by logging in with your personal account (to view all licenses generated using that account) or by entering a proof-of-license (POL) or proof-of-serial (POS) code (to view information related to a particular license).

If automatic license upgrades have been enabled in the Management Server properties, your licenses are kept up-to-date automatically. Otherwise, you can upgrade licenses manually in the following ways:

- When you log in to the online License Center, you can upgrade the license for the displayed component(s) through the link provided and save the license as a file that you can import in the system as explained in Installing Licenses (page 1050).
- You can export information on licenses through the Management Client and use the resulting text file to upgrade the licenses as explained below.

▼ To upgrade multiple licenses

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Expand the **Licenses** branch and browse to the type of Licenses that you want to upgrade.
3. Ctrl-select or Shift-select the licenses you want to upgrade.

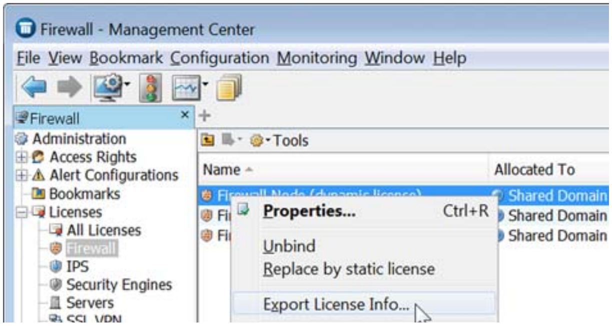

4. Right-click one of the selected items and select **Export License Info**. The Save License Upgrade Request dialog opens.

5. Select the location at which to save the license file in the dialog that opens. You are prompted to request a license upgrade.
6. Click **Yes**. The Stonesoft web site opens.
7. Browse to **Customer Care→Licenses**.
8. Enter the POL or POS code in the **License Identification** field and click **Submit**. The License Center page opens.
9. Click the **Multi-Upgrade Licenses** link on the right. The Upload Multi-Upgrade Licenses page opens.
10. Enter the information needed for the upgrade request and select or upload the license file(s) to update.
11. Click **Submit** to upload the license request. A confirmation page opens, showing the details of your request.
    - The upgraded licenses are e-mailed to you in a .zip file.

What's Next?

▶ To import the upgraded license(s), proceed to Replacing the License of a Previously Licensed Component (page 1051).

Changing License Binding Details

Prerequisites: None

Licenses that are bound to an IP address must be changed if the IP address of the component changes. Licenses that are bound to the POL code of the Management Server must be changed if you want to transfer the licenses to a different Management Server or if you replace the Management Server's license with a license that has a different POL code.

Note – Starting from SMC version 5.0, only licenses for Management Center servers can be bound to an IP address. If you have IP-address-bound licenses for engine components, you must switch to a Management Server POL-bound license if you change the engine's control IP address.

Stonesoft appliances use POS-based licenses. The licenses are bound to the serial number of the appliance hardware, and are automatically bound to the correct element.

You must change IP-address binding and POL-based binding manually. To view, change, and download your current licenses at the Stonesoft License Center at http://www.stonesoft.com/en/customer_care/licenses/, log in with your personal account (for all licenses that your account is authorized to view) or by entering a POL or POS code (to view information related to a particular license).

What's Next?

▶ If automatic license updates have been enabled in the Management Server properties, changed licenses are automatically downloaded and bound to the correct element as long as the license's identification code remains the same.

▶ To manually bind new licenses, proceed to Replacing the License of a Previously Licensed Component (page 1051).

Installing a License for an Unlicensed Component

If your system is divided into administrative Domains, licenses can be viewed only in the Shared Domain. When the license is bound to an element (automatically or manually), the license is shown only in the Shared Domain, regardless of the Domain to which the element belongs. Licenses that are not allocated to any Domain can be manually allocated to a particular Domain through their right-click menu.

> **Caution – If you use Management Server POL code-based licenses for Stonesoft appliances, make sure the licenses are bound to the correct engine elements. Remember to change the license if you change the appliance hardware.**

▼ To install a new license

1. Select **File→System Tools→Install License**. The Install License File(s) dialog opens.
2. Select one or more license files and click **Install**. The licenses are imported and installed. Most types of licenses are also automatically bound to the correct components.
3. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
4. Expand the **Licenses** branch and browse to the type of Licenses you installed.
5. If you installed a dynamic engine license (bound to the Management Server's POL code), right-click the newly-imported license and select **Bind** to select the correct engine.
   - If you bound the license to an incorrect element, right-click the license and select **Unbind**.

> **Caution – The license is permanently bound to the engine and cannot be bound to another engine after you install or refresh the policy on the engine. Permanently bound licenses cannot be rebound unless you relicense or delete the engine element that the license is currently bound to. When unbound, a permanently bound license is shown as *Retained*.**

6. Check that components are now licensed as intended. See Checking License Validity and Status (page 1053).
   - The list of Unlicensed Components may show engines that have a license bound to the POS code of an appliance, because POS-bound licenses are bound to the correct engines automatically when the engine is installed and makes initial contact with the Management Server. The engine is moved to the correct type of Licenses branch after initial contact with the Management Server.
   - Engine licenses are applied when you upload or refresh the policy on the engines.
   - If any engines are not correctly licensed, you may need to upgrade or generate the licenses again.

Related Tasks

- ▶ Checking If All Components Are Licensed (page 1052)
- ▶ Checking License Validity and Status (page 1053)

Replacing the License of a Previously Licensed Component

If your system is divided into administrative Domains, licenses can be viewed only in the Shared Domain. When the license is bound to an element (automatically or manually), the license is shown only in the Shared Domain, regardless of the Domain to which the element belongs. Licenses that are not allocated to any Domain can be manually allocated to a particular Domain through their right-click menu.

Caution – If you use Management Server POL code-based licenses for Stonesoft appliances, make sure the licenses are bound to the correct engine elements. Remember to change the license if you change the appliance hardware.

Invalid or missing licenses may prevent components from working. If you are manually replacing working licenses with new ones, we recommend that you take a backup of the Management Server before you make changes so that you can easily roll back the changes. See Creating Backups (page 1009).

▼ To install an updated license for an already licensed component

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Select **File→System Tools→Install Licenses**. The Install License File(s) dialog opens.
3. Select one or more license files and click **Install**. The licenses are imported and installed.
   - If the licenses only upgrade the software version in a previously imported license (license identifier does not change), the previous license is replaced with the new one automatically and no further action is needed.
   - An IP-address-bound license is automatically bound to the component with the correct IP address.
   - A POS-bound license is automatically bound to the appliance with the correct POS code.

Note – If you bind an IP-address-bound license or a POS-bound license to a component that already has a Management Server POL-bound license, the component has two licenses. Replace the Management Server POL-bound license as instructed in Step 4.

4. Replace the license in one of the following ways:

Table 61.3 Replacing Management Server POL-bound Licenses

Old License Binding	New License Binding	Configuration
Management Server POL-bound license	IP-address-bound license	1. Expand the **Licenses** branch and browse to the correct type of license. 2. Right-click the Management Server POL-bound license that is currently bound to the element and select **Replace by Static License**. 3. Select the license to bind to the engine.
	POS-bound license	

Table 61.3 Replacing Management Server POL-bound Licenses

Old License Binding	New License Binding	Configuration
Management Server POL-bound license that is not bound to the Management Server	Management Server POL-bound license that is bound to the Management Server	1. Expand the **Licenses** branch and browse to the correct type of license. 2. If there is a previous Management Server POL-bound license bound to the element, right-click the old license and select **Unbind**. 3. Right-click the new license and select **Bind**. 4. Select the engine to bind to the license.

Caution – After you install or refresh the policy on the engine, the license is permanently bound to the engine and cannot be bound to any other engine. Permanently bound licenses cannot be rebound unless you relicense or delete the engine element that the license is currently bound to. When unbound, a permanently bound license is shown as *Retained*.

5. Check the displayed license information, and make sure that all the components you meant to license have the correct new license.
6. Right-click any old licenses that may still be displayed and select **Delete**.
7. Refresh or install the policy to transfer the license changes to the engine.

Related Tasks

- Checking License Validity and Status (page 1053)

Checking If All Components Are Licensed

Prerequisites: None

Each Management Center Server, Security Engine, and Master Engine must have its own license. There is no difference between licenses for nodes in a cluster and licenses for single engines.

▼ To check if all components are licensed

1. Select **Configuration→Configuration→Administration**.
2. Browse to **Licenses→Unlicensed Components→All**. This view displays all elements in your system that require a license but do not currently have a valid license.
3. If any components are shown as unlicensed check the following:
   - Engine licenses generated based on a Stonesoft appliance POS code are bound when the engine makes initial contact with the Management Server. It is normal to see the corresponding elements as unlicensed until initial contact is made.
   - If you have already generated and installed licenses, check that the binding details are correct (POS code, POL code, or IP address).
   - To generate and install new licenses for the components listed, see Generating New Licenses (page 1047).

Checking License Validity and Status

Prerequisites: None

▼ To check license validity and status

1. Select **Configuration→Configuration→Administration**.
2. Browse to **Licenses→All Licenses**.

The view displays each license and the component it is currently bound to, along with the newest software version that the license allows you to install. Licenses are valid for all minor releases within the displayed major version (for example, a 5.1 license allows installing 5.1.0, 5.1.1, 5.1.2, etc.) and for any previous software version.

It is not possible to create new IP-address-bound licenses for engine components. Previously created IP-address-bound engine licenses may be present in your system and can be upgraded to new versions. However, if a license is changed in any other way, the binding must also be changed.

The information shown in the Binding column depends on the type of license:

Table 61.4 License Binding Information

License Type	Explanation
Static License	The IP address(es) to which the license is bound.
Dynamic License	The POL code of the Management Server to which the license is bound.
POS-Bound License	The POS code of the appliance to which the license is bound.

The information shown in the Status column is explained in the table below.

Table 61.5 License Status Information

Status	Explanation
Bound	The license is permanently bound to some component. Can be unbound to the “Retained” status.
Assigned	The license is assigned to a component, but is not yet permanently bound to it. This may be because the component's policy has not been installed after the license binding, or because the component is a type of element that does not have a policy. Can be unbound to the “Unassigned” status or deleted.
Retained	The license has been unbound from the “Bound” status. Can be rebound to the same component with the “Cancel Unbind” action. Management Server POL-bound licenses can be changed to the “Unassigned” status by completing one of the procedures in License Is Shown as Retained (page 1106).
Unassigned	The license is not bound to any component. Can be bound to a component or deleted. To bind the license to a component, right-click the license and select **Bind**.

Table 61.5 License Status Information (Continued)

Status	Explanation
Expired	The license has reached the end of its validity period. Only evaluation licenses and licenses for engines installed on virtualization platforms have a validity period. Other licenses are valid for the software version that is indicated in the license without any time restrictions.
Invalid (*Management Server POL-bound licenses only*)	The license was generated using a POL code that does not belong to the Management Server to which you are currently connected.

CHAPTER 62

UPGRADING THE MANAGEMENT CENTER

This section explains how you can upgrade the Management Servers, Management Clients, Log Servers, Authentication Servers, and Web Portal Servers in your Management Center.

The following sections are included:

Getting Started with Upgrading the SMC

Prerequisites: None

What SMC Upgrades Do

In addition to updating the Stonesoft Management Center (SMC) software, the upgrade also makes other changes in your system:

- New system elements and policies may be added and obsolete system elements may be removed. Elements that are in use are not deleted, but instead converted from system elements to regular elements when they have no default role anymore.
- Any elements may be updated with new types of options (related to new/changed features), and occasionally obsolete options may be removed or changed.
- A new dynamic update package is activated in your system, unless you have already installed the same or a newer update package before the installation (this may be cause of some, but not necessarily all of the changes listed above).
- The Management Client's Online Help system may be updated with new or corrected information.

A summary of changes to elements is created during each upgrade; you are presented a link to see these HTML reports when the Management Server upgrade is finished.

Limitations

All Management Center components (Management Server, Management Client, Log Server, and the optional Web Portal Server) must have the same software version. All other components try to connect the Management Server when they start. They do not start up if their software version does not match with the Management Server software version. If you have multiple Management Servers or Log Servers, you must upgrade each server separately so that they have the same software version.

What Do I Need to Know before I Begin

Although the need to do so is unlikely, the upgrade can be easily reversed if you take the correct precautions. If the Management Center upgrade fails, you can automatically revert to the previous installation if you select the option in the installer before the upgrade starts. In any case, we recommend that you take a backup of the Management Server using the SMC's internal backup tool before you upgrade. The backup contains all necessary information to restore the configurations (including the engine configurations). The backup does not contain software version or operating system specific information and can always be restored from an older-version system to a newer-version system if a direct upgrade is supported between the software versions involved. For more information, see Backing Up and Restoring System Configurations (page 1007).

The SMC is offline during the upgrade. The engines continue to operate normally and store their generated log data in their local spool while the Management Server(s) and Log Server(s) are offline. Once connectivity is restored, the spooled log data is transferred from the engines to the Log Server(s).

To check which version of the Management Center you are currently using, select **Help→About** in the Management Client. Additionally, the Management Client's version is displayed in the Management Client's login dialog.

Configuration Overview

1. Prepare the installation files as explained in Obtaining the SMC Installation Files.
2. *(If automatic license updates have been disabled)* Update the licenses as explained in Upgrading Licenses Manually (page 1048).
3. Upgrade all components that work as parts of the same SMC as explained in Upgrading Management Center Servers (page 1058).
4. (*Multiple Management Servers only*) Synchronize the Management database between the Management Servers. See Synchronizing Management Databases Manually (page 339).
5. If you distribute Web Start Management Clients from some external server (instead of the Management Server), update the Web Start distribution as explained in Distributing Web Start from External Servers (page 310).

Obtaining the SMC Installation Files

Prerequisites: None

Before running the installer, check the installation file integrity using the MD5 or SHA-1 file checksums. The checksums are on the installation DVD and on the product-specific download pages at the Stonesoft website.

Windows does not have MD5 or SHA-1 checksum programs by default, but there are several third-party programs available.

Stonesoft provides only recent versions of the software for download. We recommend you store the upgrade files yourself to make sure you can install the exact same version at a later time (for example, in case of a hardware failure), especially if your organization’s policies mandate lengthy testing periods that limit the speed of adopting new versions.

▼ To obtain the Management Center installer

1. Download the installation file from http://www.stonesoft.com/en/customer_care/downloads/. There are several packages available:
   - The .iso download allows you to create an installation DVD that can install or upgrade the software on all supported platforms.
   - Separate .zip packages are available for downloading installation files for all supported platforms or just one supported platform.

2. Change to the directory that contains the file(s) to be checked.

3. Generate a checksum of the file using the command **`md5sum`** *`filename`* or **`sha1sum`** *`filename`*, where *`filename`* is the name of the installation file.

Example
```
md5sum sg_engine_1.0.0.1000.iso
869aecd7dc39321aa2e0cfaf7fafdb8f sg_engine_1.0.0.1000.iso
```

4. Compare the displayed output to the checksum on the website.

Caution – Do not use files that have invalid checksums. If downloading the files again does not help, contact Stonesoft technical support to resolve the issue.

If you are installing from the `.iso` file, create the installation DVD using a DVD-burning application that can correctly convert the DVD structure from the `.iso` image to a DVD with several files and folders. If the end result is a DVD with the original `.iso` file on it, the DVD cannot be used for installation. If you are installing from a .zip file, unzip all folders and files in the archive to the server you want to upgrade.

Upgrading Management Center Servers

Prerequisites: Upgrading Licenses Manually (if automatic license upgrades are not active)

You can upgrade without uninstalling the previous version. If you want to reinstall in a new operating system or on different hardware, see Changing the Management Platform (page 344).

Caution – All the Management Center components (Management Server, Management Client, Log Server, and the optional Web Portal Server and Web Portal) must use the exact same SMC software version to work together. If you have multiple Management Servers of Log Servers, you must upgrade each server separately.

The same installer works with all Management Center components, including locally installed Management Clients.

If you have multiple Management Servers or Log Servers, you can upgrade them in any order. Management Servers are automatically isolated from database replication during the upgrade. There is no need to explicitly isolate the Management Servers before upgrading.

If you are upgrading from a very old version of the SMC to a new version, you may have to upgrade to an intermediate version first before upgrading to the latest version (check the *Release Notes* at http://www.stonesoft.com/en/customer_care/kb/.

▼ To upgrade the Management Center components

1. Start the installation in one of the following ways:
   - **From a .zip file:** unzip the file and run **`setup.exe`** on Windows or **`setup.sh`** on Linux.
   - **From a DVD**: insert the installation DVD and run the `setup` executable from the DVD:

Operating System	Path to Executable
Windows 32-bit	\Stonesoft_SW_Installer\Windows\setup.exe
Windows 64-bit	\Stonesoft_SW_Installer\Windows-x64\setup.exe
Linux 32-bit	/Stonesoft_SW_Installer/Linux/setup.sh
Linux 64-bit	/Stonesoft_SW_Installer/Linux-x64/setup.sh

Note – If the DVD is not automatically mounted in Linux, mount the DVD with “`mount /dev/cdrom /mnt/cdrom`”.

2. Read and accept the License Agreement to continue with the installation. The Installation Wizard automatically detects the previous installation directory.
3. Click **Next** to accept the installation directory. The Installation Wizard displays the components to be upgraded.

Illustration 62.1 Select SMC Components

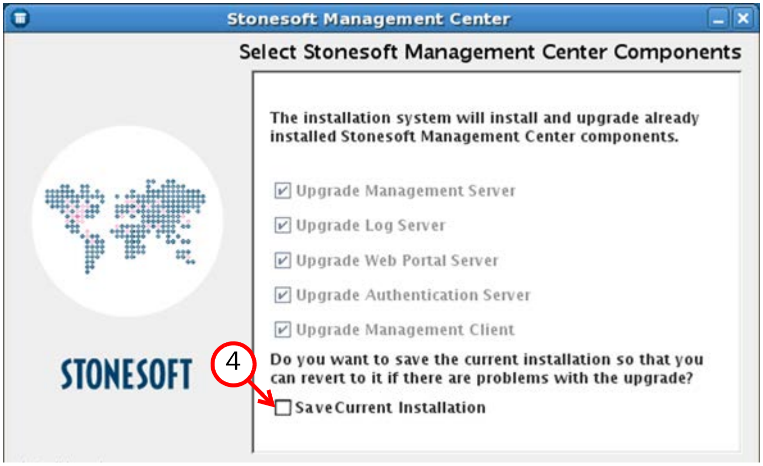

Stonesoft Management Center

Select Stonesoft Management Center Components

The installation system will install and upgrade already installed Stonesoft Management Center components.

☑ Upgrade Management Server

☑ Upgrade Log Server

☑ Upgrade Web Portal Server

☑ Upgrade Authentication Server

☑ Upgrade Management Client

Do you want to save the current installation so that you can revert to it if there are problems with the upgrade?

□ Save Current Installation

4. (*Management Server only, optional*) Select **Save Current Installation** if you want to save a copy of the current installation that you can revert to at any time after the upgrade.

5. Click **Next**.

6. (*Management Server and Authentication Server only*) Select whether to back up the server and click **Next**:
   - Select **Yes** to create a backup that can be used and viewed without a password.
   - Select **Yes, encrypt the backup** to create a password-protected backup. You are prompted for the password as you confirm the selection.
   - Select **No** if you already have a recent backup of the Management Server or Authentication Server.

7. Check the pre-installation summary and click **Install**. The upgrade begins.

8. When the upgrade is complete, click the link(s) in the notification to view the report(s) of changes the installer has made to your system configurations in your web browser.

9. Click **Done** to close the installer.

10. Upgrade any SMC components that run on other computers (for example, additional Management Servers or Log Servers) in the same way.

11. (*Multiple Management Servers only*) Synchronize the Management database between the Management Servers. See Synchronizing Management Databases Manually (page 339).

What's Next?

▶ If you distribute Web Start Management Clients from some external server (instead of the Management Server), update the Web Start distribution as explained in Getting Started with Web Start Distribution (page 308).

▶ Otherwise, the Management Center upgrade is now complete.

Default Installation Directories for SMC

Prerequisites: None

Default installation directory on Windows:

- `C:\Stonesoft\Management Center`

> **Note – If you installed the Management Server in the `C:\Program Files\Stonesoft\Management Center` directory in Windows, some program data may be stored in the `C:\ProgramData\Stonesoft\Management Center` directory.**

Default installation directory on Linux:

- `/usr/local/stonesoft/management_center`

Under the installation directory, note especially the following folders:

- `/backups/` stores Management Server (`sgm_`) and Log Server (`sgl_`) backups. The backups must be in this directory to be listed in the Management Client and when running scripts without specifying a backup file.
- `/bin/` contains the SMC command line tools (see Management Center Commands (page 1146)) as well as some additional scripts that are used by the Installation Wizard.

CHAPTER 63

UPGRADING THE ENGINES

This section explains how you can upgrade Firewalls, IPS engines, Layer 2 Firewalls, and Master Engines.

The following sections are included:

Getting Started with Upgrading Engines

Prerequisites: None

This chapter describes how to remotely upgrade engines using the Management Client. The local upgrade option is not covered here (see the product-specific *Installation Guide* instead).

How Engine Upgrades Work

The upgrade package is imported to the Management Server manually or automatically. You then apply it to selected engines through the Management Client.

The engines have two alternative partitions for the software. When you install a new software version, it is installed on the inactive partition and the current version is preserved to allow rollback to the previous version in case the installation is interrupted or other problems arise. If the engine is not able to return to operation after the upgrade, it automatically switches back to the previous software version at the next reboot. You can also switch the active partition manually.

You can upload and activate the new software separately. For example, you can upload the upgrade during office hours but activate it during a service window.

The currently installed working configuration (routing, policies, etc.) is stored separately and is not changed in an upgrade or a rollback. Although parts of the configuration may be version-specific (for example, if system communications ports are changed), the new software version can use the existing configuration. Possible version-specific adjustments are made when you refresh the policy after the upgrade.

Limitations

It is not possible to upgrade between a 32-bit version and a 64-bit version of the software. If you are running the software on a compatible standard server, you can reinstall the software using the other version. In clusters, 32-bit and 64-bit nodes cannot be online simultaneously. Appliances support only the software architecture version that they are pre-installed with.

Due to changes in the IPS components, additional steps are required for upgrading legacy Sensors, Sensor Clusters, and combined Sensor-Analyzers to version 5.4 or higher. See Upgrading Legacy IPS Engines (page 1066).

You cannot upgrade Virtual Security Engines directly. To upgrade Virtual Security Engines, you must upgrade the Master Engine that hosts the Virtual Security Engines.

What Do I Need to Know before I Begin

The Management Center must be up-to-date before you upgrade the engines. An old Management Center version may not be able to recognize the new version engines and may generate an invalid configuration for them. The Management Server can control several older versions of engines. See the *Release Notes* for version-specific compatibility information.

During a cluster upgrade, it is possible to have the upgraded nodes online and operational side by side with the older version nodes. This way, you can upgrade the nodes one by one while the other nodes handle the traffic. However, you must upgrade all the nodes to the same version as soon as possible, as prolonged use with mismatched versions is not supported.

The current engine version is displayed on the **General** tab in the Info panel when you select the engine. If the Info panel is not shown, select **View→Info**.

Configuration Overview

1. (*Manual download of engine upgrade files*) Prepare the installation files as explained in Obtaining Engine Upgrade Files (page 1063).
2. *(Manual license updates)* Update the licenses as explained in Upgrading Licenses Manually (page 1048).
3. Upgrade the engines as explained in Upgrading Engines Remotely (page 1064).

Obtaining Engine Upgrade Files

Prerequisites: None

If the Management Server is not set up to download engine upgrades automatically (see Configuring Automatic Software Updates (page 241)) or if you want to perform a local upgrade, you must download the installation files manually and check the installation file integrity using the MD5 or SHA-1 file checksums. Windows does not have MD5 or SHA-1 checksum programs by default, but there are several third-party programs available.

The remote upgrade is recommended in most cases. See the product-specific *Installation Guide* for detailed instructions if you want to upgrade engines locally.

▼ To manually download an engine upgrade file

1. Download the .zip installation file for your architecture (32-bit or 64-bit) from www.stonesoft.com/en/customer_care/downloads/.
2. Change to the directory that contains the file(s) to be checked.
3. Generate a checksum of the file using the command **`md5sum`** *`filename`* or **`sha1sum`** *`filename`*, where *`filename`* is the name of the installation file.

Example

```
md5sum sg_engine_1.0.0.1000.zip
869aecd7dc39321aa2e0cfaf7fafdb8f sg_engine_1.0.0.1000.zip
```

4. Compare the displayed output to the checksum on the Stonesoft website.

> **Caution – Do not use files that have invalid checksums. If downloading the files again does not help, contact Stonesoft Support to resolve the issue.**

5. Log in to the Management Client and select **File→Import→Import Engine Upgrades**. The Import Engine Upgrade dialog opens.
6. Select the engine upgrade (sg_engine_*`version_platform`*`.zip` file) and click **Import**. The import takes a while. You can see the related messages in the status bar at the bottom of the Management Client window.

What's Next?

▶ Proceed to Upgrading Engines Remotely (page 1064).

Upgrading Engines Remotely

Prerequisites: (*Possibly*) Upgrading the Management Center, Obtaining Engine Upgrade Files

The Management Server can remotely upgrade engine components that it manages. You can upgrade several engines of the same type in the same operation. However, we recommend that you upgrade clusters one node at a time and wait until an upgraded node is back online before you upgrade the other node(s).

Note – Clusters operate normally throughout the upgrade when the upgrade is done in stages. However, it is recommended to upgrade all nodes in the cluster to the same version as soon as possible. Prolonged use with mismatched versions is not supported. It is not possible to have 32-bit and 64-bit engines online in the cluster at the same time.

▼ To upgrade an engine remotely

1. Read the *Release Notes* for the new version, especially the required Management Center version and any other version-specific upgrade issues that may be listed.
   - Select **Configuration→Configuration→Administration** and browse to **Other Elements→Engine Upgrades**. Select the type of engine you are upgrading. A link to the *Release Notes* is included in the upgrade file’s information.
   - If the Management Server has no Internet connectivity, the *Release Notes* can be found through the Technical Documentation search at www.stonesoft.com/en/customer_care/support/.
2. Select **Monitoring→System Status**. The System Status view opens.
3. Right-click the node you wish to upgrade and select **Commands→Go Offline**.
4. Right-click the node you wish to upgrade and select **Upgrade Software** or **Configuration→Upgrade Software** depending on your selection. The Remote Upgrade Task Properties dialog opens.

Illustration 63.1 Remote Upgrade Task Properties

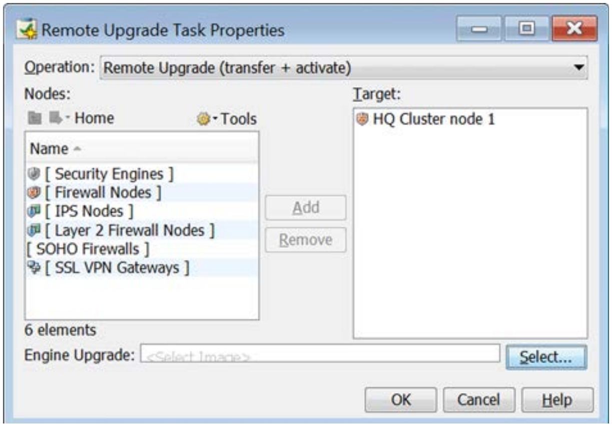

Remote Upgrade Task Properties

Operation: Remote Upgrade (transfer + activate)

Nodes:

Name
[Security Engines]
[Firewall Nodes]
[IPS Nodes]
[Layer 2 Firewall Nodes]
[SOHO Firewalls]
[SSL VPN Gateways]

6 elements

Target: HQ Cluster node 1

Engine Upgrade: <Select Image>

5. Select the type of **Operation** you want to perform:
   - Select **Remote Upgrade (transfer + activate)** to install the new software and reboot the node with the new version of the software.
   - Select **Remote Upgrade (transfer)** to install the new software on the node without an immediate reboot and activation. The node continues to operate with the currently installed version until you choose to activate the new version.
   - Select **Remote Upgrade (activate)** to reboot the node and activate the new version of the software that has been installed earlier.

Caution – To avoid an outage, do not activate the new configuration simultaneously on all the nodes of a cluster. Activate the new configuration one node at a time, and proceed to the next node only after the previous node is back online.

6. If necessary, add or remove Target engines. All engines in the same Upgrade Task must be of the same type.
7. Select the correct engine version and click **OK**. If you chose to activate the new configuration, you are prompted to acknowledge a warning that the node will be rebooted. A new tab opens showing the progress of the upgrade.
   - The time the upgrade takes varies depending on the performance of your machine and the network environment.
   - If you chose to activate the new configuration, the engine is automatically rebooted and brought back online.

The upgrade overwrites the inactive partition and then switches the active partition. To undo the upgrade, use `sg-toggle-active` command or the engine's boot menu to switch back to the previous software version on the other partition. This switch may also happen automatically at the next reboot if the engine is not able to successfully return to operation when it boots up after the upgrade.

What's Next?

▶ When the upgrades are finished, refresh the policy of upgraded engines to make sure any possible changes specific to the new software version are transferred to the engines.

Related Tasks

▶ Accessing the Engine Command Line (page 232)

▶ Creating Remote Upgrade Tasks (page 1038)

▶ Engine Commands (page 1157)

Upgrading Legacy IPS Engines

Prerequisites: (*Possibly*) Upgrading the Management Center, Obtaining Engine Upgrade Files

Prior to version 5.4, IPS engines consisted either of separate Sensor and Analyzer engines, or combined Sensor-Analyzer engines. In version 5.4, the Analyzer functionalities have been transferred to the Log Server and to Security Engines, and the Analyzer is no longer used. Because of this change, additional steps are required for upgrading legacy Sensors, Sensor Clusters, and combined Sensor-Analyzers to version 5.4 or higher. To begin the upgrade, proceed to the relevant section below:

What's Next?

- Upgrading Sensors and Sensor Clusters to IPS Engines (page 1066)
- Upgrading a Legacy Sensor-Analyzer to a Single IPS Engine (page 1067)

Upgrading Sensors and Sensor Clusters to IPS Engines

To upgrade Sensors and Sensor Clusters to IPS engines

1. Upgrade the engine software to Security Engine version 5.4 or higher as instructed in Upgrading Engines Remotely (page 1064).

Note – If you are upgrading a legacy Sensor Cluster, upgrade all nodes of the cluster before proceeding to Step 2.

2. Open the properties of the upgraded engine or engine cluster.

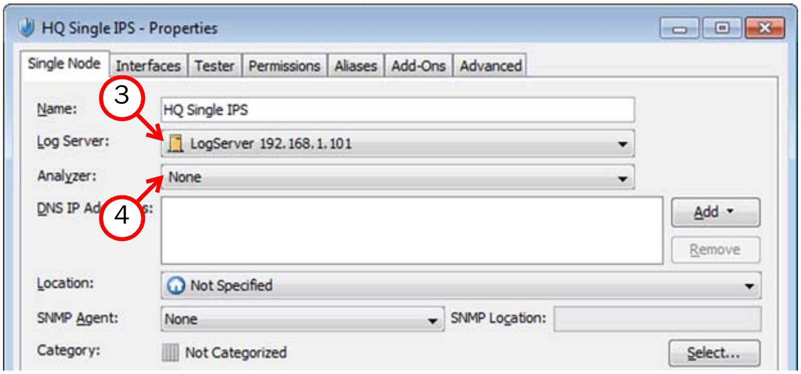

HQ Single IPS - Properties

Name: HQ Single IPS

Log Server: LogServer 192.168.1.101

Analyzer: None

Location: Not Specified

SNMP Agent: None

SNMP Location:

Category: Not Categorized

3. Make sure a **Log Server** is selected.
4. Select **None** for the **Analyzer**.
5. Click **OK**.
6. Refresh the policy to transfer the configuration changes to the engine.

What's Next?

- Upgrade all legacy Sensors or Sensor Clusters in the same way, then proceed to Removing Unused Analyzers (page 1068).

Upgrading a Legacy Sensor-Analyzer to a Single IPS Engine

When you upgrade a legacy Sensor-Analyzer engine, you convert the combined Sensor-Analyzer to a Single IPS engine. The Analyzer node is automatically removed during the conversion.

▼ To upgrade a legacy Sensor-Analyzer to a Single IPS engine

1. Select **Monitoring→System Status**. The System Status view opens.
2. Expand the Sensor-Analyzer element until you can see the Sensor and Analyzer nodes.
3. Right-click the Analyzer node and select **Tools→References**. If there are any references to the Analyzer, remove the references before the upgrade.
   - See Searching for Element References (page 46) for more information.
4. Upgrade the engine software to Security Engine version 5.4 or higher as instructed in Upgrading Engines Remotely (page 1064).

Note – The engine does not go online and process traffic until you refresh the policy after the conversion. The engine status is shown as red/gray with the message *No Policy Installed* after the upgrade, and as yellow/gray after the conversion.

5. Right-click the Sensor-Analyzer element and select **Configuration→Upgrade to Single IPS**. The Sensor-Analyzer properties dialog opens.
6. Select the **Log Server** to which event data is sent.
7. Select the **Failure Mode** to define how traffic through the Inline Interface is handled if the IPS engine goes offline.
   - **Bypass**: traffic is allowed through the Inline Interface without inspection.
   - **Normal**: traffic is not allowed through the Inline Interface.

Caution – Using Bypass mode requires a fail-open network interface card. If the ports that represent the pair of Inline Interfaces cannot fail open, policy installation fails on the engine.

8. Click **OK**. The conversion begins.
   - If the conversion fails, switch the active partition to restore the previous engine version and downgrade the Single IPS engine to a Sensor-Analyzer through the element's right-click menu.
9. Refresh the policy to transfer the configuration changes to the engine.

Removing Unused Analyzers

When you upgrade legacy Sensors or Sensor Clusters to version 5.4 IPS engines, existing Analyzer elements are kept in the system but are no longer used. After all legacy Sensors or Sensor Clusters have been upgraded, you can safely remove any unused Analyzer elements.

▼ To remove an unused Analyzer element

1. Select **Monitoring→System Status**. The System Status view opens.
2. Right-click the Analyzer and select **Tools→References**. If there are any references to the Analyzer, remove the references before deleting the element.
   - See Searching for Element References (page 46) for more information.
3. Right-click the Analyzer element and select **Mark as Obsolete**. You are prompted to confirm that you want to mark the element as obsolete.
4. Click **Yes**. The element is marked as obsolete and moved to the Obsolete Elements branch.
5. (*Optional)* Select **View→Obsolete Elements** and permanently delete the element from the Obsolete Elements branch as instructed in Deleting Obsolete Elements (page 89).

Chapter 64

Manual Dynamic Updates

Dynamic Update packages include changes and additions to the system Policies, Situations, and other elements of the Management Center.

The following sections are included:

Getting Started with Manual Dynamic Updates

Prerequisites: None

It is very important to keep the system policies and situations up-to-date so that newly discovered vulnerabilities can be detected. Changes and additions are provided in dynamic update packages available at the Stonesoft website.

What Dynamic Updates Do

Dynamic updates provide updates particularly to the deep inspection features on the Security Engines. For example, new threat patterns and changes in the system Templates and Policies are introduced in dynamic updates for up-to-date detection. They may also revise the default elements you use to configure the system.

Limitations

- You may need to upgrade first before you can use a certain dynamic update package. See the *Release Notes* for the update packages at the Stonesoft web site for more information.
- If there are several Domains defined in the system, manual dynamic updates can only be installed in the Shared Domain.

What Do I Need to Know Before I Begin?

As an alternative to downloading the updates manually as explained here, you can configure the dynamic updates to be downloaded and optionally activated automatically in your system. See Configuring Automatic Software Updates (page 241).

Virus database updates for Firewalls are always done automatically and directly by the engines. Updates are always active when the anti-virus feature is active. The anti-virus feature is not available for IPS engines or for Layer 2 Firewalls. For additional settings, see Configuring Anti-Virus Settings (page 540).

Configuration Overview

1. Download the latest dynamic update package from the Stonesoft web site and import it in the Management Client. See Importing a Dynamic Update Package (page 1071).
2. Activate the dynamic update package in the Management Client. See Activating a Dynamic Update Package (page 1072).

Importing a Dynamic Update Package

Prerequisites: None

▼ To import a dynamic update package

1. Go to the Stonesoft License Center at https://my.stonesoft.com/download.do.
2. Enter your proof-of-license (POL) code, proof-of-serial (POS) number, or license key in the **License Identification** field and click **Submit**.
3. Click **Stonesoft Dynamic Updates Downloads** in the list of Stonesoft Downloads.
4. Download the latest `.jar` dynamic update package file.
   - For details about the dynamic update package, click *Release Notes* under the `.jar` file.
5. Save the update package file to a location accessible from the computer you use to run the Management Client.

Note – Make sure that the MD5 checksums for the original files verified by Stonesoft and the files that you have downloaded match.

6. Select **File→Import→Import Update Packages**. A file browser dialog opens.
7. Browse to the correct file, select it, and click **Import**. The import takes some time, and the completion of the import is displayed in the status bar of the Management Client window.

What's Next?

▶ Proceed to Activating a Dynamic Update Package (page 1072).

Activating a Dynamic Update Package

Prerequisites: Importing a Dynamic Update Package

Activating a dynamic update package introduces the changes from an imported dynamic update package into your system.

▼ To activate a dynamic update package

1. Select **Configuration→Configuration→Administration**.
2. Expand the **Other Elements** tree and select **Updates**. A list of all imported and activated dynamic update packages opens.
3. Right-click the new dynamic update package and select **Activate**. The Management Server Backup dialog opens.
4. Unless you have a fresh installation, click **Yes**. After the backup is taken, the dynamic update package activation starts. The progress of the import and the items included are shown on a new tab.
   - The included items are shown so that you can view them after the activation is complete.
5. Click **Close** when the activation is finished.

What's Next?

▶ Refresh the policies on all engines to activate the changes.

TROUBLESHOOTING

In this section:

Chapter 65

General Troubleshooting Tips

This section contains tips for how to troubleshoot situations that are not covered by more specific troubleshooting sections.

The following sections are included:

If Your Problem Is Not Listed

When having problems with your system, you should first make sure you have followed the relevant instructions.

Some problems you are having may be related to known issues, which you can view at the Stonesoft web site at http://www.stonesoft.com/en/customer_care/support/.

If your organization is entitled to technical support, contact Stonesoft Support (for contact details, see Technical Support (page 29)). Note that issues related to generating licenses or updating the information contained in the licenses are handled by Stonesoft Order Services (order@stonesoft.com).

Tools For Further Troubleshooting

Logs and alerts provide useful information on what the components do and what is happening in your system. You can increase the detail level of logs on certain areas of operation. See Enabling/Disabling Firewall/VPN Diagnostics (page 222). To gain useful information from the logs produced, you must be able to filter the logs efficiently. See Filtering Logs in the Logs View (page 156).

There are Stonesoft-specific as well as standard networking tools available on the engines. See Working on the Engine Command Line (page 231) and Engine Commands (page 1157).

If you need a deeper understanding of how the Stonesoft components function, consult the relevant section(s) in the *Reference Guides*.

Chapter 66

Troubleshooting Accounts and Passwords

This section concentrates on troubleshooting password-related issues.

The following sections are included:

Forgotten Passwords

Problem description: You or someone else in your organization forgets one of the passwords related to the Stonesoft Management Center.

Solution: You can regain access by changing the password. An administrator who has unrestricted permissions can change any password in the Management Center. The password recovery procedures for the different passwords are as follows:

- **Management Client** login passwords can be changed in the Administrator elements. Administrator elements can be found in the Administration Configuration view under **Access Rights→Administrators**. See Creating a New Administrator Element (page 253) for more information.
- **Web Portal** login passwords can be changed in the Web Portal User elements. Web Portal elements can be found in the Administration Configuration view under **Access Rights→Web Portal Users**. See Defining Web Portal User Accounts (page 298) for more information.
- The **Engine** Root account password (for command line access) can be changed by right-clicking the individual engine node and selecting **Commands→Change Password**. If the engine is not connected to the Management Server (because, for example, it is a spare appliance), you can reset all of the engine's settings through a boot menu option in the local console accessible through a serial connection or through a directly connected monitor and keyboard.

Caution – Resetting the engine through the boot menu stops the engine from processing traffic, since all configurations are cleared.

- A **User** password used for end-user authentication can be changed in the User element. User elements are stored in the User Authentication Configuration view under **Users** (if the user is stored in the internal LDAP database or an external LDAP database that the Management Center is configured to use).
- The default **Management Server Database** user account is **dba**, and the password is created automatically. If you do not know the password but you need it for some operation, you can change the password through **File→System Tools→Password→Change Database Password**.

If none of the administrators can log in due to account issues, see Creating an Emergency Administrator Account (page 1079).

User Account Changes Have no Effect

Problem description: You add a User element (end-user account for authentication) or change the password in a User element, but the new user account or new password are not accepted when the end-user tries to authenticate. Previously created and unmodified user accounts work as expected. If you changed the password, the previous password is still accepted.

Reason: There may be a replication problem that prevents synchronizing the user database information from the Management Server to the local database on the Firewalls.

Solution: Reset the user database by right-clicking an individual Firewall node (not the upper-level Single Firewall/Firewall Cluster element) and selecting **Commands→Reset User Database**. This action copies all user information from the Management Server to the engine.

Also, make sure **User DB Replication** (automatic user database replication) is active under **Options** in the right-click menu for the Single Firewall/Firewall Cluster (top-level) element.

Creating an Emergency Administrator Account

Problem description: None of the administrators can log in to the Management Server because of account-related issues.

Solution: A new account with unrestricted permissions can be created in an emergency using the `sgCreateAdmin` script on the Management Server (located in the `<`*`installation directory`*`>/bin` folder) with the Management Server service stopped.

Note – If you installed the Management Server in the `C:\Program Files\Stonesoft\Management Center` **directory in Windows, some of the program data is stored in the** `C:\ProgramData\Stonesoft\Management Center` **directory. Command line scripts may be found in the** `C:\Program Files\Stonesoft\Management Center\bin` **and/or the** `C:\ProgramData\Stonesoft\Management Center\bin` **directory.**

You can then log in to the Management Client using this account, change passwords, and create new accounts normally as explained in Defining Administrator Accounts (page 253).

Chapter 67

Troubleshooting Alert, Log, and Error Messages

This chapter explains some common alert and log messages that you may see in the Logs view and gives hints on how to proceed when you see such messages.

The following sections are included:

Certificate Authority Expired/Expiring Alerts

See Dealing with Expiring Certificate Authorities (page 1097).

Certificate Expired/Expiring Alerts

See Replacing Expired/Missing Certificates (page 1094).

Log Spool Filling

Problem description: the “Log spool filling” alert is triggered.

Reason: The logs are not transferred from the engine at all or the engine is generating logs more quickly than can be transferred to the log server.

Solution: See Logs Are Filling up the Storage Space (page 1110).

Status Surveillance: Inoperative Security Engines

Problem description: the Inoperative Security Engines alert is triggered.

Reason: the Management Server did not receive expected status updates from a Security Engine and the Status Surveillance option is selected for the engine (in the engine’s right-click menu under Options).

If you see these alerts, one of the problems listed below may exist or may have existed temporarily:

- The connection between the engine and the Management Server may have been lost due to network connectivity problems or due to a technical issue on the Management Server. Problems that affect only management communications do not interfere with the operation of the engines - the engines continue processing traffic.
- The engine may be experiencing technical problems.

▼ To troubleshoot inoperative security engines

1. Check if the status of the engine or the system connections (shown in the info view when the engine is selected) still shows problems.
2. Check if there is a steady log stream from the affected engine and if there are any further alerts or logs from the engine that could explain the reason for the message.
3. Check if the engine is actually processing traffic normally even if the Management Server is not able to monitor the engine and/or show the log stream.
4. A console connection to the affected engine is recommended, if possible, when you suspect that the engine may not be operating properly. This allows you to see any possible error messages printed out to the console before you take corrective actions, such as reboot the node.

If you suspect technical problems on the engine, run the sginfo script on the engine before rebooting it (if possible) and contact Stonesoft support.

Related Tasks

- For information on command line tools, see Command Line Tools (page 1145).

System Alert

Problem description: you receive an alert with the title “System Alert”.

Reason: “System Alert” is a general category for all alert messages that are generated because something in the operation of the Management Center components or Security Engines requires your attention.

Solution: Select the alert entry in the Logs view and click **Details** in the toolbar to view the alert entry information. Some of the most frequently seen messages are also explained in this section.

Test Failed

Problem description: test-related alerts are triggered (for example, “Tester: Link Status test failed”).

Reason: tester alerts indicate that the automatic tester system running on the engines has detected a failure in one of the tests that the tester is configured to run. The Tester is configured to run some tests by default and administrators may configure additional tests. The tester is configured in the properties of each engine.

Solution: make sure the condition that caused the test to fail is resolved.

Related Tasks

- Getting Started with the Engine Tester (page 520).

Throughput License Exceeded

Problem description: throughput based license exceeded alerts are triggered.

Reason: Some licenses limit the throughput of the engine to a certain fixed value. If the throughput limit is reached at any particular moment, the exceeding traffic is dropped and an alert is created to notify you how many packets have been dropped. The throughput limit is counted as the total throughput of all traffic handled by the engine at any one moment (all traffic together regardless of type, direction, or the links used).

Solution: Usually, the messages are triggered by temporary spikes in traffic and do not cause major problems. If you see these messages often, you should take action:

- Make sure your appliance licenses are matched to the correct elements according to the type of appliance. If the license generated with the POS code of a lower-throughput appliance is applied to a higher-throughput appliance, the throughput is needlessly limited.
- If the hardware can handle a higher throughput than what it is licensed for, you may switch to a higher-throughput license (contact your reseller).
- If the license throughput corresponds to the maximum throughput achievable with your hardware, you may be able to install an additional cluster node or you can switch to hardware with a higher maximum throughput (contact your reseller).
- You can restrict the traffic in a more controlled way, for example, using the traffic management features of a Firewall. See Getting Started with QoS (page 782) for more information.

Connection Closed/Reset by Client/Server

Problem description: Connection fails and a connection closed by client/by server or connection reset by client/by server message is shown in the logs. If the connection closing does not occur in the expected order of a normal TCP connection, the message may state “connection closed abnormally”.

Reason: These messages inform you about an event that the Firewall has detected in the network. Connection closing is an expected event at the end of each standard TCP connection. The logging settings in the Access rules determine whether connection closing is logged. Frequent abnormal connection closing or resets may indicate problems in the network, such as an overloaded server.

Solution: If connection problems affect services, perform standard network troubleshooting steps along the whole communication path. See Taking a Traffic Capture (page 126) for information on how to generate a tcpdump file from the Stonesoft Management Center.

Connection Removed During Connection Setup

Problem description: You see “Connection removed during connection setup” messages in conjunction with connection closing logs.

Reason: This message notifies you of a connection that was abnormally cut during the TCP connection setup phase because of a RST (reset) sent by one of the communicating parties.

Solution: Check why the connection is torn down, for example, if the server is overloaded or if the service is down or using a non-standard port.

Connection State Might Be Too Large

Problem description: You see “error when serializing for state sync” messages with a “connection state might be too large” clarification log entries for a Firewall Cluster. You may also experience intermittent or continuous problems with clustering and traffic flow, which are typically alleviated for some time by rebooting all clustered nodes.

Reason: The Firewall keeps a record of all connections that are handled statefully to be able to track the connection. When the Firewall is clustered, this connection table must be synchronized between the nodes to allow connections to continue if a node goes down. When the state table grows excessively large, the Firewall engines can no longer effectively use it.

This message is usually caused by a misconfiguration. Typical configuration problems include:

- Using the Oracle Protocol Agent (PA) on the actual database connections between the client and the server. The Oracle Protocol Agent is meant for cases where TCP port 1521 is used only for negotiating the port number for Oracle database connections, and the port number for the actual connection is assigned dynamically. It must not be used in any other cases.
- Excessive idle timeouts defined in Access Rules. All TCP connections are normally explicitly closed by the communicating parties and can therefore be cleared from the state table based on actual connection state. Non-TCP protocols do not establish connections, but the communications are still handled as virtual connections on the Firewall to allow all Firewall features to be used on the traffic. Since the communicating parties do not have a closing

mechanism, these virtual connections are never cleared from the Firewalls' connection records before the communications are left idle (unused) for the duration of the defined timeout. If access rules define excessively long timeouts for such traffic between many different hosts, the connection state table can grow extremely large.

Solution: If the Oracle Protocol Agent is in use, make sure it is not applied incorrectly. If necessary, replace the default service that has a Protocol Agent attached with a custom service that matches the correct port without a Protocol Agent.

Check the Access Rules to see if there are rules that override the default idle timeout value for non-TCP traffic. Make sure the override is not applied to any traffic that does not absolutely need a longer timeout (make the rule as specific as possible) or try reducing the timeout (generally, the idle timeout should not be more than a few minutes). In some cases, allowing both communications directions separately may remove the need for long timeouts.

Connection Timeout

Problem description: You see “connection closed” log entries with a “connection timeout” clarification.

Reason: Connection timeout log messages are generated for inactive connections that the Firewall clears out from its connection tracking table. For some reason, the hosts that were communicating within this allowed connection have stopped transmitting packets between each other.

Solution: Most connection timeouts are normal and necessary to ensure that the Firewall cleans up inactive connections from its records, freeing up the resources. However, sometimes the timeout may prevent communications from continuing:

- If there is some application in your network that leaves connections inactive for long periods of time before continuing again, you can increase the timeout for those connections in the Action options for the Access rule that allows the connection. The rule-specific timeouts override the global timeouts that are set per connection state in the Firewall element's properties (Advanced Settings).

Caution – Setting long timeouts for a high number of connections considerably increases the resource consumption of the Firewall and may even lead to performance problems. This applies especially to non-TCP protocols that do not include connection closing messages, because such virtual connections are never closed before the timeout is reached.

- If the protocol is not actually connection-oriented (for example, if the protocol is SNMP), you can disable connection tracking for the traffic in the Access rule's Action options. This requires that you explicitly allow both directions of the communications in the rule, since without connection tracking, reply packets cannot be automatically allowed. NAT rules are not applied to connections that are not tracked. We recommend that you deactivate logging in rules that have connection tracking off, since these rules will create a separate log entry for each packet transmitted. This greatly increases the number of log entries generated and can potentially lead to an unmanageable level of logging traffic.

Incomplete Connection Closed

Problem description: You see “incomplete connection closed” messages in the logs.

Reason: This information message indicates that the Firewall allowed a connection and passed the first packet of a connection (the SYN packet). However, the reply packet (SYN/ACK) from the destination host did not arrive at the Firewall, so the Firewall determined that the connection was unsuccessful and removed it from its records.

This can happen in one of the following situations:

- The SYN packet did not reach the destination.
- The SYN packet reached its destination, but the destination host did not send any reply.
- The SYN packet reached its destination and the destination host replied, but the reply packet did not reach the Firewall.

Solution: It is normal to see a few of these messages in the log from time to time, but a higher number of these messages may indicate problems in your network or the communicating applications.

If this message appears in the logs often for legitimate traffic, there is a networking problem that you must address. Use normal network troubleshooting tools to find out where the packets are lost. See Taking a Traffic Capture (page 126) for information on how to generate a tcpdump file from the Stonesoft Management Center.

In some cases SYN packets may be sent maliciously to random hosts as an attempt to find out your network structure. These attempts may sometimes be seen as SYN packets to hosts that do not exist, which may trigger the Incomplete Connection Closed messages if access to those addresses is allowed and routable. The possibility of successful scans can be reduced by using dynamic NAT on the Firewall. See Editing Firewall NAT Rules (page 687).

NAT Balance: Remote Host Does Not Respond

Problem description: Connections are dropped with NAT balance messages in the logs.

Reason: A connection has been allowed, the Firewall has applied a NAT rule that defines source and/or destination translation, and the traffic has been forwarded according to the Firewall’s routing configuration, but a reply is never received.

Solution:

1. If NAT is applied to the connection in error, adjust your NAT rules accordingly. It is also possible to create a NAT rule that defines no translation to disable NAT for any matching connection.
2. Check that the Firewall routes the traffic correctly. The routing decision is made based on the translated destination IP address.
3. Make sure the destination host is up and providing the requested service, and that any intermediary Firewall allows the connection.
4. Try to trace the path that the communications take and use traffic captures as necessary to find the point of failure.

Not a Valid SYN Packet

Problem description: The “Not a Valid SYN Packet” message appears in logs in conjunction with entries on discarded packets.

Reason: “Not a Valid SYN Packet” is a TCP packet that is not the first packet of a TCP connection (the packet does not have the SYN flag set), but is not part of an existing connection either (there is no connection tracking entry on the Firewall matching this packet). This packet would be allowed by the policy if it was part of an existing tracked connection.

The messages usually also contains a code inside square brackets that indicates the flags set in the discarded packet (A=Ack, F=FIN, R=RST, P=Push, S=SYN).

Some examples of situations, where “Not a Valid SYN packet” messages can be seen:

- Asymmetric routing, which means that the opening packet does not go through the Firewall, but the reply (the SYN/ACK) does. If this is the case, there is a configuration error in the routing of the surrounding network that must be fixed.
- Connections that are idle for more than the defined connection timeout (connection has been erased from the Firewall records). This will always happen from time to time, but if necessary, the timeout can be increased or you can use the TCP Proxy Protocol Agent, which can reset the connection at the host and server when the connection is left idle (if the application in question does not properly close connections).
- Connections that have been made to look like TCP connections while they are not. If necessary, these can be allowed as individual packets without connection tracking.
- Network scans or attacks that use ACK packets.
- Heavily loaded server or client that sends a packet after the host at the other end of the connection has already timed out and closed the connection.

Solution: It is normal to see some messages like this in the logs. If a certain type of communication that you want to allow is always prevented because of connection tracking:

- If there are connections that are left idle for a long time, you can modify the idle timeout value for the Access rule that allows that specific traffic. See Defining Firewall Allow Action Options (page 678). There are also default values that you can set globally for different TCP connection states in the Firewall element’s properties (Advanced settings). This solution does not usually apply to non-TCP connections, so take care that the rule only matches the specific connections involved.

Caution – Setting long idle timeouts for a high number of connections considerably increases the resource consumption of the Firewall and may even lead to performance problems. Especially, non-TCP protocols do not include connection closing messages, so such virtual connections are never closed before the timeout is reached.

- You may have to disable connection tracking in the rule allowing the connection. We recommend that you deactivate logging in rules that have connection tracking off, since these rules will create a separate log entry for each packet transmitted. This greatly increases the number of log entries generated. NAT cannot be applied to traffic that is allowed without connection tracking and both communication directions must be explicitly allowed in the Access rules (replies are not automatically allowed).
- For some types of connections, problems may be solved by using a Service that includes a special Protocol Agent for that kind of traffic. See Using Protocol Elements (page 738) for a list of Protocol Agents.

Requested NAT Cannot Be Done

Problem description: The logs show the error message “Requested NAT cannot be done”

Solution:

- A Dynamic NAT operation may be applied to the wrong type of traffic. Dynamic (many-to-one) NAT is done by assigning different hosts the same IP address, but different ports. For this reason, dynamic NAT does not work when the protocol in question does not use ports. Only the TCP and UDP transport protocols use ports. See the **TCP** and **UDP** branches in the **Services** tree in the Management Client to check which protocols are transported over TCP or UDP.
- Dynamic NAT may run out of ports if there are too many simultaneous connections in relation to the IP addresses and the port range you have configured for dynamic NAT. You can increase the available ports for translation by adding a new IP address for your dynamic NAT rule or by expanding the port range, if the rule does not currently use the whole range of high ports.
- If the Server Pool element is used, check the NAT rules. Because the Server Pool element always does NAT, errors may occur when the Server Pool element is used and the same connection matches an overlapping NAT rule.
- Check if the information message in the log states that dynamic NAT is denied due to excessive number of connections. This may happen when a single host is opening connections at an excessive rate to a single destination IP address and port through dynamic source NAT. This message indicates the triggering of a self-protection mechanism, which prevents excessive use of processing resources to dynamic NAT operations. Set up a static NAT rule to allow these types of connections if it is not possible to adjust the connection settings of the application.

Spoofed Packets

See Packets Are Dropped as Spoofed (page 1128).

IPsec VPN Log Messages

See Reading IPsec VPN-Related Logs (page 1138) for general information and IPsec VPN Log Messages (page 1241) for a listing of messages with descriptions.

Command Failed/Connect Timed out

Problem description: You try to command an engine using the Management Client and receive an error.

Solution:

- Make sure that there is connectivity between the Management Server and the engine. Most common connectivity problems include traffic filtering by an intermediate firewall and incorrectly configured NAT configuration (a required NAT rule for the connection is missing or the engine's contact address is missing). The engine sends its status reports through the Log Server, so a green operating status is no guarantee that the Management Server can reach the engine.
- Try to refresh the engine's policy. Read all messages displayed and make sure none of the nodes perform a rollback to the previous policy. If policy installation fails. See Troubleshooting Policy Installation (page 1124). If this is a cluster that contains nodes that are out of commission.

PKIX Validation Failed

Problem description: A Java-related error appears with various cryptic messages about "PKIX".

Reason: One or both of the certificates needed for securing the internal connection has expired or is damaged.

Solution: See Replacing Expired/Missing Certificates (page 1094).

Policy Installation Errors

Problem description: Policy installation displays an error in the main policy installation panel that logs the installation progress.

Solution: See Troubleshooting Policy Installation (page 1124) or Troubleshooting Rules (page 1126).

Unexpected Error

Problem description: Some action in the Management Client displays a pop-up message stating that an "unexpected error" has occurred**.**

Solution: Exit and relaunch the Management Client. If the condition persists, contact technical support.

Chapter 68

Troubleshooting Certificates

Digital certificates allow components to prove their identity to each other. Stonesoft Management Center components use certificates in system communications and in VPNs.

The following sections are included:

Note – Do not confuse certificates with licenses. Certificates are proof of identity that components use to authenticate themselves in communications. Licenses are a proof of purchase used for ensuring that your organization is a legal license holder of the software. For troubleshooting licenses, see Troubleshooting Licensing (page 1105).

There are two types of certificates that are used in the Stonesoft Management Center:

- Components use certificates to identify each other in system communications. When a component has no valid certificate, contact between that component and other Management Center components is not possible. Depending on the components, policy installation and log traffic may be affected. The certificates used in system communications are always generated by the Internal Certificate Authority that runs on the Management Server.
- Certificates are also used in VPNs for authentication between remote gateways. These certificates may be generated by any internal or external certificate authorities that both gateways are configured to trust.

Problems Related to Certificate Validity Time

All certificates have a validity start date (“not before”) and a validity end date (“not after”). In the Stonesoft Management Center, internally generated certificates are valid for three years from their creation. Expired certificates must be replaced with new ones. Note that since each component interprets the certificate validity according to their own internal clock, incorrect date or time settings may lead to unexpected certificate rejections.

Example **An external VPN gateway device defaults to 1970 when an error occurs in setting its time. In this state, the device thinks a certificate issued today is not valid until decades into the future and certificate authentication fails.**

You can check the validity of internal certificates (of SMC servers, Firewalls, IPS engines, and Layer 2 Firewalls) and VPN Gateway Certificates in the Management Client:

- Internal certificates: Select **Configuration→Configuration→Administration**, navigate to **Other Elements→Internal Certificates**, and open the properties of the certificate.
- VPN Gateway Certificates: Select **Configuration→Configuration→VPN**, navigate to **Other Elements→Certificates→Gateway Certificates**, and open the properties of the certificate.

Problems Related to Certificate Authority Validity Time

The Management Center’s internal Certificate Authorities, Internal CA and internal VPN CA, are both valid for ten years. A new Internal CA and a new VPN CA are automatically created six months prior to the internal CA’s expiration date. The components that use certificates signed by the internal CAs must receive new certificates that have been signed by the new internal CAs. Otherwise, the system communications and VPN connections cannot function.

You can check the validity of internal CAs in the Management Client:

- Internal Certificate Authorities: Select **Configuration→Configuration→Administration**, navigate to **Other Elements→Internal Certificate Authorities**.
- Internal VPN CAs: Select **Configuration→Configuration→VPN**, navigate to **Other Elements→Certificates→VPN Certificate Authorities**, and open the properties of the certificate authority.

When the system has created a new Internal CA, system components gradually start using the new Internal CA to sign certificates. Initially, the new Internal CA is in the "Ready to Use" state, and only Management Server certificates are signed by the new Internal CA. Certificates for other components are signed by whichever Internal CA is currently used by the Management Server. When all Management Servers are using the new Internal CA, the new Internal CA changes to the "Active" state.

Each component must receive a new certificate signed by the new internal CA. The system automatically tries to create new certificates for Security Engines, Master Engines, and SSL VPN gateways. For other components, you must always manually create new certificates. If the automatic certificate creation fails, you must create new certificates manually for Security Engines, Master Engines, and SSL VPN gateways.

When a new internal VPN CA has been created, the VPN gateways that trust the old VPN CA must be made to trust the new VPN CA. VPN clients that use certificates for user authentication also require new certificates signed by the new VPN CA.

Other Problems

Common certificate-related issues not related to the validity time include:

- Certificates used in system communications become invalid when the certificate authority changes.
  - This happens if the Management Server is reinstalled and the configuration is recreated manually or by importing elements instead of importing a backup (backups contain certificate authority information). If backup restoration is not an option, all system components must receive a new certificate for system communications.
  - In some cases, also restoring the Management Server backup may result in the internal certificate authority being different from the certificate authority that was used to create certificates for some components. The invalid certificates must be replaced with new ones.
  - Note that in this case, internally signed VPN certificates can be made valid by configuring the issuing certificate authority as an additional trusted certificate authority for the VPNs.
- Certificates created for the security gateways for establishing the VPN during the VPN configuration are stored on the gateway devices (Firewall/VPN engines). These certificates are not included in the Management Server backup, and are not changed in any way when a Management Server backup is restored.
- Certificates may become unusable if the private key for that certificate is lost. This can happen, for example, if the Firewall engine hardware fails and needs to be replaced. Firewall Clusters share each VPN certificate and can synchronize the private key from node-to-node as needed. If the private key is erased from a Single Firewall or from all the nodes of a Firewall Cluster, a new certificate must be created.

- Externally issued VPN certificates may be revoked by the certificate authority that issued them. This safety measure is used when the certificate is suspected to be compromised. This problem can only be solved by contacting the certificate issuer.

What's Next?

- ▶ To recreate a certificate for an SMC server component, see Renewing SMC Server Certificates (page 1094).
- ▶ To recreate a system communications or VPN certificate for an engine, see Renewing Engine Certificates (page 1096).
- ▶ To make Firewall/VPN engines accept VPN certificates signed by a new or different certificate authority, see Defining a VPN Certificate Authority (page 973).

Related Tasks

- ▶ Checking When Internal Certificates or Internal CAs Expire (page 130)
- ▶ Checking When an Internal VPN CA Expires (page 982)

Replacing Expired/Missing Certificates

Renewing SMC Server Certificates

Problem descriptions:

- The system indicates that the certificate of a Management Server, Log Server, or Web Portal Server has expired.
- The system indicates that the Certificate Authority that signed the certificate of a Management Server, Log Server, or Web Portal Server is about to expire, a new Certificate Authority has been created, and the server requires a new certificate.
- The components refuse communication attempts with each other.

If the Management Server certificate expires or is lost, it is not possible to log in using Management Clients. Log Server certificate expiration or loss prevents log browsing, reporting, and status monitoring from working correctly, and forces the engines to spool logs locally.

Solution: The certificates of any of the SMC server components can be renewed without impacting any of the other components. The certificate generation requires authentication in the form of an administrator login (unrestricted administrator permissions required).

▼ To generate new certificates

1. Stop the SMC server you want to recertify.

Note – To certify a Log Server or a Web Portal Server, the Management Server must be running and accessible through the network.

2. On the server you want to certify, open the `<installation directory>/bin` folder.

Note – If you installed the Management Server in the `C:\Program Files\Stonesoft\Management Center` directory in Windows, some of the program data may be stored in the `C:\ProgramData\Stonesoft\Management Center` directory. Command line scripts may be found in the `C:\Program Files\Stonesoft\Management Center\bin` and/or the `C:\ProgramData\Stonesoft\Management Center\bin` directory.

3. Run the correct script for the server type:
   - Management Server: `sgCertifyMgtSrv.[bat|sh]`
   - Additional Management Server: `sgOnlineReplication.[bat|sh]`
   - Log Server: `sgCertifyLogSrv.[bat|sh]`
   - Web Portal Server: `sgCertifyWebPortalServer.[bat|sh]`
4. When prompted, log in using an administrator account with unrestricted permissions.
   - If Domains are in use, you can also specify the **Domain** the Log Server or the Web Portal Server belongs to. If you do not specify the Domain, the Shared Domain is used.
5. Make sure the **Recertify an Existing Server** option is selected and that the correct server is selected in the list below.
6. Click **OK** and wait for a confirmation.
7. Start the SMC server.

No further action is required. When you restart the server, all other components accept the new certificate, since it is issued by a certificate authority they trust. In this context, components only trust the Internal Certificate Authority that issued also their own certificate. Administrators who log in to the Management Client or to the Stonesoft Web Portal receive a notification of the certificate fingerprint change on the Management or Web Portal Server when they log in for the first time after the server's certificate is changed. If you want to check the certificate fingerprint before accepting it, run the `sgShowFingerprint` command on the server. See Command Line Tools (page 1145).

Related Tasks

- ▶ See Management Center Commands (page 1146) for more information on the command line tools on the servers.
- ▶ Checking When Internal Certificates or Internal CAs Expire (page 130)

Renewing Engine Certificates

Problem descriptions:

- The system indicates that an engine component's certificate has expired.
- The system indicates that the Certificate Authority that signed the component's certificate is about to expire or has expired, a new Certificate Authority has been created, and the engine requires a new certificate.
- Components refuse connection attempts with each other.

If the certificate for system communications expires or is deleted, the Firewalls, IPS engines, and Layer 2 Firewalls continue processing traffic normally but all communications with other components will cease. In the case of a cluster, traffic may be disturbed if expired certificates prevent nodes from synchronizing information. The same also happens if the Internal Certificate Authority that signs the certificates for system communications is about to expire, and engines do not have new certificates signed by the new Internal Certificate Authority that the system has automatically created. See Dealing with Expiring Certificate Authorities (page 1097).

If a certificate for VPN authentication expires or is deleted, VPNs that are configured to use certificates (this includes any VPN involving Stonesoft IPsec VPN Clients) cannot be established before the certificate is renewed. A VPN certificate renewal or change may cause a security warning to appear to VPN client users connecting to the gateway.

Reason: Automatic certificate renewal may be disabled in your system or may fail due to the lack of required information or due to information loss or incompatibility issues that depend on the software versions used when the system was originally installed.

Solution: Renew the certificates manually. Renewing the certificate requires a reboot, but otherwise your system and the network services are not affected by the certificate change.

To generate new certificates:

- For the engine certificates used for SMC system communications, renew the contact between the engine and the Management Server using a new one-time password. See Connecting Engines to the Management Center (page 513) and Reconfiguring Basic Engine Settings (page 233).
- The VPN gateway certificates are created and managed separately from the certificates used in system communications. If a VPN certificate has expired, renew the certificate as explained in Renewing VPN Certificates (page 979). If the VPN certificate is invalid for some other reason, delete the Gateway Certificate element in the Management Client and then create a completely new certificate as explained in Creating and Signing VPN Certificates (page 975).

Related Tasks

- ▶ See Engine Commands (page 1157) for a command line tool listing.
- ▶ Checking When Internal Certificates or Internal CAs Expire (page 130)
- ▶ Checking When Gateway Certificates Expire (page 981)

Dealing with Expiring Certificate Authorities

Problem description:

- The system indicates that the Certificate Authority that signed the certificate of a component is about to expire, a new Certificate Authority has been created, and a new certificate must be generated for the component because the Certificate Authority has been renewed.
- The system indicates that a VPN gateway must be set to trust a new internal VPN Certificate Authority that was created automatically.
- Components refuse connection attempts with each other.

Reason: The messages are triggered to show that the expiration date of a Certificate Authority approaches, a new Certificate Authority has been automatically created, or that a Certificate Authority has expired. The messages may also indicate that administrator actions are required so that the system components can start using the new Certificate Authority.

Solution: If a Certificate Authority is about to expire, the components that use certificates signed by the Certificate Authority require new certificates that are signed by a valid Certificate Authority.

- Renewing an external certificate authority used in VPN configurations:
  - Configure a new certificate authority and make sure it is a trusted certificate authority in the VPN configurations.
  - Create new certificates for the components involved in the VPN configuration, signed by the new certificate authority. See Managing VPN Certificates (page 971) for more information.
- Renewing internal certificate authorities used in securing system communications and in VPNs:
  - The system automatically generates a new Internal Certificate Authority and a new internal VPN Certificate Authority six months before their expiration dates.
  - Each component that uses certificates signed by the Internal Certificate Authority or the internal VPN Certificate Authority requires a new certificate that is signed by the new Internal Certificate Authority or internal VPN Certificate Authority.
  - See the sections that follow for information on renewing certificates for different types of components.

System Communications Certificates for the SMC Servers

You must manually create a new certificate for each Management Server, Log Server, and Web Portal Server using command line tools. See Renewing SMC Server Certificates (page 1094).

System Communications Certificates for the Engines

No administrator actions are usually needed to create a new certificate for the engines, including SSL VPN Gateways. Once the new Internal CA has been created, new certificates are automatically created for the engines. However, if the automatic certificate renewal fails for certificates used in system communications, you must make initial contact from the engine to the Management Server. See Connecting Engines to the Management Center (page 513). After successful initial contact, a new certificate is automatically created for the engine.

Internal VPN Gateways

A new certificate is automatically created for the internal gateway. You only need to refresh the policy on the engine.

External VPN Gateways

If an external gateway trusts the internal VPN CA and the internal VPN CA has been renewed, you must create a new certificate for the external gateway and sign it with the new internal VPN CA. You must also set the new VPN CA as a trusted CA in the External Gateway's properties and also in the properties of the VPN Profile element that is used in the VPN configuration. See Managing VPN Certificates (page 971) for more information.

VPN Clients

When a VPN client connects to a gateway that trusts the new internal VPN CA, the VPN client users must accept the fingerprint of the gateway's certificate before they can connect to the gateway. Inform the users that the fingerprint of the gateway's certificate has changed and provide the information for checking the fingerprint in the VPN client.

If certificates are used to authenticate VPN client users and the certificates have been signed with the internal VPN CA, you must create new certificates for the VPN client users using the new VPN CA. See Renewing VPN Certificates (page 979) for more information.

Reverting to an Internal RSA Certificate Authority

Problem Description:

You created a new Internal ECDSA Certificate Authority as explained in Creating a New Internal ECDSA Certificate Authority (page 348), but you want to go back to using an Internal RSA Certificate Authority.

Solution:

1. Select **Configuration→Configuration→Administration**. The Administration Configuration view opens.
2. Browse to **Other Elements→Internal Certificate Authorities**.
3. Click the Tools icon in the toolbar and select **Create New Internal RSA Certificate Authority**. You are prompted to confirm that you want to create a new Internal RSA Certificate Authority.
4. Click **Yes**.
   - The Internal RSA Certificate Authority creation process begins and a new tab opens to show the progress of the process.
   - When the process is finished, the progress report shows the steps you must take next.
   - The status of the RSA CA is **Created for Different Certificate Type**.
5. Restart the Log Server, Authentication Server, and Web Portal Server. See Management Center Commands (page 1146) for more information.
6. Start the **Renew Internal Certificate Authorities Task**. See Starting Tasks Manually (page 1040).
   - When the Task is finished, the status of the RSA CA is **Ready to Use for Different Certificate Type**.
7. Right-click the Renew Internal Certificate Authorities Task in the History branch and select **Show Progress Report**. The progress report shows which steps you must take next.
   - Follow the instructions to resolve any issues. For example, you may be prompted to check the status or connectivity of some engines.
8. Recertify the Management Server as explained in Renewing SMC Server Certificates (page 1094).
9. Start the **Renew Internal Certificate Authorities Task** again.
   - When the Task is finished, the status of the RSA CA is **Active**.
10. Recertify the Log Server, Authentication Sever, and Web Portal Server as explained in Renewing SMC Server Certificates (page 1094).

What’s Next?

▶ Disable 256-bit encryption on the engine. See Enabling 256-bit Security Strength for Engines (page 234).

Chapter 69

Troubleshooting Engine Operation

Issues covered here concentrate on errors and problems that are directly related to the operation of Firewalls, IPS engines, Layer 2 Firewalls, and Master Engines. Issues related to configuring individual features, such as issues related to designing and installing policies, are covered in separate sections.

The following sections are included:

Node Does not Go or Stay Online

Problem description: when you command a node online, it does not go online or turns itself offline shortly after going online.

Solution:

If you have just updated the engines or you are using an evaluation license:

- Open the Administration Configuration view and browse to **Licenses** to check that your licenses are valid for the version of engines you are using. If you need new licenses, generate them at the Stonesoft website. See Generating New Licenses (page 1047).

If the nodes are in a cluster, and only one node at a time stays online:

- Check whether the cluster is in Standby mode. Standby mode keeps one node online at a time and uses the other nodes as backups in case the online node fails. See Adjusting Firewall Clustering Options (page 556), and Adjusting IPS Clustering Options (page 562).
- Refresh the policy of the cluster and check that the installation is successful so that no nodes roll back to the previous configuration. All nodes in the cluster must have the same configuration that has been installed in the same policy installation operation. You may have to adjust the rollback timeout in the cluster's properties if policy rollback on some node is the problem. See Adjusting Firewall System Parameters (page 553), Adjusting IPS Engine System Parameters (page 560), and Adjusting Layer 2 Firewall System Parameters (page 564). If policy installation fails, see Troubleshooting Policy Installation (page 1124).
- Check for alerts in the Logs view about tests failing. Check if any of the failed tests are configured to turn the node offline when they fail. The tester will leave one node in a cluster online even if the test fails on all nodes. If you see a test failure, it may indicate a genuine problem you need to solve or the test may be misconfigured and may need to be disabled or reconfigured. See more information in Getting Started with the Engine Tester (page 520).

See the Logs view for any alerts or logs regarding the functioning of the nodes.

- Certain internal error conditions (for example, heartbeat connection failures or missing certificates) may cause nodes to go offline. These events are shown as logs and alerts.

Error Commanding an Engine

Problem description: you try to command an engine using the Management Client and receive an error.

Solution:

- Refresh the engine's policy. If there are disabled nodes in the same cluster, enable them before refreshing the policy. Read all messages displayed and make sure none of the nodes roll back to the previous policy. If policy installation fails, see Troubleshooting Policy Installation (page 1124).
- If policy installation fails and the status and statistics information displayed in the Management Client is missing (status is displayed as unknown) even though the engine is online and processing traffic, the problem may be due to an expired certificate. See Understanding Certificate-Related Problems (page 1092).

Errors with Heartbeat and Synchronization

Problem description: you receive alerts regarding the heartbeat connection between the nodes in a cluster.

Solution:

- Apply normal network troubleshooting (check speed and duplex settings, cabling, etc.) to ensure that the heartbeat link works reliably. See Taking a Traffic Capture (page 126) for information on how to generate a tcpdump file from the Stonesoft Management Center.
- You should have a primary and a backup heartbeat connection using completely separate physical links.
- It is highly recommended that you use a dedicated link for both the primary and the backup heartbeat. The heartbeat and state synchronization are time-critical communications. The heartbeat connection is absolutely critical to the operation of a cluster. The cluster cannot work without a reliable heartbeat connection.

If you have installed two or more clusters with a single LAN as a shared heartbeat, and you see extra log entries about unauthenticated heartbeat messages:

- Change the Heartbeat IP and the Synchronization IP so that each cluster uses a different address. See Adjusting Firewall Clustering Options (page 556), Adjusting IPS Clustering Options (page 562), and Adjusting Layer 2 Firewall Clustering Options (page 566).

Problems Contacting the Management Server

Problem description: the engine cannot establish initial contact to the Management Server, or all subsequent attempts to command the engine through the Management Client fail.

Solution:

- Apply normal network troubleshooting (check speed and duplex settings, cabling, etc.). See Taking a Traffic Capture (page 126) for information on how to generate a tcpdump file from the Stonesoft Management Center.
- If there is a local Firewall between a remote site Firewall and the Management Server, the local Firewall blocks the communication.
- For a full list of all system communications in all configurations, see Default Communication Ports (page 1167).

Example **A Firewall with reversed management connections (for example, because it has a dynamic IP address) contacts the Management Server on port 8906. You must create an Access rule in the policy of the main site Firewall to allow the connection:**

- **Source**: remote site Firewall
- **Destination**: contact address of the Management Server
- **Service**: SG-dynamic-control
- **Action**: Allow

CHAPTER 70

TROUBLESHOOTING LICENSING

Licenses are a proof of purchase used for ensuring that your organization is a legal license holder of the software.

The following sections are included:

Troubleshooting Licensing

Note – Do not confuse Licenses with Certificates. Licenses are a proof of purchase (given to you by Stonesoft) used for ensuring that your organization is a legal license holder of the software. Certificates are proof of identity (generated in your installation of the system) that components use to authenticate themselves in system communications. For troubleshooting certificates, see Troubleshooting Certificates (page 1091).

Licenses are introduced to the system as elements that you install in your Management Server. You must install licenses to set up the components your organization has purchased. Components that do not have a license do not work. However, on current appliance models running Firewall/VPN 5.0 and IPS 5.0 or above, the licenses may be generated automatically.

If the Management Server does not have a valid license, you are shown a license-related dialog each time you log in using the Management Client. You cannot create a working configuration without a Management Server license, since most controls are disabled. If an engine component is missing a license, you can create a configuration for it, but you cannot transfer that configuration to the engine.

Basic troubleshooting information and solutions for the most common problems with licenses are provided in the following topics:

- License Is Shown as Retained
- License Is Shown as Unassigned (page 1107)
- Throughput License Exceeded (page 1083)

Related Tasks

▶ Getting Started with Licenses (page 1044)

License Is Shown as Retained

Problem description: when you view the information for a license in the Management Client, the State column displays “retained”.

Reason: Licenses can be generated based on the Management Server's proof-of-license (POL) code, the appliance’s proof-of-serial (POS) code, or a fixed IP address. When you start using a Management Server POL-based license, you bind it to the correct component, and the binding is fixed when you install the component's policy. Because of this fixed binding, if you unbind a Management Server POL-based license, it is shown as ‘Retained’ and it is not possible to reuse the Management Server POL-based license on another component or delete the Management Server POL-based license.

Solution: complete one of the procedures listed below:

- Rebind the license to the component it was previously bound to.
  - Right-click the license and select **Cancel Unbind**.
- Replace the component's license with a new Management Server POL-based license.
  - Install a new Management Server POL-based license and bind it to the component that the retained license is bound to. The state of the previous Management Server POL-

based license changes to **Unassigned** and it can be rebound to some other component or deleted.

- Replace the component's license with a new IP-address-bound license or POS-bound license:
  - You can replace a Management Server POL-based license directly with an IP-address-bound license or POS-bound license. Right-click the Management Server POL-based license and select **Replace by Static License**. The IP-address-bound license or POS-bound must be installed before you select the **Replace by Static License** option.
- Delete the element to which the Management Server POL-based license is bound.
  - After the element is deleted, the state of the Management Server POL-based license changes to **Unassigned** and it can be rebound to some other component or deleted.

License Is Shown as Unassigned

Problem description: when you import a new license, it does not attach itself to any component and is shown as ‘unassigned’.

Reason: The license does not contain identifying information that the Management Server could use to attach the license to a component. This is normal for Management Server POL-based licenses (for more information on license types and states, see Checking License Validity and Status (page 1053)).

Solution: Right-click the license and select **Bind** to select which component you want to license. See Installing Licenses (page 1050) for more information.

Throughput License Exceeded Alerts

See Throughput License Exceeded (page 1083).

CHAPTER 71

TROUBLESHOOTING LOGGING

This section covers some common problems you may have with viewing logs or when performing tasks related to the log files. For information about particular messages you may see in the logs, see Log Messages (page 1084).

The following sections are included:

Problems With Viewing Logs

If the Logs view is unable to contact some Log Server:

- Check that the Log Server is running, and that the Log Server is reachable from the computer used for running the Management Client. Logged data is not routed through the Management Server, so a green status shown for a Log Server in the Management Client (based on information that the Management Server has) is not an indication of whether the Log Server is reachable for log browsing or not.
- If there is a NAT device between some Management Client and a Log Server, administrators must select the correct Location for the Management Client in the status bar at the bottom right corner of the Management Client window. See Getting Started with System Communications (page 64).

If some or all logs are not visible in the Logs view:

- Check that the filtering, time range, and source settings in the Logs view are correct and make sure you have clicked **Apply** after making the latest changes in these settings.
- Check the logging options in your policies. Not all connections automatically create log entries. The Alert, Stored, and Essential options create permanent log entries. The Transient option means that logs are not stored and they can only be viewed in the Current Events mode in the Logs view when they reach the Log Server.
- Check that logs you want to keep are not being pruned. See Pruning Log Data (page 1024). The Log Server deletes selected logs according to how pruning is configured.
- Check that the logs are being transferred from the engines to the Log Server. The log entries are spooled on the engines if a connection to the Log Server is unavailable. Connections between engines and the Log Server should be shown as green in the System Status view.
- Depending on your logging settings and the number of engines that send data to the same Log Server, the logging process may slow down due to a lack of resources on the engine, in the network, or on the Log Server. See Logs Are Filling up the Storage Space.

Logs Are Filling up the Storage Space

Problem description: the hard disk of an engine or a Log Server is filling up with log files as indicated by alerts you receive.

Solution:

If an engine is filling up with logs:

- Check that the Log Server is running. If it is not running, try to start it (see Log Server Does not Run (page 1111) if the Log server does not stay running). If the log server is running, check for network problems between the engine and the Log Server. The log entries are spooled on the engines if they cannot be sent to the Log Server. Stopping and restarting the Log Server process may help in resetting the connection.
- If the volume of logs is extremely high, they may not be transferred quick enough, and logs have to be spooled even though they are being transferred. If you suspect this is the case, turn off all diagnostics logs for all engines that you are not actively troubleshooting (see Enabling/Disabling Firewall/VPN Diagnostics (page 222)), turn off logging for all rules that have connection tracking set to off (because these rules log each packet individually), and check if logs that are currently pruned could be prevented from being generated in the first place.

If the Log Server is filling up with logs:

- Set up log management tasks that archive and remove the oldest Log entries from the Log Server hard disk, see Getting Started with Log Data Management (page 1016). To avoid problems in the future, set up tasks to run automatically at regular intervals.
- The logs are stored on the Log Server machine under a folder structure based on dates and times (default location is `<installation directory>/data/storage`). In an emergency, you can also move or delete old log entries manually. You should always avoid manual handling of the newest entries.

Log Server Does not Run

Problem description: the Log Server is not running and does not stay running when you try to start it.

Solution:

- To get more information, try to start the log server using the script `<intallation directory>/bin/sgStartLogSrv[.bat|.sh]`. Startup messages are shown on the command line and you may gain more information if the startup fails.

Possible issues and solutions include:

- The software version may be incorrect. The Log Server must have the exact same software version as the Management Server. Upgrade components as necessary.
- The Log Server's certificate for system communications may have expired, been deleted, or become otherwise invalid, see Understanding Certificate-Related Problems (page 1092).
- The Log Server may not have a license, the license may be bound to the wrong IP address, or the IP address the license is bound to may not be active on the server, see Troubleshooting Licensing (page 1106).
- There may not be enough space for logs on the hard disk, see Logs Are Filling up the Storage Space (page 1110).
- Sometimes files that are necessary for the Log Server to run may be moved or deleted by external processes, such as a virus scanners, or lost due to hard disk errors. In these types of cases, Java shows an error message stating "Could not find the main class". Reinstall the same software version and check the configuration of the host computer to prevent the same from occurring in the future.

Chapter 72

Troubleshooting the Management Client

This section concentrates on general problems you may encounter when using the Management Client.

The following sections are included:

Some Options Are Disabled

Problem description: Some options are disabled (you can see them, but they are grayed out or otherwise uneditable).

Solution:

- If nearly all controls are disabled, the Management Server license may not be installed or valid. See Generating New Licenses (page 1047) for instructions on how to create and import a license file. Make sure that the IP address in the license is correct and that the network interface with that IP address is active on the Management Server.
- Your administrator permissions may restrict your actions. For example, Operator administrators are only allowed to view the contents of the elements granted to them. They are not allowed to change the elements' properties. See Administrator Accounts (page 247) for more information. If you want to edit an element that is locked, you must first unlock the element. See Locking and Unlocking Elements (page 88).
- The disabled options may be part of an element or object created by the installation. These system objects cannot be edited, but usually you can make copies of them and edit and use the copies instead. Select the element in a Configuration view and check the **History** tab in the Info panel. System elements display "System" as the Creator.
- The options may be part of a feature that is not supported by the engine version you are using. If this is the case, you must upgrade the engines before you can use the feature. See the Release Notes for information on new and enhanced features in different versions at http://www.stonesoft.com/en/customer_care/support/. For information on new features introduced in this version of the Management Center, see New in This Release (page 31).
- Changing the disabled options may require selecting some other option first. Look for an option to activate the feature or override the default values, or a way to change the default values instead.

Slow Startup and Use

Problem description: The Management Client login and other operations are slow (e.g., refreshing a list of network elements takes a long time).

Solution: Make sure the Management Server and Log Server hostnames can be resolved on the computer running the Management Client (even if the Management Client is running on the same computer as the server components). Add the IP address and hostname pairs into the local hosts file on the client computer:

- In Linux: `/etc/hosts`
- In Windows: `\WINNT\system32\drivers\etc\hosts`

Log browsing is also slowed down if all selected Sources (Log Servers) are not available (data is displayed only after queries time out). Exclude temporarily unavailable Log Servers from Sources and make sure that the Location is set correctly for your Management Client (in the status bar) if there is a NAT device in between your client and a Log Server.

Problems Logging in to the Management Client

If the Management Client reports a connection problem:

- Make sure you entered the address correctly on the login screen. If there is a NAT device between the Management Client and the Management Server, the IP address used must be the translated address. Locations and Contact Addresses are not used for selecting the correct address when you are just logging in.
- Verify that the Management Server is running.
- Verify that the network between the Management Client and the Management Server is working and routing the traffic correctly. See Default Communication Ports (page 1167) for the ports and protocols used for the communications. Note that these communications are not allowed in the predefined Firewall Template for Firewalls.

If you see a certificate expiration message:

- The Management Server certificate has expired. See Renewing SMC Server Certificates (page 1094).

If you have forgotten your password:

- See Forgotten Passwords (page 1078).

Problems With Layout and Views

If a view that was previously visible is now missing:

1. Select the appropriate view from the **View** menu. If the view is not listed, it is not available in this window.
2. If the view is still not visible, select **View**→**Layout**→**Reset Layout**.

If all the views that you want to be displayed on the screen are visible but you want to change the layout:

- You can drag and drop the views by their titles to different predefined positions on the screen and resize them by dragging the borders.
- You can reset the layout to the default positions, select **View**→**Layout**→**Reset Layout**.

Problems With Viewing Statistics

If you can see some types of statistics for an element, but not other types:

- This is normal. Not all statistics are compatible with all element types.

If an element does not show any statistics at all:

- If this is a server element, this is normal. Servers do not produce statistical information. Only engine elements do.
- Check that the Log Server is selected correctly in the Management Server properties, that the Log Server is running, and that the network connections are up between the engines, the Log Server, and the Management Server. Statistical information is sent from the engines to the Log Server, which relays the information to the Management Server.

Problems With Status Monitoring

If a status icon for an element is always white:

- The status monitoring for the element is turned off. To turn it on, see Monitoring Task Execution (page 125).

If you do not know what the status icon means:

- Place the mouse pointer over the element and wait for a tooltip to show to view the element's status as text. For more information, see Monitoring the System (page 93).

If the status of the Management Server indicates problems, but the engine is processing traffic normally:

- Check network connectivity: status information is normally sent from the engines to the Log Server, which then relays the information to the Management Server. There is also a backup channel directly from the engines to the Management Server. See Default Communication Ports (page 1167).

If no information is available on the Appliance Status tab in the Info panel:

- Appliance Status information is not displayed automatically. Click **Refresh** on the Appliance Status tab to view the information. Appliance Status information is available for engine versions 4.3 or higher.

Problems Installing Web Start on an External Server

Problem description: The automatic Web Start installation using the `webstart_setup.vbs` or `webstart_setup.sh` scripts fails.

Solution: Manually configure the Web Start files.

▼ To configure Web Start files manually

1. Copy the Web Start folder to your workstation. The Web Start software can be found on your product DVD in the `Stonesoft_SW_Installer/Webstart` directory.
2. Edit the following files in a text editor and enter the correct path of the Web Start folder on your web server or network drive on the codebase row at the beginning of each file:

Table 72.1 Files to be Edited

File name	Example

`smcclient.jnlp`	`codebase="http://www.example.com/webstart"`
	Or
	`codebase="file://localhost/C:/webstart"`

`bouncycastle.jnlp`	`codebase="http://www.example.com/webstart"`
	Or
	`codebase="file://localhost/C:/webstart"`

`smc_help.jnlp`	`codebase="http://www.example.com/webstart"`
	Or
	`codebase="file://localhost/C:/webstart"`

3. Save and close each file.

4. Put a link to the smcclient.jnlp file on your web site. An example start page is provided with the Web Start package, containing the link to start the application.

5. If necessary, modify the configuration of the web server to return the appropriate MIME type for`.jnlp`-files (`application/x-java-jnlp-file`). Consult the manual of your web server for instructions on how to configure the MIME type.

Problems With Management Servers

Problem description: Commands sent to Management Servers using the Control Management Servers dialog fail. An error message indicating that the Management Client is unable to connect to a Management Server may appear.

Solution:

- Make sure the Management Client Location is selected correctly so that the Management Client connects to the Management Server at the correct Contact Address. See Selecting the Management Client Location (page 73).
- Make sure the Location and Contact Address are correctly configured for the Management Server. See Defining Server Contact Addresses (page 70).

Problem description: The Status of the Management Server on the General tab in the Info panel reads “Replication error”.

Solution: Retry automatically replicating the Management databases.

▼ To retry automatic database replication

1. Select **File→System Tools→Control Management Servers**. The Control Management Servers dialog opens.

2. Select the Location from which to send the command to ensure that the command is sent to the correct Contact Address for the Management Server.

3. Right-click the Management Server and select **Replication→Retry**.
   - If retrying automatic replication fails, see Synchronizing Management Databases Manually (page 339).

CHAPTER 73

TROUBLESHOOTING NAT

This section concentrates on some common problems you may encounter with NAT definitions.

The following sections are included:

Troubleshooting NAT Errors

Consider the following when troubleshooting NAT issues:

- A Dynamic NAT operation cannot be applied to all types of traffic. Dynamic (many-to-one) NAT is done by assigning different hosts the same IP address but different ports, so that the subsequent replies can be recognized and forwarded correctly. For this reason, dynamic NAT does not work when the protocol in question does not use ports (only one connection can be handled at a time). Only the TCP and UDP transport protocols use ports. See the **TCP** and **UDP** branches in the **Services** tree in the Management Client to check which protocols are transported over TCP or UDP. Use static translation for other protocols, or check if the communicating application supports NAT traversal (tunneling the protocol inside TCP or UDP).
- Dynamic NAT may run out of ports if there are too many simultaneous connections in relation to the IP addresses and the port range you have configured for dynamic NAT. You can increase the available ports for translation by adding a new IP address for your dynamic NAT rule or by expanding the port range, if the rule does not currently use the whole range of high ports. The number of simultaneous NATed connections equals the number of IP addresses multiplied by the number of ports.
- Check the NAT rules for definitions that overlap with the address translation done configured in an Outbound Multi-Link or Server Pool element or the NAT pool defined for VPN clients in the Firewall element’s properties. Errors may occur when one of the listed elements is used and the same connection matches an overlapping NAT rule, because the elements also use NAT. Naturally, only one address translation operation can be done for each packet and overlapping definitions may therefore cause conflicts. Overlap within the NAT rules alone is allowed, since the rules are resolved based on their order (first matching rule is applied).
- Check that the NAT definitions do not overlap with an IP address that is used by some physical host present in the network (this configuration error is most common in the case of source address translation for a DMZ or external IP address). IP address conflicts are not any less harmful to connectivity when NAT is involved.

NAT Is Not Applied Correctly

Problem description: NAT is not applied at all or is applied incorrectly.

Solution:

- Any connection that you want to NAT must be allowed by an Access rule with the connection tracking option active. No NAT of any kind can be applied to packets that are not tracked as part of a connection.
- If the target of the translation is traffic that is entering or exiting a VPN tunnel, enable address translation for traffic transmitted over that VPN in the properties of the VPN element (default setting is to disable all address translation for tunnelled VPN traffic). The setting affects only traffic wrapped inside the VPN tunnel, not the tunnel itself (the encrypted packets).
- If traffic is not translated at all or the wrong translation is applied, check the NAT rules: make a rule search with the original (before translation) source and destination address and check if the traffic matches the wrong NAT rule higher up in the rule table. Only the first matching rule is considered. Note that NAT rules with an empty NAT cell are valid and specify that addresses are *not* translated for matching traffic.
- In addition to NAT rules, NAT definitions are also used in a NetLink or Server Pool elements, and as a NAT pool defined for VPN clients in the Firewall element’s properties. There must not be additional NAT operations defined for the same connections in the NAT rules.

NAT Is Applied When it Should Not Be

Problem description: the Firewall translates an IP address to some other IP address even though it should not.

Solution:

- The Firewall reads the NAT rules from top to bottom and only the first rule that matches is considered, so you can make exceptions to rules by placing a different, partially overlapping rule above. Leaving the **NAT** cell empty tells the Firewall that addresses in any connections that match the rule should not be translated. For more information, see Editing Firewall NAT Rules (page 687).
- For VPN traffic, you can also enable and disable address translation for all traffic transmitted over a VPN in the properties of the VPN element (default setting is to disable all address translation for tunnelled VPN traffic). The setting affects only traffic wrapped inside the VPN tunnel, not the tunnel itself (the encrypted packets).
- Alongside the NAT rules, NAT definitions are also used in NetLink and Server Pool elements, and as a NAT pool defined for VPN clients in the Firewall element's properties. In the case of NetLinks, NAT rules are used to select traffic for balancing, and only the actual IP addresses used for the translation are defined in the NetLink elements. NAT is required for the operation of these features and you must exclude the connections in question from the scope of these features to disable NAT.

CHAPTER 74

TROUBLESHOOTING POLICIES

This section covers some common problems you may encounter when working with policies and the rules that they contain.

The following sections are included:

Troubleshooting Policy Installation

Problem description: the policy installation fails, or there is a warning message indicating problems in the policy installation window.

Solution: see below for most common problems and messages:

- The Engine Performs a Roll-Back at Policy Installation
- The Management Server Contact to Nodes Times Out
- Policy Installation Fails for Some Other Reason (page 1125)
- Warning Automatic Proxy ARP Option Is Ignored (page 1125)

The Engine Performs a Roll-Back at Policy Installation

Problem description: The policy installation reports that the Management Server can contact the engines and installs the new policy successfully, but then the policy installation results in a roll-back to the previously installed policy version.

Reason: The roll-back is a safety mechanism that prevents changing the engines' policy in ways that cut the connectivity between the engines and the Management Server. After each policy installation, the engine contacts the SMC using its new configuration and automatically reverts its policy if the contact does not succeed within a time-out period.

Solution:

- Make sure the policy and the configuration changes you have made do not prevent communications between the Management Server and the engine.
  - Check the IPv4 Access Rules and NAT rules (as applicable). You can additionally validate the policy to see if there are issues in it that prevent the policy installation. See Validating Rules Automatically (page 714). The rule search is useful for finding the first rule that matches the connections. See Searching in Rules (page 666).
  - Check the Routing. You can use the Route Query tool to check where the packets will be routed after a policy installation. See Checking Routes (page 608).
  - Check the Locations and Contact Addresses of the system components, which are required if NAT is applied to these system communications. See Getting Started with System Communications (page 64).
- The rollback occurs after a timeout set in the engine element's advanced properties. If you are sure that there are no configuration or policy design issues, you can increase the timeout to allow for longer delays in contact. This may help if the timeout is caused by poor network reliability or delays caused by processing a policy that is extremely large considering the engine's available resources. See Adjusting Firewall System Parameters (page 553).

The Management Server Contact to Nodes Times Out

Problem description: The engine is up and running, but policy installation fails when the Management Server is contacting the nodes or the Management Server is left waiting for contact from a node (when node-initiated contact is active) that never happens.

Reason: There is no network-level connectivity, the engine or the Management Server uses the wrong IP address, or the engine and the Management Server reject each others' certificates.

Solution:

- Check for network problems, such as faulty or loose cables, mismatching speed/duplex settings, IP addresses, and routing.
- Check the Locations and Contact Addresses of the system components, which are required if NAT is applied to these system communications. See Getting Started with System Communications (page 64).
- In a cluster, all nodes that the Management Server tries to contact must be reachable and operational to install a policy on any of the clustered engines. If you have taken an engine down for maintenance, you must temporarily disable it to install the policy on the other cluster members. See Disabling Nodes of a Cluster Temporarily (page 224).
- If the problem seems to be related to certificates, you can recertify the engine to re-establish contact between the engine and the Management Center. See Getting Started with Connecting Engines to the SMC (page 514).
- Check the software version on your Security Engines (shown in the Info panel when you select the element in the Management Client). See the release notes for information regarding version compatibility between the engine and SMC software versions.

Policy Installation Fails for Some Other Reason

- If the policy installation ends with an *error*, read the messages both in the main window and in the Issues tab below and correct as necessary. Note that *warnings* do not prevent policy installation; you can still press **Continue** in the policy installation even if a warning is triggered. You can also validate the policy separately to see if there are issues that prevent you from installing it. See Validating Rules Automatically (page 714).
- Try installing the policy with the "**Keep Previous Configuration Definitions**" option deselected. Normally, the option should be selected, but under certain conditions, the old configuration definitions may not be compatible with the new policy, so the engine cannot fulfill this request and policy installation fails.

Caution – When the policy is installed with the "Keep Previous Configuration Definitions" option deselected, even some currently active connections that are allowed in the new policy may be cut. The applications must then reopen the connections.

- Make sure that a current dynamic update package is imported and activated on the Management Server.

Warning Automatic Proxy ARP Option Is Ignored

- When installing a Firewall policy, the "Automatic Proxy ARP option in NAT rule <rule tag> is ignored: none of the CVI interfaces are directly connected to the network in question" warning is shown when proxy ARP has been defined but there is no matching CVI network configured in the Firewall element. Automatic proxy ARP is used in NAT to handle ARP requests to the translated IP address for hosts in networks that are directly connected to the Firewall. This warning can be due to incorrect IP-address or netmask setting, the (not directly connected) Network in question missing from the Routing view, or selecting the option for a NAT rule that involves an IP address for which the Firewall cannot act as an ARP proxy.
- Related settings can be configured in NAT rules (Editing Firewall NAT Rules (page 687)), in a Server Pool element (Configuring Outbound Multi-Link Settings (page 611)), in the Firewall element (Firewall Interface Configuration (page 421)), and in the Routing view (Getting Started with Routing (page 588)).

Validating Rules

You can automatically validate a policy and check the rules for invalid configurations, for example, if policy installation fails. See Validating Rules Automatically (page 714) for more information.

Rule That Allows ANY Service Does Not Allow All Traffic

Problem description: some connection you want to allow is stopped.

Solution:

In IPS policies, Access rules allow all connections by default. If a connection you want to allow is stopped because of an IPS Access rule, your Access rules contain a specific rule that stops these connections. If your problem is related to Inspection Rules and Exceptions, see Inspection Policy Produces False Positives (page 1126).

In Firewall and Layer 2 Firewall Access rules, even if you set the Source, Destination, and Service to ANY and set the rule to allow the traffic, certain connections may still be discarded.

- Connections with a protocol that assigns ports dynamically must be allowed using the appropriate Protocol Agent, so that the Firewall can track the assigned port.
- For a Protocol Agent to be used in a rule with ANY as the Service, there must be a matching rule with Continue as the action further up in the rule table with a Service in which the correct Protocol Agent is used.
- The Firewall Template contains a rule that does this for some, but not all protocols that use a dynamic port assignment. Add your own rules as necessary.
- Connections that violate the standard TCP connection sequence are dropped. If you must allow connections in your network for some application that implements TCP incorrectly, you may need to adjust or even disable connection tracking in the Access rules for those connections. See Defining Firewall Allow Action Options (page 678). We recommend that you disable logging for rules that have connection tracking set to off, because such rules create a new log entry for each packet.

Inspection Policy Produces False Positives

Problem description: the Inspection Policy produces alerts or terminates traffic that you consider to be normal.

Solution: (do one of the following)

- If a Situation is not valid in your environment under any conditions, change the action for the Situation to **Permit** on the Rules tab.
- If a Situation is not valid between some known hosts, add an Exception for the Source, Destination, and Situation that produce false positives and set the action to **Permit**. This can be done manually or based on a log entry through its right-click menu.
- If a custom Situation produces false positives, adjust the parameters in the Situation to better match the traffic pattern that you want to detect.

Enabling Passthrough for PPTP Traffic

To allow PPTP passthrough, add matching Access rule(s) with the following two services:

- The TCP Service for PPTP. The default Service element for PPTP uses the standard destination port 1723. Check the actual port used and create a new Service with a different port, if necessary.
- The IP Service for GE (IP protocol 47).

Make sure that the GRE traffic is not matched against any dynamic NAT rule, including the dynamic NAT rule required to load-balance connections between NetLinks in a Multi-Link configuration. Use static NAT instead if IP address translation is required or configure the communicating applications to encapsulate the traffic in TCP or UDP (NAT traversal mode).

Dynamic NAT cannot be applied because it uses ports to keep track of connections using the same IP address. GRE works directly on top of IP and does not have the concept of ports, so it is not possible to do the same with GRE, requiring a static translation that forms a fixed one-to-one relationship between an original and translated IP address (use a static IP address to IP address or network to same-size network mapping in the NAT rules).

Even with static NAT, some PPTP implementations require extra setup (for example, encapsulation of the packets) to work correctly when IP addresses are translated.

Traffic I Want to Allow Is Stopped by the Firewall

▼ To troubleshoot why traffic is not passing through the Firewall

1. In the Logs view, check whether the connection is logged.

- Add a quick filter for both the source and destination IP address of the traffic you want to allow in the Query panel and click **Apply**.
- If the logs show that the connection is discarded or refused by a rule, click the Rule Tag link in the log entry to check the rule.
- See Troubleshooting Alert, Log, and Error Messages (page 1081) for some common messages you may see in the logs.

2. Check the Access rules and NAT rules of the active policy for rules that match the same source, destination, and service.

2a. Open the Search Rules panel through the policy view's toolbar and drag and drop the corresponding elements to the search fields at the bottom of the rule table.

2b. Select **Show Only Matching Rules** from **Options** in the search panel.

2c. Deselect **Do Not Match ANY** from Options in the search panel.

2d. If several rules are shown, the topmost rule is the one that is applied, unless the **Source VPN** cell (in IPv4 Access rules) has a definition that does not match. The other cells are not used for matching, but define options for what happens when traffic does match.

3. If the first matching Access rule is set to allow the traffic, check that other parts of the rule are correct:

- Some protocols require the correct Protocol Agent, which is set by including the correct Service with the correct Protocol attached. In some cases, you may need to change the options of the Protocol Agent. See Using Protocol Elements (page 738).

- ANY rules do not use most Protocol Agents by default. See Rule That Allows ANY Service Does Not Allow All Traffic (page 1126).
- You can create new Services for any source or destination port or port range as needed.
- The Connection Tracking Action options define if stateful inspection is used and how strict the checks are. Connection tracking allows NAT rules to be applied to the connection and a rule table where reply packets do not need to be separately allowed. The Firewall checks that the communications follow the standards of the protocol used and discards invalid communications. If invalid communications need to be allowed, you may need to adjust connection tracking options. See Defining Firewall Allow Action Options (page 678).

4. If there is a matching NAT rule, make sure that they are applied correctly. Particularly, dynamic NAT must only be used for protocols that work on top of TCP or UDP, since dynamic NAT uses ports to keep track of the translated connections.
5. Check your routing configuration. If Routing is incorrectly configured on the Firewall, packets may be dropped because the Firewall has no route where to send them. See Getting Started with Routing (page 588).

Packets Are Dropped as Spoofed

The antispoofing rules are automatically generated based on your routing configuration. Generally, traffic is only allowed if the IP address seen in the communications corresponds to the IP address space that is defined for routing through that interface in the Routing view. Normally, communications require this routing information in any case for any reply packets to be correctly routed, but in cases where communications are one-way, you can make exceptions to the antispoofing in the Antispoofing view.

By default, the antispoofing tree is read by picking the most specific entry defined in the view (for example, a definition of a single IP address is picked over a definition of a whole network). If some IP address must be allowed access through two or more different interfaces, the definition for each interface must be at the same level of detail for the IP address in question.

Example **If Interface A contains a Host element for 192.168.10.101 and Interface B contains a Network element for 192.168.10.0/24, connections from 192.168.10.101 are considered spoofed if they enter through Interface B. If this is not desired, the Host element must be added also under Interface B (in addition to the Network element already included).**

The behavior explained above can be changed by setting the more general setting (network) as *Absolute* through its right-click menu in the antispoofing tree. This allows the address through the interface even if there is a more specific definition attached to some other interface.

Antispoofing also discards packets that are in a routing loop: if the Firewall accepts a packet, but then receives the exact same packet again through a different interface, the Firewall drops it. This makes no difference to the success of communications, but saves the Firewall and other equipment in your network from handling the same packet over and over again until it finally expires. If this is the case, you must correct the routing in your network. Often, routing loops are indicated by “NIC index changed” information in logs that discard the connection (the same packet enters the Firewall a second time, but through a different interface - usually because the device where the Firewall is configured to send the packet routes the packet right back to the Firewall).

Unsupported Definitions in IPv6 Access Rules

Elements used in the Source and Destination fields of the IPv6 Access rules must contain an IPv6 address. Elements that contain only an IPv4 address cannot be used, and cause an error message indicating the rule is invalid. Make sure that all elements used in the Source and Destination field of the IPv6 Access Rules contain an IPv6 address.

CHAPTER 75

TROUBLESHOOTING REPORTING

This section concentrates on common problems you may encounter when generating reports from raw statistical and log data stored on the Log Server.

The following sections are included:

Troubleshooting Reporting

Error messages for reports are shown in the Comment column of the Stored Reports view. Check the status of the report task there before you proceed with the troubleshooting.

This subject is covered in the following topics:

- No Report is Generated at All
- Empty Report Sections or Incomplete Data (page 1133)

No Report is Generated at All

Problem description: you try to run a report but no report is generated. There may be an error message displayed.

Solution:

- If the report seems to generate, but you cannot find it in the Stored Reports view, you may have inadvertently created the report without selecting the Stored option. Try creating the report again and make sure you have selected all the outputs you want to produce.
- If the report generation does not begin, and there is no error indicated, you may have accidentally defined an end time that is in the future (remember the delay in the Start Earliest field). Check the start time indicated for the report task. The time is interpreted according to the clock of the computer you are using to run the Management Client; make sure the time and timezone settings are correct.
- If the “out of memory” error appears, check if you have placed one or more IP address-based top rate items under progress sections in the Report Design:
  - Generating reports with such a design consumes large amounts of memory, as you are collecting full progress information for every IP address that appears in your logs over the chosen period of time (the unnecessary data is discarded only after the top items are selected at the completion of the report task).
  - To reduce the memory load, use a Drill-down top rate design, which first finds the top IP addresses and then gets the progress information on those.
  - Memory consumption can also be reduced by restricting the amount of data included in the report (for example, set a shorter time range).
- If the “unreachable server” error appears on the report task’s row after the date and time, some Log Server defined in the system is not running or is not reachable from the Management Server. If you have defined Log Server elements that do not currently represent any physical, running Log Server, check the **Exclude This Log Server from Statistics and Reporting** option in the properties of such Log Server elements.

Caution – Be very careful when excluding Log Servers from reporting. If you inadvertently select this setting for a Log Server that is in use, there is no warning that the generated reports are missing parts of the log data.

Empty Report Sections or Incomplete Data

Problem description: your report generates, but does not contain all the data that you think it should (there are empty sections or the data displayed seems incorrect).

Solution:

- If you get sections that say “No Data”, check that the log data that the section requires is in the active log data directory of your Log Server(s) (archived logs are not included in reporting). Open the Logs view, make sure only **Active Log Data** is selected for all Log Servers on the **Sources** tab, select the same time period and filter that you used in the report, and click **Apply**. The logs shown correspond to the log data available for generating the logs. Also take note of the next point below.
- If you are only missing information on traffic volumes (for example, Traffic by Source IP) check your Access rules. Traffic information is available only on connections that are handled by rules for which you have decided to collect accounting information
  - The collected accounting information is shown in the Logs view for log entries that have “Connection Closed” as their Event (see the Details panel for those entries). Note that connection closing may not be logged at all, depending on the logging options of your Access rules.
  - If accounting data has not been collected at all, traffic volumes are unknown and the report sections on traffic volumes return “No Data”.
  - If only some rules collect accounting data, only the traffic that matches those rules is included in the report sections on traffic volumes.
  - Alternatively, the report items listed under **Counters** can be used to generate reports on traffic volumes. The data for these items comes from stored summaries of the statistical data that you can view as live statistics in an Overview. This data is always stored and includes information on all traffic.
  - To start collecting information on traffic volumes from now on, see Defining Access Rule Logging Options (page 684).
- If you are missing a large amount of data (all data from one or more Log Servers), check in the properties of all Log Servers in use that the **Exclude This Log Server from Statistics and Reporting** option is not selected. Log Servers with this option selected are completely ignored in all reporting.
- The time range you enter is interpreted according to the clock and time zone setting of the computer you are using to run the Management Client. If these are not correctly set in the operating system, the report period may be different from what you intend.

Chapter 76

Troubleshooting Upgrades

This section concentrates on common problems you may encounter when upgrading the system components.

The following sections are included:

Upgrade Fails Because of Running Services

Problem description: you cannot upgrade because the upgrade reports that some services are still running.

Solution:

- Check the Services window in Windows and stop any SMC services that are still running (Management Server, Log Server, Authentication Server, or Web Portal Server).
- If all Services are stopped in the Windows Services window, but the upgrade still reports that services are still running, set the services to Manual startup and reboot the computer.

Stonesoft Management Center Installation Failed

Problem description: you see the following error message during an installation or upgrade: “Stonesoft Management Center installation failed”.

Reason: this message is shown if any anomaly is detected during the installation.

Solution:

To find the cause of the problem:

- Check the installation log in the installation directory for messages regarding the upgrade.

> **Note – If you installed the Management Server in the** `C:\Program Files\Stonesoft\Management Center` **directory in Windows, some of the program data may be stored in the** `C:\ProgramData\Stonesoft\Management Center` **directory.**

- If you are skipping versions, check the Release Notes for the version you are installing and make sure the upgrade you attempted is possible. The Release Notes will also list any known issues that may cause errors during the upgrade.

To solve problems indicated:

- Missing files are one of the most common errors. This can happen if you copy installation source files manually. Check that you have copied all necessary installation source files and folders an run the installation again. If you have not checked the integrity of the installation files, compare the checksum of your local files to the correct checksum at the Stonesoft website as instructed in Obtaining the SMC Installation Files (page 1057).
- Start the component in question. The message is shown for many types of errors, and the component may still be able to run without problems.
- If you are unable to get the component running, uninstall the existing installation and install the component as a new installation. In the case of an upgrade, you can restore all elements, licenses, certificate, etc. from a backup you have taken through the Management Client, the command line script, or in the Installation Wizard when you did the upgrade. The backup taken with the previous version is valid for restoration on the upgraded system (see the *Release Notes* for any version-specific limitations or exceptions).

CHAPTER 77

TROUBLESHOOTING IPSEC VPNS

This section concentrates on some common problems you may encounter when creating and managing IPsec VPNs.

The following sections are included:

Checking Automatic IPsec VPN Validation Results

The Management Client has automatic IPsec VPN validation that checks the settings you have selected are compatible with each other.

▼ To check IPsec VPN issues

1. Right-click the VPN element or the Route-Based VPN element and select **Preview VPN** or **Edit VPN**. The VPN editor opens in preview or edit mode.
2. Switch to the **Issues** tab in the info panel. If the Issues tab is not visible, select the **View→Issues** item in the menu.
3. If issues are displayed, read the descriptions and fix the problems that are described.

Reading IPsec VPN-Related Logs

The Firewall logs contain information on IPsec VPN negotiations. This section helps you find the relevant logs. For reference information on the most common IPsec VPN-related messages in the logs, see IPsec VPN Log Messages (page 1241). To view logs of IPsec VPN traffic, see Monitoring VPNs (page 969).

Log messages related to IPsec VPN negotiations contain the value *IPsec* in the **Facility** field, which is a good filtering criterion for viewing IPsec messages. The **IKE Cookie** and **IPsec SPI** fields contain identifiers related to each particular IPsec VPN instance, which helps further in reading and filtering the logs. The **Situation** and **Information Message** fields include the actual VPN-related events. If possible, examine logs from the devices at both ends of the IPsec VPN tunnel for more complete information.

Tip – Right-click a VPN log entry and select Search Related Events to see logs related to the same IPsec VPN negotiation.

You can collect more detailed information by enabling the IPsec diagnostics. For VPN clients, you should also enable authentication and DHCP relay diagnostics. See Enabling/Disabling Firewall/VPN Diagnostics (page 222) for instructions.

Additionally, log messages generated by Access rules (which are not included if filtering specifically for IPsec logs) may contain relevant information on the connections that the gateway processes, and whether policy-based VPN traffic is directed correctly to VPN tunnels by the policy.

A normal IPsec VPN tunnel negotiation proceeds as follows:

1. The negotiations start when a connection matches a rule in the Firewall Policy that triggers the VPN negotiation (or a similar mechanism at the other end).
2. The gateway at the source end of the connection or the VPN client (the *initiator*, I) contacts the gateway at the other end (the *responder*, R) to establish trust and exchange keys in the IKE Phase 1 negotiations.
3. If Phase 1 negotiation succeeds, IKE Phase 2 negotiations begin. At this stage, the gateways agree on further settings used for handling the connection.
4. If Phase 2 negotiations succeed, the VPN tunnel is ready and ESP or AH packet(s) (the actual traffic) can be seen in the logs. New connections that are opened through the VPN are logged using a VPN-specific log message “New Connection Through VPN”.

VPN Certificate Issues

Your internal gateway always needs a certificate if VPN clients connect to it. In a gateway-to-gateway VPN, certificates are required when the VPN Profile used includes a certificate-based authentication method (RSS Signatures or DSS Signatures).

Certificate Acceptance

By default, the gateways only accept certificates signed by your Management Server. To accept certificates from other sources, you must define the Certificate Authority (CA) that signed the certificate as trusted. See Defining a VPN Certificate Authority (page 973). By default, all Gateways and all VPNs accept any valid CA that you have configured in the system. You can configure the trusted CAs at the Gateway level (Defining Trusted CAs for a Gateway (page 929)), and at the VPN level (Defining Trusted CAs for a VPN (page 948)). A CA must be trusted on both levels to be accepted as a trusted CA for a VPN.

Creating, Signing, Renewing, Transferring to Gateways

Internally signed certificates are created, uploaded to the engines, and renewed automatically if automatic certificate management is selected in the properties of the internal Gateway element.

You can manually create certificate requests, import certificates, and sign certificate requests in the VPN Configuration view under the **Other Elements→Certificates** branch of the tree. Any certificate request you create is, by default, also immediately signed using the internal CA and uploaded to the engine. To disable this action (for example, to sign the certificate using an external CA), you must deactivate this option in the new certificate request you create.

To sign or upload a certificate, first display the certificates and then select the corresponding item in the Tools menu above the panel that displays the certificates.

For detailed instructions, see Getting Started With VPN Certificates (page 972). For information on troubleshooting VPN certificates, see Understanding Certificate-Related Problems (page 1092).

Problems with Internal to External Gateway VPN

Both policy-based VPNs and the Route-Based VPN are IPsec compliant, and can form a VPN with any other fully IPsec compliant device. Lists of VPN solutions that have been tested to be compatible are published by the VPN Consortium (www.vpnc.org) and ICSA labs (www.icsalabs.com).

Note – Make sure that you have successfully installed or refreshed the policy on all affected Firewall/VPN engines after you have made changes to any part of the VPN configuration.

When creating a VPN with an external gateway:

- There are no settings that would always work with a device of a certain brand and model. Most IPsec settings depend on user preference and there are many alternative settings that you can use regardless of the type of gateway.
- Although the settings can in large part be selected based on preference, always take special care to ensure that all VPN settings are exactly the same at both ends (for both gateways at both ends: typically, four definitions in all).

- Ensure that also matching networks and netmasks are defined at both ends. In the Stonesoft Management Center, all networks you want to be accessible through the VPN must be placed in a Site element attached to the correct Gateway element. The networks defined must be identical at the other end.
- One commonly missed setting is the SA (Security Association) setting, which can be per net or per host. Some gateways may not have an explicit setting for this. Find out the setting used.
- For third-party devices, check for parameters that are explicitly set in the VPN configuration in the Management Client but not on the other device (find out the default settings used).
- If the VPN works when the connection is initiated from one end, but not when initiated from the other, even though the Firewall's policy has rules for both ways, the problem may be due to overlapping, but mismatching lifetime or encryption domain (in the Management Center, the IP address definitions in Site elements).
- Matching the settings of different devices may sometimes be difficult because of differences in the user interfaces and terminology used. The Virtual Private Network Consortium tests the interoperability of IPsec compliant devices using a certain profile. Various manufacturers (including Stonesoft) have produced step-by-step instructions for setting up a VPN matching this profile. See http://www.vpnc.org/InteropProfiles/. Consult these documents to see how the same VPN is configured in various different devices.

Problems Connecting With a Stonesoft IPsec VPN Client

Client-to-gateway VPNs are only supported in policy-based VPNs created using VPN elements. It is not possible to create a client-to-gateway VPN in the Route-Based VPN.

If NAT is used and the configuration download succeeds, but the client cannot connect to the VPN gateway after that:

- If NAT is done between the Stonesoft IPsec VPN Client and the Firewall, you must set the Contact Address for interface(s) that are used as a VPN endpoint. The Contact Address tells the IPsec VPN Clients the external NATed address they must contact. See Defining Contact IP Addresses (page 66). After the changes, refresh the Firewall Policy and make sure the IPsec VPN Clients download a new configuration from the engine.

If NAT should be used to translate the Stonesoft IPsec VPN Client address, but NAT is not done:

- Check the VPN properties and see if the “Enable NAT with this VPN” option is selected. NAT is only done if the option is selected.
- Check that the NAT rules are correct. In most cases, NAT is performed using the NAT Pool address range defined in the Firewall element's properties (Advanced Settings tab) and the same traffic should not match any of the NAT rules in the Firewall's policy (except in some cases a rule that specifically defines that no NAT is performed on this traffic to prevent subsequent NAT rules from matching).

For any general problems:

- Make sure the IPsec VPN Client version is up to date. Older clients may have known issues preventing correct operation and may not support all features configured for the gateway.
- Check for any VPN-capable devices between the Stonesoft Firewall/VPN device and the Stonesoft IPsec VPN Client; these may sometimes attempt to take part in the VPN negotiations.
- Check the Firewall logs for clues. See Reading IPsec VPN-Related Logs (page 1138).

Traffic Does Not Use the Route-Based VPN

It is possible to create a half-configured Route-Based VPN by configuring only the Tunnel Interfaces and the Routing. This creates a black-hole routing situation in which traffic routed to the Tunnel Interfaces is silently discarded. No warnings are given when you install the Firewall policy, as the configuration is considered valid. Traffic is only sent into the Route-Based VPN after you define the Route-Based VPN tunnels. To complete the Route-Based VPN configuration, see Editing the Route-Based VPN (page 964).

Reference

In this section:

APPENDIX A

COMMAND LINE TOOLS

This appendix describes the command line tools for Stonesoft Management Center and the engines.

Note – Using the Management Client is the recommended configuration method, as most of the same tasks can be done through it.

The following sections are included:

Management Center Commands

Management Center commands include commands for the Management Server, Log Server, Web Portal Server, and Authentication Server. Most of the commands are found in the `<installation directory>/bin/` directory. In Windows, the command line tools are `*.bat` script files. In Linux, the files are `*.sh` scripts.

Note – If you installed the Management Server in the `C:\Program Files\Stonesoft\Management Center` directory in Windows, some of the program data is stored in the `C:\ProgramData\Stonesoft\Management Center` directory. Command line tools may be found in the `C:\Program Files\Stonesoft\Management Center\bin` and/or the `C:\ProgramData\Stonesoft\Management Center\bin` directory.

Commands that require parameters must be run through the command line (cmd.exe in Windows). Commands that do not require parameters can alternatively be run through a graphical user interface, and may be added as shortcuts during installation.

Caution – `login` and `password` parameters are optional. Giving them as Command Line parameters may pose a security vulnerability. Do not enter login and password information unless explicitly prompted to do so by a Command Line tool.

Table A.1 Management Center Command Line Tools

Command	Description
sgArchiveExport [**host=**<*Management Server Address* [\Domain]>] [**login**=<*login name*>] [**pass**=<*password*>] [**format**=**<*exporter format: CSV or*** *XML*>] **i=**<*input files and/or directories*> [**o=**<*output file name*>] [**f=**<*filter file name*>] [**e=**<*filter expression*>] [**-h** | **-help** | **-?**] [**-v**]	Displays or exports logs from archive. This command is only available on the Log Server. The operation checks privileges for the supplied administrator account from the Management Server to prevent unauthorized access to the logs. Enclose details in double quotes if they contain spaces. **Host** specifies the address of the Management Server. If the parameter is not defined, the loopback address (localhost) is used. **login** defines the username for the account that is used for this operation. If this parameter is not defined, the username `root` is used. **pass** defines the password for the user account. **format** defines the file format for the output file. If this parameter is not defined, the XML format is used. **i** defines the source from which the logs will be exported. Can be a folder or a file. The processing recurses into subfolders. **o** defines the destination file where the logs will be exported. If this parameter is not defined, the output is displayed on screen. **f** defines a file that contains the filtering criteria you want to use for filtering the log data. You can export log filters individually in the Management Client through **Tools→Save for Command Line Tools** in the filter's right-click menu. **e** allows you to type in a filter expression manually (using the same syntax as exported filter files). **-h**, **-help**, or **-?** displays information on using the script. **-v** displays verbose output on the command execution. **Example** (exports logs from one full day to a file using a filter): `sgArchiveExport login=admin pass=abc123 i=c:/stonesoft/Stonesoft/data/archive/firewall/year2011/month12/./sgB.day01/ f=c:/stonesoft/Stonesoft/export/MyExportedFilter.flp format=CSV o=MyExportedLogs.csv`

Table A.1 Management Center Command Line Tools (Continued)

Command	Description
sgBackupAuthSrv **[pwd=***<password>***]** **[path=***<destpath>***]** **[nodiskcheck]** **[comment=***<comment>***]** **[-h | --help]**	Creates a backup of Authentication Server user information. The backup file is stored in the *<installation directory>*/backups/ directory. Backing up the Authentication only backs up Users, not the configuration of the Authentication Server. The Authentication Server configuration is included in the Management Server backup. **pwd** enables encryption. **path** defines the destination path. **nodiskcheck** ignores free disk check before creating the backup. **comment** allows you to enter a comment for the backup. The maximum length of a comment is 60 characters. **-h** or **--help** displays information on using the script. Also see **sgRestoreAuthBackup**.
sgBackupLogSrv **[pwd=***<password>***]** **[path=***<destpath>***]** **[nodiskcheck]** **[comment=***<comment>***]** **[nofsstorage]** **[-h | --help]**	Creates a backup of Log Server configuration data. The backup file is stored in the *<installation directory>*/backups/ directory. Twice the size of log database is required on the destination drive. Otherwise, the operation fails. **pwd** entering a password enables encryption. **path** defines the destination path. **nodiskcheck** ignores free disk check before creating the backup. **comment** allows you to enter a comment for the backup. The maximum length of a comment is 60 characters. **nofsstorage** creates a backup only of the log server configuration without the log data. **-h** or **--help** displays information on using the script. Also see **sgRestoreLogBackup**.
sgBackupMgtSrv **[pwd=***<password>***]** **[path=***<destpath>***]** **[nodiskcheck]** **[comment=***<comment>***]** **[-h | --help]**	Creates a complete backup of the Management Server (including both the local configuration and the stored information in the configuration database). The backup file is stored in the *<installation directory>*/backups/ directory. Twice the size of the Management Server database is required on the destination drive. Otherwise, the operation fails. **pwd** entering a password enables encryption. **path** defines the destination path. **nodiskcheck** ignores free disk check before creating the backup. **comment** allows you to enter a comment for the backup. The maximum length of a comment is 60 characters. **-h** or **--help** displays information on using the script. Also see **sgRestoreMgtBackup** and **sgRecoverMgtDatabase**.

Table A.1 Management Center Command Line Tools (Continued)

Command	Description
sgCertifyAuthSrv	Contacts the Management Server and creates a new certificate for the Authentication Server to allow secure communications with other system components. Renewing an existing certificate does not require changing the configuration of any other system components.
sgCertifyLogSrv [**host=**<*Management Server Address* [\Domain]>]	Contacts the Management Server and creates a new certificate for the Log Server to allow secure communications with other system components. Renewing an existing certificate does not require changing the configuration of any other system components. **host** specifies the address of the Management Server. If the parameter is not defined, the loopback address (localhost) is used. **Domain** specifies the administrative Domain the Log Server belongs to if the system is divided in administrative Domains. If the Domain is not specified, the Shared Domain is used. The Log Server needs to be shut down before running this command. Restart the server after running this command.
sgCertifyMgtSrv	Creates a new certificate for the Management Server to allow secure communications between the Stonesoft system components. Renewing an existing certificate does not require changes on any other system components. The Management Server needs to be shut down before running this command. Restart the server after running this command
sgCertifyWebPortalSrv [**host=**<*Management Server Address* [\Domain]>]	Contacts the Management Server and creates a new certificate for the Web Portal Server to allow secure communications with other system components. Renewing an existing certificate does not require changing the configuration of any other system components. **host** specifies the address of the Management Server. If the parameter is not defined, the loopback address (localhost) is used. **Domain** specifies the administrative Domain the Web Portal Server belongs to if the system is divided in administrative Domains. If the Domain is not specified, the Shared Domain is used. The Web Portal Server needs to be shut down before running this command. Restart the server after running this command.
sgChangeMgtIPOnAuthSrv <*IP address*>	Changes the Management Server's IP address in the Authentication Server's local configuration to the IP address you give as a parameter. Use this command if you change the Management Server's IP address. Restart the Authentication Server after running this command.

Table A.1 Management Center Command Line Tools (Continued)

Command	Description
sgChangeMgtIPOnLogSrv *<IP address>*	Changes the Management Server's IP address in the Log Server's local configuration to the IP address you give as a parameter. Use this command if you change the Management Server's IP address. Restart the Log Server service after running this command.
sgChangeMgtIPOnMgtSrv *<IP address>*	Changes the Management Server's IP address in the local configuration to the IP address you give as a parameter. Use this command if you change the Management Server's IP address. Restart the Management Server service after running this command.
sgClient	Starts a locally installed Stonesoft Management Client.
sgCreateAdmin	Creates an unrestricted (superuser) administrator account. The Management Server needs to be stopped before running this command.
sgExport [**host=**<*Management Server Address* [\Domain]>] [**login**=<*login name*>] [**pass**=<*password*>] **file=**<*file path and name*> [**type**=<*all*|*nw*|*ips*|*sv*|*rb*|*al*> [**name**= <*element name 1, element name 2, ...*>] [**recursion**] [**-system**] [**-h** | **-help** | -?]	Exports elements stored on the Management Server to an XML file. Enclose details in double quotes if they contain spaces. **host** specifies the address of the Management Server. If the parameter is not defined, the loopback address (localhost) is used. **Domain** specifies the administrative Domain for this operation if the system is divided in administrative Domains. If the Domain is not specified, the Shared Domain is used. **login** defines the username for the account that is used for this operation. If this parameter is not defined, the username `root` is used. **pass** defines the password for the user account. **file** defines the name and location of the export ZIP file. **type** specifies which types of elements are included in the export file: `all` for all exportable elements `nw` for network elements `ips` for IPS elements `sv` for services `rb` for security policies `al` for alerts `vpn` for VPN elements. `name` allows you to specify by name the element(s) that you want to export. **recursion** includes referenced elements in the export, for example, the network elements used in a policy that you export. **-system** includes any system elements that are referenced by the other elements in the export. **-h**, **-help**, or **-?** displays information on using the script.

Table A.1 Management Center Command Line Tools (Continued)

Command	Description
sgHA [**host=**<*Management Server Address* [\Domain]>] [**login**=<*login name*>] [**pass**=<*password*>] [**master**=<*management server*>] [**-set-active**] [**-set-standby**] [**-sync**] [**-fullsync**] [**-check**] [**-retry**] [**-isolate**] [**-force**] [**-h**|**-help**|**-?**]	Controls highly available (active and standby) Management Servers. **host** specifies the address of the Management Server. If the parameter is not defined, the loopback address (localhost) is used. **Domain** specifies the administrative Domain for this operation if the system is divided in administrative Domains. If the Domain is not specified, the Shared Domain is used. **login** defines the username for the account that is used for this operation. If this parameter is not defined, the username `root` is used. **pass** defines the password for the user account. **master** defines the Management Server used as a master server for the operation. **-set-active** activates and locks all domains. **-set-standby** deactivates and unlocks all domains. **-sync** performs full database replication. This functions differently on a standby Management Server and an active Management Server. If you run it on an active Management Server, it replicates the active database to every standby Management Server that has the **Include in Database Replication** option selected in its properties. If you run it on a standby Management Server, it replicates the active database from the active Management Server only to this standby Management Server (regardless of whether the **Include in Database Replication** option is selected in the standby Management Server's properties). **-fullsync** performs full database replication with active server backup. **-check** checks that the database is in sync with the active server. **-retry** retries replication if this has been stopped due to a recoverable error.m **-isolate** isolates the management server from database replication. This is an initial requirement for synchronization. **-force** enforces the operation even if all Management Servers are not in sync. Note that using this option may cause instability if used carelessly. **-h**, **-help**, or **-?** displays information on using the script.

Table A.1 Management Center Command Line Tools (Continued)

Command	Description
sgImport [**host=**<*Management Server Address* [\Domain]>] [**login**=<*login name*>] [**pass**=<*password*>] **file=**<*file path and name*> [**-replace_all**] [**-h**|**-help**|-?]	Imports Stonesoft Management Server database elements from a Stonesoft XML file. When importing, existing (non-default) elements are overwritten if both the name and type match. **host** specifies the address of the Management Server. If the parameter is not defined, the loopback address (localhost) is used. **Domain** specifies the administrative Domain for this operation if the system is divided in administrative Domains. If the Domain is not specified, the Shared Domain is used. **login** defines the username for the account that is used for this operation. If this parameter is not defined, the username root is used. **pass** defines the password for the user account. **file** defines the ZIP file whose contents you want to import. **-replace_all** ignores all conflicts by replacing all existing elements with new ones. **-h**, **-help**, or **-?** displays information on using the script.
sgImportExportUser [**host=**<*Management Server Address* [\Domain]>] [**login**=<*login name*>] [**pass**=<*password*>] **action=**<*import|export*> **file=**<*file path and name*> [**-h**|**-help**|**-?**]	Imports and exports a list of Users and User Groups in an LDIF file from/to a Stonesoft Management Server's internal LDAP database. To import User Groups, all User Groups in the LDIF file must be directly under the **stonesoft** top-level group (dc=stonesoft). **The user information in the export file is stored as plaintext. Handle the file securely.** **host** specifies the address of the Management Server. If the parameter is not defined, the loopback address (localhost) is used. **Domain** specifies the administrative Domain for this operation if the system is divided in administrative Domains. If the Domain is not specified, the Shared Domain is used. **login** defines the username for the account that is used for this operation. If this parameter is not defined, the username root is used. **pass** defines the password for the user account. **action** defines whether users are imported or exported. **file** defines the file that is used for the operation. **Example**: sgImportExportUser login=admin pass=abc123 action=export file=c:\temp\exportedusers.ldif **-h**, **-help**, or **-?** displays information on using the script.

Table A.1 Management Center Command Line Tools (Continued)

Command	Description
sgInfo **SG_ROOT_DIR** **FILENAME** [**fast**] [**-nolog**] [**-client**] [**-h**|**-help**|-?]	Creates a ZIP file that contains copies of configuration files and the system trace files. The resulting ZIP file is stored in the logged in user's home directory. The file location is displayed on the last line of screen output. Provide the generated file to Stonesoft support for troubleshooting purposes. **SG_ROOT_DIR** Stonesoft Management Center installation directory. **FILENAME** name of output file. **-nolog** extended log server information is NOT collected. **-client** collects traces only from the Management Client. **-h**, **-help**, or **-?** displays information on using the script.
sgOnlineReplication [**login**=<*login name*>] [**pass**=<*password*>] [**active-server=**<*name of active server*>] [**standby-server=**<*name of standby server*>] [**standby-server-address=**<*IP address of standby server*>] [**-nodisplay**] [**-h**|**-help**|**-?**]	Replicates the Management Server's database from the active Management Server to the standby Management Server. **Note!** Use this script only if the secondary Management Server's configuration has been corrupted, the secondary Management Server's certificate has expired, or in new SMC installations if the automatic database replication between the Management Servers has not succeeded. Otherwise, synchronize the database through the Management Client (see Synchronizing Management Databases Manually (page 339) or use the **sgHA** command. **pass** defines the password for the user account. **active-server** option specifies the IP address of the active Management Server from which the Management database is replicated. **standby-server** option specifies the name of the standby Management Server to which the Management database is replicated. **standby-server-address** option specifies the IP address of the standby Management Server to which the Management database is replicated. **-nodisplay** sets a text only console. **-h**, **-help**, or **-?** displays information on using the script. The return values are: **0** OK **8** sgOnlineReplication.sh failed to initialize properly **9** login failed **11** unknown error **12** bad command line arguments **13** replication canceled by user.
sgReinitializeLogServer	**Note!** This script is located in `<installation directory>/bin/install`. Creates a new Log Server configuration if the configuration file has been lost.

Table A.1 Management Center Command Line Tools (Continued)

Command	Description
sgRestoreArchive <ARCHIVE_DIR>	Restores logs from archive files to the Log Server. This command is available only on the Log Server. **ARCHIVE_DIR** is the number of the archive directory (0 – 31) from where the logs will be restored. By default, only archive directory 0 is defined. The archive directories can be defined in the <*installation directory*>/data/LogServerConfiguration.txt file: ARCHIVE_DIR_*xx*=*PATH*.
sgRestoreAuthBackup [**-pwd=**<*password*>] [**-backup=**<*backup file name*>] [**-nodiskcheck**] [**-h**|**-help**]	Restores the Authentication Server user information from a backup file in the <*installation directory*>/backups/ directory. Apply the Authentication Server's configuration after this command. **-pwd** defines a password for encrypted backup. **-backup** defines a name for the backup file. **-nodiskcheck** ignores free disk check before backup restoration. **-h** or **-help** displays information on using the script.
sgRestoreLogBackup [**-pwd=**<*password*>] [**-backup=**<*backup file name*>] [**-nodiskcheck**] [**-overwrite-syslog-template**] [**-h**|**-help**]	Restores the Log Server (logs and/or configuration files) from a backup file in the <*installation directory*>/backups/ directory. Apply the Authentication Server's configuration after this command. **-pwd** defines a password for encrypted backup. **-backup** defines a name for the backup file. **-nodiskcheck** ignores free disk check before backup restoration. **-overwrite-syslog-template** overwrites a syslog template file if found in the backup. **-h** or **-help** displays information on using the script.
sgRestoreMgtBackup [**-pwd**=<*password*>] [**-backup**=<*backup file name*>] [**-nodiskcheck**] [**-h**|**-help**]	Restores the Management Server (database and/or configuration files) from a backup file in the <*installation directory*>/backups/ directory. **-pwd** defines a password for encrypted backup. **-backup** defines a name for the backup file. **-nodiskcheck** ignores free disk check before backup restoration. **-h** or **-help** displays information on using the script.
sgRevert	**Note!** This script is located in <installation directory>/bin/uninstall. Reverts to the previous installation saved during the upgrade process. The previous installation can be restored at any time, even after a successful upgrade.

Table A.1 Management Center Command Line Tools (Continued)

Command	Description
sgShowFingerPrint	Displays the CA certificate’s fingerprint on the Management Server.
sgStartAuthSrv	Starts the Authentication Server.
sgStartLogSrv	Starts the Log Server and its database.
sgStartMgtDatabase	Starts the Management Server’s database. There is usually no need to use this script.
sgStartMgtSrv	Starts the Management Server and its database.
sgStartWebPortalSrv	Starts the Web Portal Server.
sgStopLogSrv	Stops the Log Server.
sgStopMgtSrv	Stops the Management Server and its database.
sgStopMgtDatabase	Stops the Management Server’s database. There is usually no need to use this script.
sgStopWebPortalSrv	Stops the Web Portal Server.
sgStopRemoteMgtSrv [**host=<***Management Server Host Name***>**] [**login=<***login name*>] [**pass=<***password*>] [**-h**|**-help**|**-?**]	Stops the Management Server service when run without arguments. To stop a remote Management Server service, provide the arguments to connect to the Management Server. **host** is the Management Server’s host name if not localhost. **login** is a Stonesoft administrator account for the login. **pass** is the password for the administrator account. **-h**, **-help**, or **-?** displays information on using the script.

Table A.1 Management Center Command Line Tools (Continued)

Command	Description
sgTextBrowser [**host=**<*Management Server address* [\Domain]>] [**login**=<*login name*>] [**pass**=<*password*>] [**format**=<*CSV*|*XML*>] [**o=**<*output file*>] [**f=**<*filter file*>] [**e=**<*filter expression*>] [**m=**<*current*|*stored*>] [**limit=**<*maximum number of unique records to fetch*>] [**-h**|**-help**|**-?**]	Displays or exports current or stored logs. This command is available on the Log Server. Enclose the file and filter names in double quotes if they contain spaces. **host** defines the address of the Management Server used for checking the login information. If this parameter is not defined, Management Server is expected to be on the same host where the script is run. If Domains are in use, you can specify the Domain the Log Server belongs to. If domain is not specified, the Shared Domain is used. **login** defines the username for the account that is used for this export. If this parameter is not defined, the username root is used. **pass** defines the password for the user account used for this operation. **format** defines the file format for the output file. If this parameter is not defined, the XML format is used. **o** defines the destination output file where the logs will be exported. If this parameter is not defined, the output is displayed on screen. **f** defines the Stonesoft exported filter file that you want to use for filtering the log data. **e** defines the filter that you want to use for filtering the log data. Type the name as shown in the Management Client. **m** defines whether you want to view or export logs as they arrive on the Log Server (current) or logs stored in the active storage directory (stored). If this option is not defined, the current logs are used. **limit** defines the maximum number of unique records to be fetched. The default value is unlimited. **-h**, **-help**, or **-?** displays information on using the script.

Engine Commands

The commands in the following two tables can be run on the command line on Firewall, Virtual Firewall, Layer 2 Firewall, and/or IPS engines.

Note – All command line tools that are available in the Firewall role are also available for Virtual Firewalls. However, there is no direct access to the command line of Virtual Firewalls. Commands to Virtual Firewalls must be sent from the command line of the Master Engine using the `se-virtual-engine` command.

Table A.2 Stonesoft Engine Command Line Tools

Command	Engine Role	Description
`se-virtual-engine` **`-l	--list`** **`-v`** *`<virtual engine ID>`* **`-e	--enter`** **`-E “`***`<command [options]>`***`”`** **`-h	--help`**	Firewall (*Master Engine only)*	Used to send commands to Virtual Firewalls from the command line of the Master Engine. All commands that can be used for the Firewall role can also be used for Virtual Firewalls. **`-l`** or **`--list`** list the active Virtual Security Engines. **`- v <virtual engine ID>`** specifies the ID of the Virtual Security Engine on which to execute the command. **`-e`** or **`--enter`** enters the command shell for the Virtual Security Engine specified with the `-v` option. To exit the command shell, type `exit`. **`-E “`***`<command [options]>`***`”`** executes the specified command on the Virtual Security Engine specified with the `-v` option. **`-h`** or **`--help`** shows the help message for the se-virtual-engine command.

Table A.2 Stonesoft Engine Command Line Tools (Continued)

Command	Engine Role	Description
sg-blacklist **show** [-v] [-f *FILENAME*] | **add [** [**-i** *FILENAME*] | [**src** *IP_ADDRESS/MASK*] [src6 *IPv6_ADDRESS/PREFIX*] [dst *IP_ADDRESS/MASK*] [dst6 *IPv6_ADDRESS/PREFIX*] [**proto** *{tcp|udp|icmp|NUM}*] [**srcport** *PORT{-PORT}*] [**dstport** *PORT{-PORT}*] [**duration** *NUM*] **]** | **del [** [**-i** *FILENAME*] | [**src** *IP_ADDRESS/MASK*] [src6 *IPv6_ADDRESS/PREFIX*] [dst *IP_ADDRESS/MASK*] [dst6 *IPv6_ADDRESS/PREFIX*] [**proto** *{tcp|udp|icmp|NUM}*] [**srcport** *PORT{-PORT}*] [**dstport** *PORT{-PORT}*] [**duration** *NUM*] **]** | **iddel** *NODE_ID ID* | **flush**	Firewall, Layer 2 Firewall, IPS	Used to view, add, or delete active blacklist entries. The blacklist is applied as defined in Access Rules. **Commands:** **show** displays the current active blacklist entries in format: engine node ID | blacklist entry ID | (internal) | entry creation time | (internal) | address and port match | originally set duration | (internal) | (internal). Use the -f option to specify a storage file to view (/data/blacklist/db_<number>). The **-v** option adds operation's details to the output. **add** creates a new blacklist entry. Enter the parameters (see below) or use the **-i** option to import parameters from a file. **del** deletes the first matching blacklist entry. Enter the parameters (see below) or use the **-i** option to import parameters from a file. **iddel** *NODE_ID ID* removes one specific blacklist entry on one specific engine. NODE_ID is the engine's ID, ID is the blacklist entry's ID (as shown by the show command). **flush** deletes all blacklist entries. **Add/Del Parameters:** Enter at least one parameter. The default value is used for the parameters that you omit. You can also save parameters in a text file; each line in the file is read as one blacklist entry. **src** *IP_ADDRESS/MASK* defines the source IP address and netmask to match. Matches any IP address by default. **src6** *IPv6_ADDRESS/PREFIX* defines the source IPv6 and prefix length to match. Matches any IPv6 address by default. **dst** *IP_ADDRESS/MASK* defines the destination IP address and netmask to match. Matches any IP address by default. **dst6** *IPv6_ADDRESS/PREFIX* defines the destination IPv6 address and prefix length to match. Matches any IPv6 address by default. **proto** *{tcp|udp|icmp|NUM}* defines the protocol to match by name or protocol number. Matches all IP traffic by default. **srcport** *PORT[-PORT]* defines the TCP/UDP source port or range to match. Matches any port by default. **dstport** *PORT[-PORT]* defines the TCP/UDP destination port or range to match. Matches any port by default. **duration** *NUM* defines in seconds how long the entry is kept. Default is 0, which cuts current connections, but is not kept. **Examples:** sg-blacklist add src 192.168.0.2/32 proto tcp dstport 80 duration 60 sg-blacklist add -i myblacklist.txt sg-blacklist del dst 192.168.1.0/24 proto 47

Table A.2 Stonesoft Engine Command Line Tools (Continued)

Command	Engine Role	Description
sg-bootconfig [**--primary-console**=*tty0|ttyS PORT,SPEED*] [**--secondary-console**=[*tty0|ttyS PORT,SPEED*]] [**--flavor**=*up|smp*] [**--initrd**=*yes|no*] [**--crashdump**=*yes|no|Y@X*] [**--append**=*kernel options*] [**--help**] **apply**	Firewall, Layer 2 Firewall, IPS	Used to edit boot command parameters for future bootups. **--primary-console**=*tty0|ttyS PORT,SPEED* parameter defines the terminal settings for the primary console. **--secondary-console**= [*tty0|ttyS PORT,SPEED*] parameter defines the terminal settings for the secondary console. **--flavor**=*up|smp* [*-kdb*] parameter defines whether the kernel is uniprocessor or multiprocessor. **--initrd**=*yes|no* parameter defines whether Ramdisk is enabled or disabled. **--crashdump**=*yes|no|Y@X* parameter defines whether kernel crashdump is enabled or disabled, and how much memory is allocated to the crash dump kernel (Y). The default is 24M. X must always be 16M. **--append**=*kernel options* parameter defines any other boot options to add to the configuration. **--help** parameter displays usage information. **apply** command applies the specified configuration options.
sg-clear-all	Firewall, Layer 2 Firewall, IPS	**Note! Use this only if you want to clear all configuration information from the engine.** This command resets all configuration information from the engine. It does not remove the engine software. After using this command, you must reconfigure the engine using the sg-reconfigure command.
sg-cluster **[-v *<virtual engine ID>*]** **[status [-c *SECONDS*]]** **[versions]** **[online]** **[lock-online]** **[offline]** **[lock-offline]** **[standby]** **[safe-offline]** **[force-offline]**	Firewall, Layer 2 Firewall, IPS	Used to display or change the status of the node. **-v *<virtual engine ID>*** (*Master Engine only*) option specifies the ID of the Virtual Security Engine on which to execute the command. **status** [**-c** *SECONDS*] command displays cluster status. When **-c** *SECONDS* is used, status is shown continuously with the specified number of seconds between updates. **version** command displays the engine software versions of the nodes in the cluster. **online** command sends the node online. **lock-online** command sends the node online and keeps it online even if another process tries to change its state. **offline** command sends the node offline. **lock-offline** command sends the node offline and keeps it offline even if another process tries to change its state. **standby** command sets an active node to standby. **safe-offline** command sets the node to offline only if there is another online node. **force-offline** command sets the node online regardless of state or any limitations. Also sets all other nodes offline.

Table A.2 Stonesoft Engine Command Line Tools (Continued)

Command	Engine Role	Description
sg-contact-mgmt	Firewall, Layer 2 Firewall, IPS	Used for establishing a trust relationship with the Management Server as part of engine installation or reconfiguration (see `sg-reconfigure` below). The engine contacts the Management Server using the one-time password created when the engine's initial configuration is saved.
sg-dynamic-routing **[start]** **[stop]** **[restart]** **[force-reload]** **[backup** *<file>*] **[restore** *<file>*] **[sample-config]** **[route-table]** **[info]**	Firewall	`start` starts the Quagga routing suite. `stop` stops the Quagga routing suite and flushes all routes made by zebra. `restart` restarts the Quagga routing suite. `force-reload` forces reload of the saved configuration. `backup` *<file>* backs up the current configuration to a compressed file. `restore` *<file>* restores the configuration from the specified file. `sample-config` creates a basic configuration for Quagga. `route-table` prints the current routing table. `info` displays the help information for the sg-dynamic-routing command, and detailed information about Quagga suite configuration with vtysh.
sg-ipsec -d [**-u** *<username[@domain]>* | **-si** *<session id>* | **-ck** *<ike cookie>* | **-tri** *<transform id>* **-ri** *<remote ip>* | **-ci** *<connection id>*]	Firewall	Deletes VPN-related information (use `vpninfo` command to view the information). Option **-d** (for delete) is mandatory. **-u** deletes the VPN session of the named VPN client user. You can enter the user account in the form <username@domain> if there are several user storage locations (LDAP domains). **-si** deletes the VPN session of a VPN client user based on session identifier. **-ck** deletes the IKE SA (Phase one security association) based on IKE cookie. **-tri** deletes the IPSEC SAs (Phase two security associations) for both communication directions based on transform identifier. **-ri** deletes all SAs related to a remote IP address in gateway-to-gateway VPNs. **-ci** deletes all SAs related to a connection identifier in gateway-to-gateway VPNs.

Table A.2 Stonesoft Engine Command Line Tools (Continued)

Command	Engine Role	Description
sg-logger **-f** *FACILITY_NUMBER* **-t** *TYPE_NUMBER* [**-e** *EVENT_NUMBER*] [**-i** *"INFO_STRING"*] [**-s**] [**-h**]	Firewall, Layer 2 Firewall, IPS	Used in scripts to create log messages with the specified properties. **-f** *FACILITY_NUMBER* parameter defines the facility for the log message. **-t** *TYPE_NUMBER* parameter defines the type for the log message. **-e** *EVENT_NUMBER* parameter defines the log event for the log message. The default is 0 (H2A_LOG_EVENT_UNDEFINED). **-i** *"INFO_STRING"* parameter defines the information string for the log message. **-s** parameter dumps information on option numbers to stdout **-h** parameter displays usage information.
sg-raid [**-status**] [**-add**] [**-re-add**] [**-force**] [**-help**]	Firewall, Layer 2 Firewall, IPS	Configures a new hard drive. This command is only for Stonesoft appliances that support RAID (Redundant Array of Independent Disks) and have two hard drives. **-status** option displays the status of the hard drive. **-add** options adds a new empty hard drive. Use **-add -force** if you want to add a hard drive that already contains data and you want to overwrite it. **-re-add** adds a hard drive that is already partitioned. This command prompts for the drive and partition for each degraded array. Use **-re-add -force** if you want to check all the arrays. **-help** option option displays usage information.
sg-reconfigure [**--boot**] [**--maybe-contact**] [**--no-shutdown**]	Firewall, Layer 2 Firewall, IPS	Used for reconfiguring the node manually. **--boot** option applies bootup behavior. Do not use this option unless you have a specific need to do so. **--maybe-contact** option contacts the Management Server if requested. This option is only available on firewall engines. **--no-shutdown** option allows you to make limited configuration changes on the node without shutting it down. Some changes may not be applied until the node is rebooted.
sg-selftest [**-d**] [**-h**]	Firewall	Runs cryptography tests on the engine. **-d** option runs the tests in debug mode. **-h** option displays usage information.
sg-status [**-l**] [**-h**]	Firewall, Layer 2 Firewall, IPS	Displays information on the engine's status. **-l** option displays all available information on engine status. **-h** option displays usage information.

Table A.2 Stonesoft Engine Command Line Tools (Continued)

Command	Engine Role	Description
sg-toggle-active *SHA1 SIZE* | **--force** [**--debug**]	Firewall, Layer 2 Firewall, IPS	Switches the engine between the active and the inactive partition. This change takes effect when you reboot the engine. You can use this command, for example, if you have upgraded an engine and want to switch back to the earlier engine version. When you upgrade the engine, the active partition is switched. The earlier configuration remains on the inactive partition. To see the currently active (and inactive) partition, see the directory listing of `/var/run/stonegate` (**`ls -l /var/run/stonegate)`**. The *SHA1 SIZE* option is used to verify the signature of the inactive partition before changing it to active. If you downgrade the engine, check the checksum and the size of the earlier upgrade package by extracting the signature and size files from the sg_engine_[version.build]_i386.zip file. **--debug** option reboots the engine with the debug kernel. **--force** option switches the active configuration without first verifying the signature of the inactive partition.
sg-upgrade	Firewall	Upgrades the node by rebooting from the installation DVD. Alternatively, the node can be upgraded remotely using the Management Client.
sg-version	Firewall, Layer 2 Firewall, IPS	Displays the software version and build number for the node.
sginfo [**-f**] [**-d**] [**-s**] [**-p**] [**--**] [**--help**]	Firewall, Layer 2 Firewall, IPS	Gathers system information you can send to Stonesoft support if you are having problems. Use this command only when instructed to do so by Stonesoft support. **-f** option forces sgInfo even if the configuration is encrypted. **-d** option includes core dumps in the sgInfo file. **-s** option includes slapcat output in the sgInfo file. **-p** option includes passwords in the sgInfo file (by default passwords are erased from the output). **--** option creates the sgInfo file without displaying the progress **--help** option displays usage information.

The table below lists some general Linux operating system commands that may be useful in running your engines. Some commands can be stopped by pressing `Ctrl+c`.

Table A.3 General Command Line Tools on Engines

Command	Description
dmesg	Shows system logs and other information. Use the -h option to see usage.
halt	Shuts down the system.
ip	Displays IP address information. Type the command without options to see usage. **Example:** type **`ip addr`** for basic information on all interfaces.
ping	Tests connectivity with ICMP echo requests. Type the command without options to see usage.
ps	Reports the status of running processes.
reboot	Reboots the system.
scp	Secure copy. Type the command without options to see usage.
sftp	Secure FTP. Type the command without options to see usage.
ssh	SSH client (for opening a terminal connection to other hosts). Type the command without options to see usage.
tcpdump	Gives information on network traffic. Use the **-h** option to see usage. You can also analyze network traffic by creating tcpdump files from the Stonesoft Management Client GUI with the Traffic Capture feature. See Taking a Traffic Capture (page 126).
top	Displays the top CPU processes taking most processor time. Use the **-h** option to see usage.
traceroute	Traces the route packets take to the specified destination. Type the command without options to see usage.
vpninfo	Displays VPN information and allows you to issue some basic commands. Type the command without options to see usage.

Server Pool Monitoring Agent Commands

You can test and monitor the Server Pool Monitoring Agents on the command line with the commands described in the table below.

Table A.4 Server Pool Monitoring Agent Commands

Command	Description
agent [-v *level*] [-c *path*] [test [*files*]] [syntax [*files*]]	(*Windows only*) Allows you to test different configurations before activating them. -v *level* Set the verbosity level. The default level is 5. Levels 6-8 are for debugging where available. -c *path* Use the specified path as the first search directory for the configuration. test [*files*] Run in the test mode - status queries do not receive a response. If you specify the files, they are used for reading the configuration instead of the default files. The output is directed to syslog or eventlog instead of the console where the command was run unless you use the -d option. syntax [*files*] Check the syntax in the configuration file. If no files are specified, the default configuration files are checked.
sgagentd [-d] [-v *level*] [-c *path*] [test [*files*]] [syntax [*files*]]	(*Linux only*) Allows you to test different configurations before activating them. -d Don't Fork as a daemon. All log messages are printed to stdout or stderr only. -v *level* Set the verbosity level. The default level is 5. Levels 6-8 are for debugging where available. -c *path* Use the specified path as the first search directory for the configuration. test [*files*] Run in the test mode - status queries do not receive a response. If you specify the files, they are used for reading the configuration instead of the default files. The output is directed to syslog or eventlog instead of the console where the command was run unless you use the -d option. syntax [*files*] Check the syntax in the configuration file. If no files are specified, the default configuration files are checked. The output is directed to syslog or eventlog instead of the console where the command was run unless you use the -d option.

Table A.4 Server Pool Monitoring Agent Commands (Continued)

Command	Description
sgmon [*status*|*info*|*proto*] [-p *port*] [-t *timeout*] [-a *id*] *host*	Sends a UDP query to the specified host and waits for a response until received, or until the timeout limit is reached. The request type can be defined as a parameter. If no parameter is given, `status` is requested. The commands are: `status` - query the status. `info` - query the agent version. `proto` - query the highest supported protocol version. `-p` *port* Connect to the specified port instead of the default port. `-t` *timeout* Set the timeout (in seconds) to wait for a response. `-a` *id* Acknowledge the received log messages up to the specified id. Each response message has an id, and you may acknowledge more than one message at a given time by using the id parameter. Note that messages acknowledged by sgmon will no longer appear in the firewall logs. *host* The IP address of the host to connect to. To get the status locally, you may give `localhost` as the host argument. This parameter is mandatory. **Return value**: 0 if the response was received 1 if the query timed out -1 in case of an error

APPENDIX B

DEFAULT COMMUNICATION PORTS

This chapter lists the default ports used in connections between Stonesoft components and the default ports Stonesoft components use with external components.

The following sections are included:

Management Center Ports

The illustrations below present an overview to the most important default ports used in communications between the Stonesoft Management Center (SMC) components and from the SMC to external services. See the table below for a complete list of default ports.

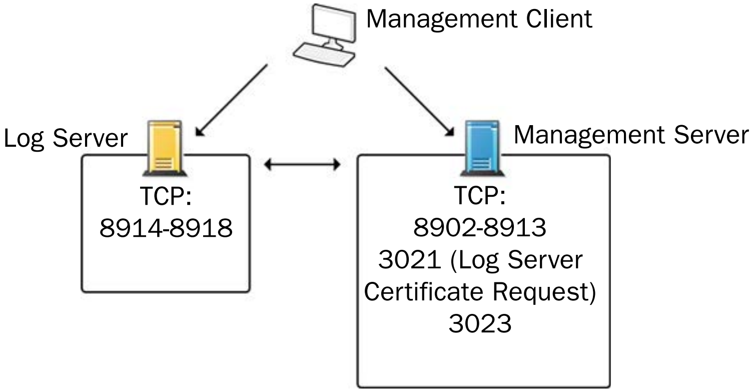

Illustration B.1 Destination Ports for Basic Communications Within SMC

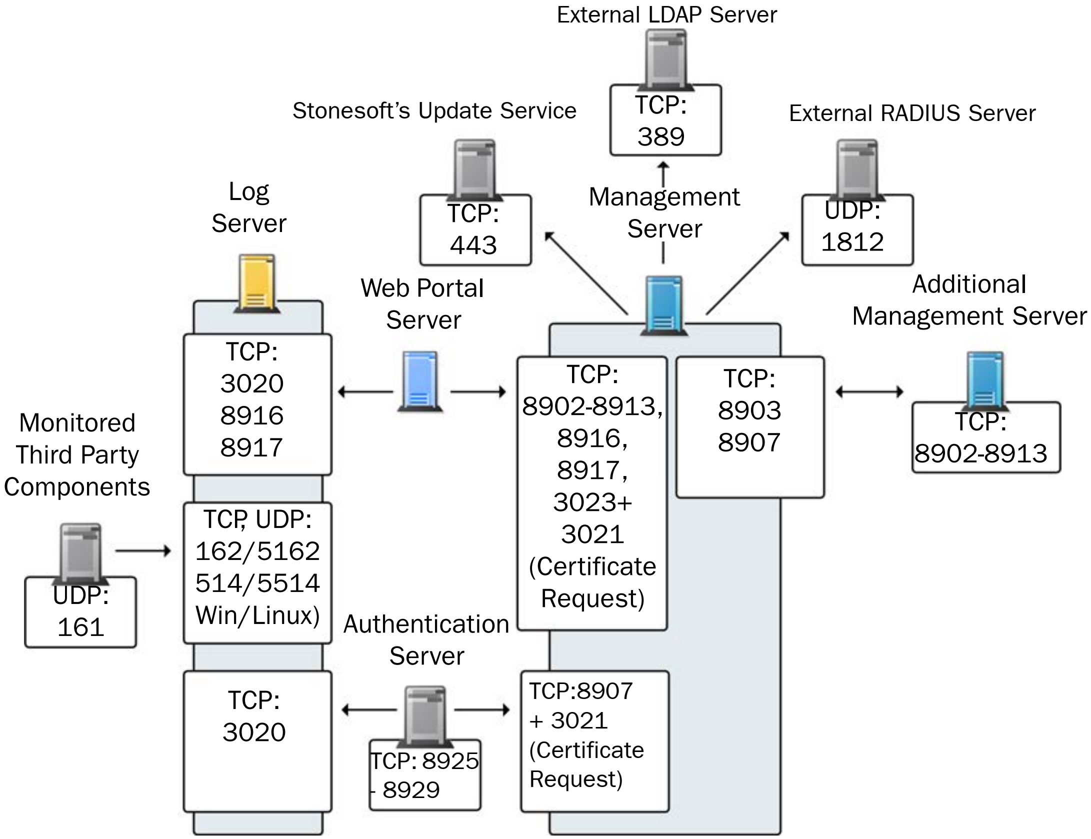

Illustration B.2 Default Destination Ports for Optional SMC Components and Features

The table below lists all default ports SMC uses internally and with external components. Many of these ports can be changed. The name of corresponding default Service elements are also included for your reference. For information on communications between SMC components and the engines, see the separate listings.

Table B.1 Management Center Default Ports

Listening Host	Port/ Protocol	Contacting Hosts	Service Description	Service Element Name
Additional Management Servers	8902-8913/TCP	Management Server	Database replication (push) to the additional Management Server.	SG Control
Authentication Server	8925-8929/TCP	Management Server	Stonesoft Management Server commands to Authentication Server.	SG Authentication Commands
Authentication Server node	8988-8989/TCP	Authentication Server node	Data synchronization between Authentication Server nodes.	SG Authentication Sync
DNS server	53/UDP, 53/TCP	Management Client, Management Server, Log Server	DNS queries.	DNS (UDP)
LDAP server	389/TCP	Management Server	External LDAP queries for display/ editing in the Management Client.	LDAP (TCP)
Log Server	162/UDP, 5162/UDP	Monitored third party components	SNMPv1 trap reception from third party components. Port 162 is used if installed on Windows, port 5162 if installed on Linux.	SNMP (UDP)
Log Server	514/TCP, 514/UDP, 5514/TCP, 5514/UDP	Monitored third party components	Syslog reception from third party components. Port 514 is used if installed on Windows, port 5514 if installed on Linux.	Syslog (UDP) [*Partial match*]
Log Server	3020/TCP	Authentication Server, Log Server, Web Portal Server, Security Engines	Alert sending from the Authentication Server, Log Server, and Web Portal Server. Log and alert messages; monitoring of blacklists, connections, status, and statistics from Security Engines.	SG Log
Log Server	8914-8918/TCP	Management Client	Log browsing.	SG Data Browsing
Log Server	8916-8917/TCP	Web Portal Server	Log browsing.	SG Data Browsing (Web Portal Server)
Management Server	3021/TCP	Log Server, Web Portal Server	System communications certificate request/renewal.	SG Log Initial Contact

Table B.1 Management Center Default Ports (Continued)

Listening Host	Port/ Protocol	Contacting Hosts	Service Description	Service Element Name
Management Server	8902-8913/TCP	Management Client, Log Server, Web Portal Server	Monitoring and control connections.	SG Control
Management Server	3023/TCP	Management Server, Log Server, Web Portal Server	Log Server and Web Portal Server status monitoring. Status information sent from a Management Server in one Domain to a Management Server in another Domain.	SG Status Monitoring
Management Server	8903, 8907/TCP	Additional Management Servers	Database replication (pull) to the additional Management Server.	SG Control
Management Server	8907/TCP	Authentication Server	Status monitoring.	SG Control
Monitored Third Party Components	161/UDP	Log Server	SNMP status probing to external IP addresses.	SNMP (UDP)
RADIUS server	1812/UDP	Management Server	RADIUS authentication requests for administrator logins. The default ports can be modified in the properties of the RADIUS Server element.	RADIUS (Authentication)
Stonesoft servers	443/TCP	Management Server	Update packages, engine upgrades, and licenses from update.stonesoft.com and smc.stonesoft.com.	HTTPS
Syslog Server	514/UDP, 5514/UDP	Log Server	Log data export to syslog servers. The default ports can be modified in the LogServerConfiguration.txt file.	Syslog (UDP) [*Partial match*]

Security Engine Ports

The illustrations below present an overview to the most important default ports used in communications between Security Engines and the SMC and between clustered Security Engine nodes. Master Engines use the same default ports as clustered Security Engines. Virtual Security Engines do not communicate directly with other system components. See the table below for a complete list of default ports for the engines.

Illustration B.3 Destination Ports for Basic Security Engine Communications

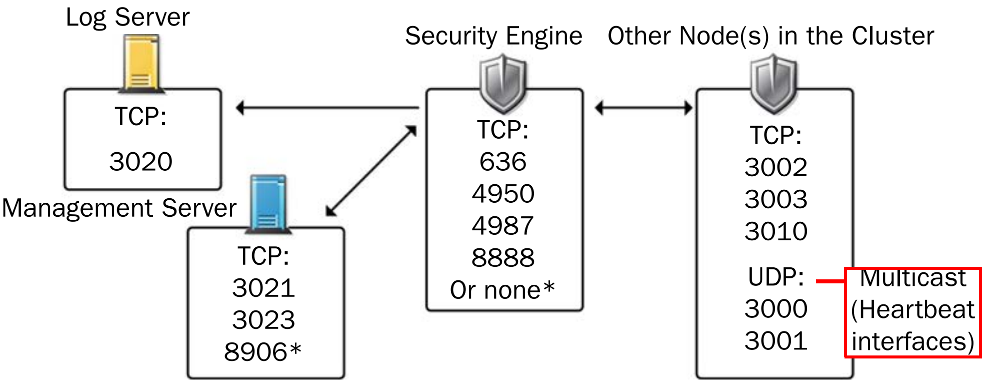

*Single engines with “Node-initiated Contact to Management Server” selected.

Illustration B.4 Default Destination Ports for Security Engine Service Communications

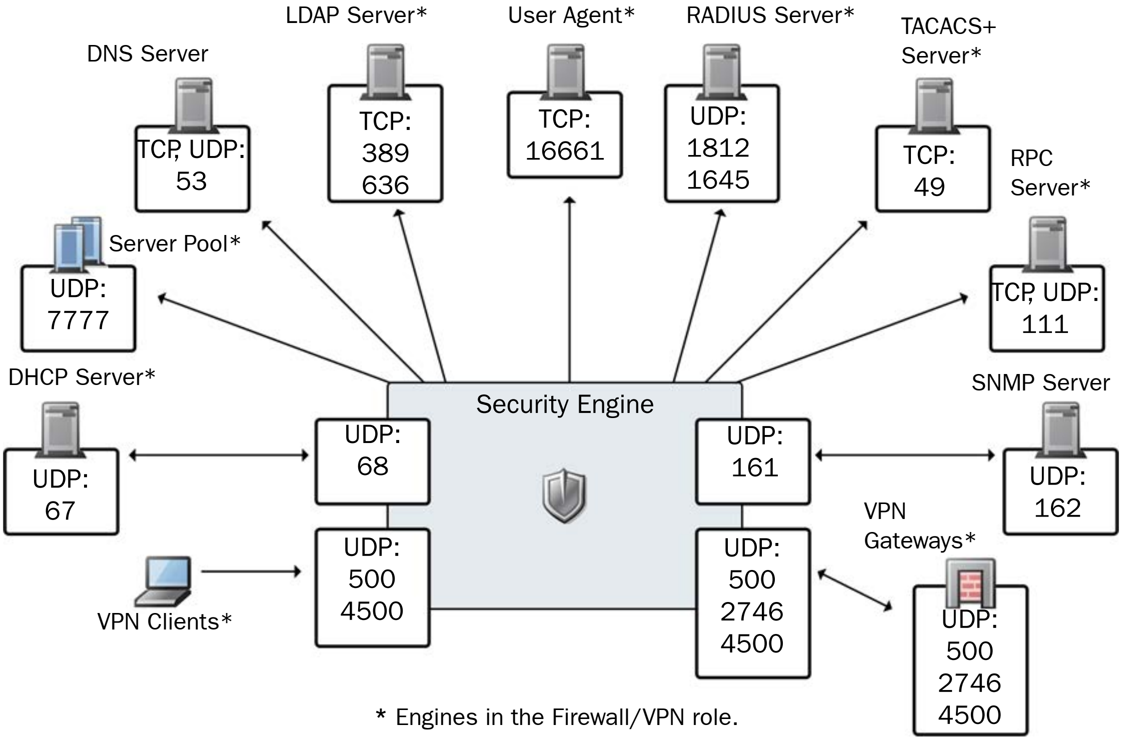

* Engines in the Firewall/VPN role.

The table below lists all default ports the Security Engines use internally and with external components. Many of these ports can be changed. The name of corresponding default Service elements are also included for your reference.

Table B.2 Security Engine Default Ports

Listening Host	Port/Protocol	Contacting Hosts	Service Description	Service Element Name
Anti-virus signature server	80/TCP	Firewall	Anti-virus signature update service.	HTTP
Authentication Server	8925-8929/ TCP	Firewall, Master Engine	User directory and authentication services.	LDAP (TCP), RADIUS (Authentication)
BrightCloud Server	2316/TCP	Firewall, Layer 2 Firewall, IPS	BrightCloud web filtering update service.	BrightCloud update
DHCP server	67/UDP	Firewall	Relayed DHCP requests and requests from a firewall that uses dynamic IP address.	BOOTPS (UDP)
DNS server	53/UDP, 53/TCP	Firewall, Master Engine	Dynamic DNS updates.	DNS (TCP)
Firewall	67/UDP	Any	DHCP relay on firewall engine.	BOOTPS (UDP)
Firewall	68/UDP	DHCP server	Replies to DHCP requests.	BOOTPC (UDP)
Firewall, Master Engine	500/UDP	VPN clients, VPN gateways	VPN negotiations, VPN traffic.	ISAKMP (UDP)
Firewall, Master Engine	636/TCP	Management Server	Internal user database replication.	LDAPS (TCP)
Firewall, Master Engine	2543/TCP	Any	User authentication (Telnet) for Access rules.	SG User Authentication
Firewall	2746/UDP	Stonesoft VPN gateways	UDP encapsulated VPN traffic (engine versions 5.1 and lower).	SG UDP Encapsulation
Firewall, Master Engine	4500/UDP	VPN client, VPN gateways	VPN traffic using NAT-traversal.	NAT-T
Firewall Cluster Node, Master Engine cluster node	3000-3001/ UDP 3002-3003, 3010/TCP	Firewall Cluster Node, Master Engine cluster node	Heartbeat and state synchronization between clustered Firewalls.	SG State Sync (Multicast), SG State Sync (Unicast), SG Data Sync
Firewall, Layer 2 Firewall, IPS, Master Engine	4950/TCP	Management Server	Remote upgrade.	SG Remote Upgrade

Table B.2 Security Engine Default Ports (Continued)

Listening Host	Port/Protocol	Contacting Hosts	Service Description	Service Element Name
Firewall, Layer 2 Firewall, IPS, Master Engine	4987/TCP	Management Server	Management Server commands and policy upload.	SG Commands
Firewall, Layer 2 Firewall, IPS	8888/TCP	Management Server	Connection monitoring for engine versions 5.1 and lower.	SG Legacy Monitoring
Firewall, Layer 2 Firewall, IPS, Master Engine	15000/TCP	Management Server, Log Server	Blacklist entries.	SG Blacklisting
Firewall, Layer 2 Firewall, IPS, Master Engine	161/UDP	SNMP server	SNMP monitoring.	SNMP (UDP)
IPS Cluster Node	3000-3001/ UDP 3002-3003, 3010/TCP	IPS Cluster Node	Heartbeat and state synchronization between clustered IPS engines.	SG State Sync (Multicast), SG State Sync (Unicast), SG Data Sync
LDAP server	389/TCP	Firewall, Master Engine	External LDAP queries, including StartTLS connections.	LDAP (TCP)
Layer 2 Firewall Cluster Node	3000-3001/ UDP 3002-3003, 3010/TCP	Layer 2 Firewall Cluster Node	Heartbeat and state synchronization between clustered Layer 2 Firewalls.	SG State Sync (Multicast), SG State Sync (Unicast), SG Data Sync
Log Server	3020/TCP	Firewall, Layer 2 Firewall, IPS, Master Engine	Log and alert messages; monitoring of blacklists, connections, status, and statistics.	SG Log
Management Server	3021/TCP	Firewall, Layer 2 Firewall, IPS, Master Engine	System communications certificate request/renewal (initial contact).	SG Initial Contact
Management Server	3023/TCP	Firewall, Layer 2 Firewall, IPS, Master Engine	Monitoring (status) connection.	SG Status Monitoring
Management Server	8906/TCP	Firewall, Layer 2 Firewall, IPS, Master Engine	Management connection for single engines with "Node-Initiated Contact to Management Server" selected.	SG Dynamic Control
RADIUS server	1812, 1645/ UDP	Firewall, Master Engine	RADIUS authentication requests.	RADIUS (Authentication), RADIUS (Old)

Table B.2 Security Engine Default Ports (Continued)

Listening Host	Port/Protocol	Contacting Hosts	Service Description	Service Element Name
RPC server	111/UDP, 111/TCP	Firewall, Master Engine	RPC number resolve.	SUNRPC (UDP), Sun RPC (TCP)
Server Pool Monitoring Agents	7777/UDP	Firewall, Master Engine	Polls to the servers' Server Pool Monitoring Agents for availability and load information.	SG Server Pool Monitoring
SNMP server	162/UDP	Firewall, Layer 2 Firewall, IPS, Master Engine	SNMP traps from the engine.	SNMP Trap (UDP)
TACACS+ server	49/TCP	Firewall, Master Engine	TACACS+ authentication requests.	TACACS (TCP)
User Agent	16661/TCP	Firewall, Master Engine	Queries for matching Users and User Groups with IP addresses.	SG Engine to User Agent
VPN gateways	500/UDP, 2746/UDP (Stonesoft gateways only), or 4500 UDP.	Firewall, Master Engine	VPN traffic. Ports 2746 and 4500 may be used depending on encapsulation options.	ISAKMP (UDP)

APPENDIX C

PREDEFINED ALIASES

This appendix lists the predefined Aliases that exist in the system. The predefined Aliases are used in the default system policies. Some of them may be useful when you create your own rules.

The following sections are included:

Pre-Defined User Aliases

User Aliases are usually created by administrators, but there are also some pre-defined user aliases in the system. User Aliases are preceded with one $ character. The table below lists all the editable automatically created user Aliases. These Aliases are used in the firewalls' default DHCP Relay Sub-Policy.

Table C.1 System-defined User Aliases

Pre-Defined User Alias	Description
$ DHCP address pools	Addresses that can be allocated by DHCP server(s).
$ DHCP address pools for IPsec VPN Clients	Address pools for assigning virtual IP addresses to VPN clients.
$ DHCP servers	All DHCP servers defined for the Firewall.
$ DHCP servers for IPsec VPN Clients	The DHCP servers defined for assigning virtual IP addresses to VPN clients.

System Aliases

System Aliases are automatically created non-editable Aliases. The System Aliases are preceded with two $$ characters. The table below lists the definitions of all the System Aliases. These Aliases are used in the Firewall's default system policies.

Table C.2 System Aliases

System Alias	Description
$$ DHCP Enabled Interface Addresses	IP addresses (of CVIs on clusters) which have DHCP relay enabled.
$$ DHCP Enabled interface addresses for IPsec VPN clients	IP addresses (of NDIs on clusters) which have DHCP relay enabled for VPN Clients.
$$ DHCP Interface X.dns	IP address of the DHCP-assigned DNS server for interface number X.
$$ DHCP Interface X.gateways	IP address of the DHCP-assigned default router for interface number X.
$$ DHCP Interface X.ip	DHCP-assigned IP address for interface number X.
$$ DHCP Interface X.net	Network behind the dynamic IP interface number X.
$$ Interface ID X.ip	First IP address (CVI) of Physical Interface ID X.
$$ Interface ID X.net	Directly connected networks behind Physical Interface ID X.
$$ Local Cluster	All addresses of the cluster.
$$ Local Cluster (CVI addresses only)	All CVI addresses of the cluster.

Table C.2 System Aliases (Continued)

System Alias	Description
$$ Local Cluster (DHCP Interface Addresses)	All DHCP-assigned IP addresses of the engine.
$$ Local Cluster (NDI addresses only)	All NDI addresses of all nodes in the cluster.
$$ Local Cluster (NDI for heartbeat addresses only)	Heartbeat NDI addresses of all nodes in the cluster.
$$ Local Cluster (NDI for management addresses only)	Management NDI address(es) of all nodes in the cluster.
$$ Log Servers	IP addresses of all Log Servers.
$$ Management Servers	IP addresses of all Management Server.
$$ Valid DHCP Address Pools for IPsec VPN clients	Address pools defined for assigning virtual IP addresses to VPN clients.
$$ Valid DHCP Servers	All DHCP servers defined for the Firewall.
$$ Valid DHCP Servers for IPsec VPN clients	The DHCP servers defined for assigning virtual IP addresses to VPN clients.

Appendix D

Regular Expression Syntax

Regular expressions are used to define patterns in traffic for custom Situations, which can be used in Inspection rules on security engines.

The following sections are included:

Syntax for Stonesoft Regular Expressions

Stonesoft custom Situations are often defined as text strings using regular expression patterns for matching byte sequences in network traffic.

The expression matching starts always at the beginning of the traffic stream. For this reason, '`.*`' is usually necessary at the beginning of a regular expression to indicate that there can be an undefined number of bytes in the traffic stream preceding the expression.

The regular expression string consists of one or more *branches* that are separated by the '|' logical OR symbol as follows: "*branch-1*|*branch-2*| . . .". A branch contains one or more regular expressions one after another. The Situation matches if any of the regular expression branches separated by '|' matches to the network traffic byte stream.

Note – Regular expressions are case sensitive. Space characters are included in the matching process unless the modifier (?S) or (?x) is used to ignore spaces.

The Stonesoft regular expressions are described in the table below.

Table D.1 Stonesoft Regular Expression Syntax

Sequence	Description	Example
<*char*>	Matches only the defined characters.	'2', 'A', '`foo`' match exactly to the defined characters: '2', 'A', and '`foo`' respectively.
`.` (dot)	matches any character, including the null character `\x00` and a missing character. Matches also other than printable characters, such as the linefeed.	'`.`' matches any single character or byte.
\x<*hex*>	Matches the hexadecimal byte value ranging from `\x00` to `\xFF`.	'`\x4d`' matches hexadecimal value '`4d`' which represents the decimal value 77 and the ASCII character 'M'.
[<*char*>**]**	Match any of the characters in the list.	'`[15aB]`' matches when any of the characters '`1`', '`5`', '`a`', or '`B`' in the matching location of the inspected string.
[^<*char*>**]**	Matches when none of the characters in the list is present.	'`[^aBc]`' matches if none of the characters '`a`', '`B`', or '`c`' is present in the matching location of the inspected string.
[<*char1*>**-**<*char2*>**]**	Matches all the characters ranging from <*char1*> to <*char2*>, these two characters included.	'`[a-f]`' matches any character within the range from '`a`' to '`f`', with '`a`' and '`f`' included.

Table D.1 Stonesoft Regular Expression Syntax (Continued)

Sequence	Description	Example
[:<*class*>**:]**	Used in bracket expression to match any character of the defined class. The <*class*> can be: **alnum** [0-9A-Za-z], **alpha** [A-Za-z], **ascii** (ascii char), **blank** (space or tab), **cntrl** (control char), **digit** [0-9], **graph** (alnum or punct), **lower** [a-z], **print** (printable char), **punct** [.,”’!?;:], **space** (any space char), **upper** [A-Z], **word** (alnum + ‘_’ char), **xdigit** [0-9A-Fa-f].	‘[[:digit:]]’ matches any digit, e.g. 1, 5, or 7.
\\<*char*>	Used for escaping special metacharacters to be interpreted as normal characters, in this case as <*char*>. The metacharacters are: \\\|)(][^-*+?.#	‘\\[’ matches the ‘[’ character instead of interpreting it as the regular expression class metacharacter.
#<*text*>	Anything starting with ‘# ’ up to the linefeed (\x0a) or the carriage return (\x0d) character is considered as a comment and not used in the matching process.	‘# my comment.’ is not used in the matching process.
(<*expr1*>| <*expr2*>**)**	Matches if either the sub-expression <*expr1*> or <*expr2*> matches.	‘a(bc|de)’ matches ‘abc’ and ‘ade’.

It is also possible to indicate repeated, consecutive regular expressions by using quantifiers as described in the table below.

Table D.2 Stonesoft Regular Expression Quantifiers

Quantifier	Description	Example
<*expr*>*****	Matches if there are zero or more consecutive <*expr*> strings.	‘a*’ matches ‘<empty>’, ‘a’, ‘aa’ and so on.
<*expr*>**+**	Matches if there are one or more consecutive <*expr*> strings.	‘a+’ matches ‘a’, ‘aa’, ‘aaa’ and so on, but not the empty string.
<*expr*>**?**	Matches if there is zero or one <*expr*> string.	‘a?’ matches ‘<empty>’ and ‘a’.
<*expr*>**{***n*,*m***}**	{*num*} matches exactly *num* times the expression. {*num*,} matches *num* or more times the expression. {*num*,*max*} matches at least *num* and no more than *max* times the expression.	“a{5,}” matches five or more consecutive ‘a’. “a{5,7}” matches five, six, or seven consecutive ‘a’.

The ‘*’ and ‘+’ wildcard characters in the middle of a regular expression pattern can easily result in an expression that has a very large number of different matching states. For this reason, they must be used with care. The computed matching pattern can grow exponentially, thus becoming too complex for efficient use on the Sensors.

Use the “{*num*,*max*}” quantifier where possible, instead of the ‘*’ and ‘+’ bounds. Variables can also be used to break down the pattern to smaller chunks as described in Bit Variable Extensions (page 1184).

The illustration below provides an example regular expression.

Illustration D.1 Example regular expression

```
# This expression matches any of the following patterns in the traffic:
# ‘/bin/{ash|bash|csh|ksh|sh|tcsh}’

# First, match ‘/bin/sh’ with zero or more characters in front of it:
.*/bin/sh|
# or match ‘/bin/’ with zero or more characters in front of it,
# followed by ‘ash’, ‘csh’, or ‘ksh’:
.*/bin/[ack]sh|
# or match ‘/bin/’ with zero or more characters in front of it,
# followed by ‘bash’ or ‘tcsh’:
.*/bin/(ba|tc)sh

# Alternatively, this expression with all the patterns can be integrated
# into one, for example: .*/bin/(ba|tc|[ack])?sh
```

Special Character Sequences

The printable characters are defined simply by typing them in the regular expression. The hexadecimal values **\x***HH* can also be used to match any byte value (for example, ASCII character). In addition, there are some shorthands for common non-printable characters and character classes. The special character sequences are listed in the table below.

Table D.3 Special Character Sequences

Sequence	Description
\a	Bell (BEL) = \x07
\t	Horizontal tab (HT) = \x09
\n	Linefeed (LF) = \x0A
\f	Formfeed (FF) = \x0C
\r	Carriage return (CR) = \x0D
\e	Escape (ESC) = \x1B
\OOO	Octal code *OOO* of the character.
\x*HH*	Hexadecimal code *HH* of the character.
\c<*char*>	Control character that corresponds to Ctrl+<*char*>
\w	"word" class character = [A-Za-z0-9_]
\W	Non-"word" class character = [^A-Za-z0-9_]
\s	Whitespace character = [\t\r\n\f]

Table D.3 Special Character Sequences (Continued)

Sequence	Description
\S	Non-whitespace character = `[^ \t\r\n\f]`
\d	Digit character = `[0-9]`
\D	Non-digit character = `[^0-9]`
\b	Backspace (BS) = `\x08` Note: allowed only in bracket expressions.
\Q *<expr>* **\E**	Quotes all metacharacters between the `\Q` and `\E`. Backslashes are considered as normal characters. For example, "`\QC:\file.exe\E`" matches the "`C:\file.exe`" string, not the "`C:\x0Cile.exe`" string where `\x0C` is the formfeed "`\f`".

Pattern-Matching Modifiers

The Stonesoft regular expression syntax has Perl-like extensions. The pattern-matching modifiers are extensions that can be used to control the matching process in more detail. The modifiers are enabled with **(?***<modifiers>***)** and disabled with a minus **(?-***<modifiers>***)**, where *<modifiers>* is a list of one or more modifiers.

The modifiers `(?C)`, `(?L)`, and `(?s)`are enabled by default. The pattern-matching modifiers are listed in the table below.

Table D.4 Pattern-Matching Modifiers

Sequence	Description
(?i)	"Case insensitive mode" When enabled, case insensitive matching is used for the uppercase and lowercase letters. Thus, a letter matches regardless of its capitalization. When disabled, the letters are matched case-sensitively so that capitalization is taken into account in the matching process.
(?s)	"Single line mode" When enabled, the dot character '`.`' matches also the null character \x00 and a missing character in addition to matching any character (including linefeed and other non-printable characters). When disabled, the linefeed or a missing character are not matched. This modifier is enabled by default. Use **(?-s)** to disable it.
(?x)	"Extended readability mode" When enabled, equals to enabling (?C), (?L), and (?S). Comments, linefeeds and spaces are not used in the matching process, allowing to use them for readability of the expression. When disabled, equals to disabling (?C), (?L), and (?S).Comments, linefeeds and spaces are used in the matching process.

Table D.4 Pattern-Matching Modifiers (Continued)

Sequence	Description
(?C)	"Allow comments mode" When enabled, anything after the hash character '# ' is considered as a comment and not included in the matching process. When disabled, the hash character '# ' and anything following are used in the matching process. This modifier is enabled by default. Use **(?-C)** to disable it.
(?L)	"Ignore linefeeds mode" When enabled, the linefeed and carriage return characters are not included in the matching process unless specifically defined (\x0A or \n for linefeed and \x0D or \r for carriage return). When disabled, the linefeeds and carriage returns are used in the matching process. This modifier is enabled by default. Use **(?-L)** to disable it.
(?S)	"Ignore spaces mode" When enabled, the space and horizontal tab characters are not used in the matching process unless specifically defined (\x20 for space and \x09 or \t for horizontal tab). When disabled, the space and horizontal tab characters are used in the matching process.
(?<*modifiers*> **:**<*sub-expr*>**)**	Applies the <*modifiers*> modifiers only to the subexpression <*sub-expr*>. These modifiers are not used in other parts of the regular expression.

Bit Variable Extensions

Variables can be used to define regular expression patterns that are related to each other. These relations can be expressed with the variables so that the regular expression matches only when all the related patterns match. Complex matching with multiple Situations is also possible, as the variables and the variable values are shared with all the Situations in a Situation Group.

A variable extension can use the following expressions:

- A *value setting expression* defines the values for one or more variables when the corresponding top-level branch matches.
  For example, **(?{var_a=1})** sets the value 1 for the variable var_a.
- A *conditional expression* inspects the values defined for one or more variables so that the corresponding top-level branch matches, and the optional variable setting expressions are processed only if the conditional expression is true.
  For example, **(?{var_b==1})** matches when the variable var_b is equal to 1.
- When using both variable expression types for the same top-level branch, the implication operator '->' must be used.
  For example, **(?{var_a==1->var_a=0})** matches when the variable var_a is equal to 1, and finally sets the value for this variable to be 0.

Each variable is unique within the Situation or Situation Group where the variable is used. The name of a Situation variable can be anything consisting of alphanumeric and underscore characters [A-Za-z_0-9]. The variable name must not begin with a digit. The variable has a boolean value that can be either 0 or 1. The variable values persist through each individual traffic stream.

Note – In variable expressions a single equal sign '=' sets a value for a variable, whereas two consecutive equal signs '==' evaluate the value of a variable.

Variables are defined with the expressions listed in the table below.

Table D.5 Variable Extensions

Sequence	Description
(?{<*var*>**=**<*value*>**})**	The expression matches and the <*var*> variable's value is set to <*value*> (0 or 1). Multiple value setting expressions can be defined by separating them with a comma '**,**'.
(?{<*var*>**=**<*value*>**,ignore})**	Sets the <*var*> variable's value to <*value*> (0 or 1). The `ignore` keyword is used to indicate a partial match that does not trigger response alone but requires another matching branch.
(?{<*var*>**==**<*value*>**})**	The expression matches only when the <*var*> variable's value is <*value*>. Multiple conditional expressions can be defined by separating them with '**&&**'.
(?{<*var1*>**==**<*value1*>**->**<*var2*>**=**<*value*>**})**	The expression matches only when the <*var1*> variable's value is <*value1*>. When the condition is true, the <*var2*> variable's value is set to <*value2*>.

(?{...}) can be used in the two top-level branches that are separated by the logical OR symbol '|'. A variable extension is processed only when its top-level branch matches.

Illustration D.2 Expression with Variables

```
# Expression matches only when 'POST /attack.asp?' string is followed
# by 'Content-Type: application/x-www-form-urlencoded' string
# with any number of bytes in between.

(?i)#case-insensitive mode

.*POST /attack.asp\?(?{match=1,ignore})|#does not trigger response alone
.*Content-Type: application/x-www-form-urlencoded(?{match==1->match=0})
```

The expression in the illustration below detects two different strings in the same connection. The variable is used so that the response is triggered only when the first branch matches, followed by the second branch match. Neither of the branches trigger the response alone.

Note – A '*' or '?' wildcard in a middle of an expression can result in a computed matching pattern that is too complex for efficient use on the Sensors. Therefore, it is recommended to divide the pattern into two branches as in the illustration above.

Variable Expression Evaluation

Variable expression evaluation is an extension to regular expression syntax that provides the ability to parse values from the input stream, perform arithmetic operations, detect large blocks of data, and use variable larger than one bit. This allows you to create more precise and reliable Situations in cases that are difficult with the traditional regular expression syntax.

Table D.6 Variable Expression Syntax

Sequence	Description
(?[<expression>])	<expression> is one or more comma-separated expressions

Variables can be 1, 8, 16, 32 or 64 bits long. By default, a variable is one bit (either 0 or 1). The default variable size in bits can be changed with a postfix that contains a "@" sign and the number of bits.

Example **test@32 means that the variable test is 32 bits long.**

If the variable name is prefixed with a dollar sign ($), the variable is matched against the current connection instead of the current stream. This matching in both client to server and server to client traffic.

Example **$command_seen@32 checks that a certain command has been issued by the client and the server has accepted the command without errors.**

Each expression has a value after evaluation. The type of the value can be a 32-bit or 64-bit unsigned integer, or a void. The results of Expressions can be used to perform basic integer arithmetic, variable assignment, and comparisons. The order of operations is similar to that of the C programming language, for example A + B * C is A + (B * C), not (A + B) * C. The '->' is lowest in precedence. Statements inside parentheses () are always evaluated first, so the order of operations can be overridden with parentheses,

Table D.7 Operations on Expression Results

Sequence	Description
`false`	Always evaluates to a false.
`true`	Always evaluates to a true.
`<number>`	A literal number in decimal, octal and hexadecimal format, for example "32" or "0x20".
`<var> = <expr>`	Sets a value returned by expression <expr> to a variable <var>. See variable syntax below.
`<var> += <expr>`	Adds the value of variable <var> with the value returned by expression <expr> and sets the result to variable <var>.
`<var> -= <expr>`	Subtracts the value from variable <var> by the value returned by expression <expr> and sets the result to variable <var>.
`<var> *= <expr>`	Multiplies the value of <var> by the value returned by expression <expr> and sets the result to variable <var>.

Table D.7 Operations on Expression Results (Continued)

Sequence	Description
`<var> /= <expr>`	Divides the value of <var> with the value returned by expression <expr> and sets the result to variable <var>.
`<var> %= <expr>`	Divides the value of <var> with the value returned by expression <expr> and sets the modulo of result to variable <var>.
`<var> <<= <expr>`	Shifts the value of <var> to left by number of steps returned by expression <expr> and sets the result to variable <var>.
`<var> >>= <expr>`	Shifts the value of <var> to right by number of steps returned by expression <expr> and sets the result to variable <var>.
`<var> &= <expr>`	Performs bitwise AND with the value of variable <var> and the value returned by expression <expr> and sets the result to variable <var>.
`<var> |= <expr>`	Performs bitwise OR with the value of variable <var> and the value returned by expression <expr> and sets the result to variable <var>.
`<var> ^= <expr>`	Performs bitwise XOR with the value of variable <var> and the value returned by expression <expr> and sets the result to variable <var>.
`<expr_a> -> <expr_b>`	Expression <expr_b> is evaluated only if <expr_a> is true.
`<expr_a> ? <expr_b> : <expr_c>`	Expression <expr_b> is evaluated only if <expr_b> is true and expression <expr_c> is evaluated if <expr_a> is false.
`<expr_a> == <expr_b>`	Test if expressions <expr_a> and <expr_b> return an equal value.
`<expr_a> != <expr_b>`	Test if expressions <expr_a> and <expr_b> do not return an equal value.
`<expr_a> < <expr_b>`	Test if expression <expr_b> returns higher value than expression <expr_a>.
`<expr_a> <= <expr_b>`	Test if expression <expr_b> returns higher or equal value than expression <expr_a>.
`<expr_a> > <expr_b>`	Test if expression <expr_a> returns higher value than expression <expr_b>.
`<expr_a> >= <expr_b>`	Test if expression <expr_a> returns higher or equal value than expression <expr_b>.
`<expr_a> & <expr_b>`	Performs bitwise AND with expressions <expr_a> and <expr_b> and returns the result.
`<expr_a> | <expr_b>`	Performs bitwise OR with expressions <expr_a> and <expr_b> and returns the result.
`<expr_a> ^ <expr_b>`	Performs bitwise XOR with expressions <expr_a> and <expr_b> and returns the result.
`<expr_a> && <expr_b>`	Performs AND with expressions <expr_a> and <expr_b> and returns the result.

Table D.7 Operations on Expression Results (Continued)

Sequence	Description
`<expr_a> || <expr_b>`	Performs OR with if expressions <expr_a> and <expr_b> and returns the result.
`<var>++, ++<var>`	Increase value of variable <var> by one.
`<var>--, --<var>`	Decrease value of variable <var> by one.
`-<expr>`	Negate the result of the expression <expr>.
`~<expr>`	Bitwise invert the result of the expression <expr>.
`!<expr>`	Perform NOT operation with the expression <expr>.

Stream Operations

The binary data from the input stream can be read into variables with the following expressions:

Table D.8 Binary Data Variable Expressions

Sequence	Description
`parse_be@<size>`	Parse big endian value. <size> is the size of the value to be read in bits, and it can be one of the following: 8, 16, 24, 32, 40, 48, 56 or 64.
`parse_le@<size>`	Parse little endian value. <size> is the size of the value to be read in bits, and it can be one of the following: 8, 16, 24, 32, 40, 48, 56 or 64.

ASCII values can be read from the input stream with the following expressions:

Table D.9 ASCII Data Variable Expressions

Sequence	Description
`parse_dec(<length>)`	Parse ASCII decimal value. <length> is the maximum number of the characters to parse. The actual number of parsed digits is available in the variable $parse_length@32. If no characters could be parsed, then the variable is set to zero.
`parse_hex(<length>)`	Parse ASCII hexadecimal value. <length> is the maximum number of the characters to parse. The actual number of parsed digits is available in the variable $parse_length@32. If no characters could be parsed, then the variable is set to zero.
`parse_int(<length>)`	Parse ASCII value; parses hexadecimal if the string starts with "0x", octal if the string starts with zero ("0") and decimal otherwise. <length> is the maximum number of the characters to parse. The actual number of parsed digits is available in the variable $parse_length@32. If no characters could be parsed, then the variable is set to zero.

Table D.9 ASCII Data Variable Expressions (Continued)

Sequence	Description
`parse_oct(<length>)`	Parse ASCII octal value. <length> is the maximum number of the characters to parse. The actual number of parsed digits is available in the variable $parse_length@32. If no characters could be parsed, then the variable is set to zero.

Miscellaneous operations with the input stream:

Table D.10 Miscellaneous Input Stream Operations

Sequence	Description
`CRC(<length>)`	Calculates a 32-bit CRC value starting from the current byte up to number of bytes specified by <length> parameter. This function is suitable to detect large binary blocks from the input stream.
`skip(<length>)`	Skip <length> number of bytes.
`regex(<regexp>)`	Launch independent subexpression. See section "Independent Subexpression" for more information.

Other Expressions

Table D.11 Other Expressions

Sequence	Description
`sid()`	Generate a situation. This expression is used to generate a situation indicating a match.
`sid(<id>)`	Generate a specific situation specified by <id>.
`cancel`	Stop matching in the current level.
`cancel_fp_ctx`	Stop matching in the current fingerprinting context.

The illustration below provides an example of a regular expression that launches a variable expression.

Illustration D.3 Expression with Variable Expression

```
# Launches a variable expression when byte 20h (space character) is seen in
# the input stream. The evaluation first assigns variable aa@16 a value of
# 8, checks if the value of variable bb@32 is 47, and if so, generates a
# situation.
.*\x20(?[aa@16=8, bb@32==47 -> sid()])
```

When `(?[])` is used (at the end of a regular expression branch or elsewhere), the situation is not reported automatically. `sid()` must be used explicitly. This is different from `(?{})`, where the situation is automatically reported if **`ignore`** is not used.

System Variables

The syntax of the system variables is the same as for other variables (see the table Table D.5, except that the variable's value is not user-changeable.

Table D.12 System Variables

Sequence	Description
$major	The major version number of the Stonesoft engine.
$minor	The minor version number of the Stonesoft engine.
$patch	The patch level number of the Stonesoft engine.
$build	The build number of the Stonesoft engine.
$fpver	The version number of the current fingerprint matching engine.
$dport	The current destination port of the connection. This value can be used to limit matching to traffic that is destined to a specific port.
$offset	The byte that is under inspection when counted from the beginning of the traffic stream.
$parse_length@32	Number of digits parsed by last parse_dec(), parse_hex(), parse_oct() or parse_in() expression. See Stream Operations below.

The illustration below provides example of an offset expression that matches when `'\xff\x53\x4d\x42\x25'` occurs in the stream, so that `\x25` is the 25th byte in the stream.

Illustration D.4 Expression with Variables

```
# Expression matches when '\xff\x53\x4d\x42\x25' occurs in the stream with
\x25 as the # 25th byte in the stream. The 25th byte in the stream has an
offset value
# of 24.

.*\xff\x53\x4d\x42\x25(?{$offset==24})
```

Independent Subexpressions

Independent subexpressions allow another regular expression to be launched independently from the main regular expression. The syntax for the independent subexpression is as follows:

Table D.13 Independent Subexpression Syntax

Sequence	Description
`(?><regular_expression>)`	<regular_expression> is a normal Stonesoft regular expression launched independently from the main regular expression.
`(?>(?{<expression>})<regular_expression>)`	<expression> is a comparison expression that is evaluated before the independent subexpression <regular_expression> is launched. <regular expression> is launched only if <expression> evaluates to true.

Within `(?[...])` independent subexpressions can be launched by `regex(<regular_expression>)`.

The illustration below provides an example of a regular expression that launches independent subexpression.

Illustration D.5 Expression with Independent Subexpressions

```
# Main expression matches when a GET string has been seen in the input stream
# and launches an independent subexpression. The independent subexperssion
# detects whether the parameter for the GET request is longer than 400
# characters.

*GET(?>\s*[^\s]{400})
```

The illustration below provides an example of an independent subexpression that includes a precondition parameter check. The independent subexpression is launched only if the precondition expression evaluates to true.

Illustration D.6 Independent Subexpression with Precondition Parameter Check

```
# Launches independent subexpression for a Content-length: HTTP header only
# if it is seen in a POST request.

(?x)
(?i)
.*POST(?{post_seen=1,ignore})|
.*\nContent-length:(?>(?{post_seen==1}[^\n]{1024})
```

Parallel Matching Groups

Stonesoft allows you to set different regular expressions to be matched in parallel groups within one Situation Context. Normally, manual situation group definitions are not needed and Stonesoft automatically compiles all your custom Situations in the same group (group 0).

Manual group definitions is needed if the IPS policy upload fails due to fingerprint/DFA compilation problems that may occur with complex regular expressions.

To use grouping, add a new preprocessing tag to the beginning of the regular expression:

Table D.14 Preprocessing Tag for Setting a Group for Matching

Syntax	Description

`#!!GROUP(X)`
`Comment`
`#!!#`
	'X' is the group number from 0 to 7. The comment is optional. If you do not specify the group with this tag, the Situation is processed in group zero.

Illustration D.7 Setting a parallel Matching Group

```
#!!GROUP(1)
This heavy regular expression is matched in parallel matching group 1.
#!!#

#Insert regular expression below
```

APPENDIX E

SNMP TRAPS AND MIBS

Firewall/VPN, IPS, and Layer 2 Firewall engines can send SNMP traps on system events. The traps are configured using SNMP Agent elements. Additionally, Tester entries can be configured to send SNMP traps. The SNMP traps are listed in the table below.

Table E.1 SNMP Traps for Stonesoft Firewall/VPN, IPS, and Layer 2 Firewalls

Trap Name	Objects Included	Description
fwPolicyInstall	fwSecurityPolicy	(*Firewall and Layer 2 Firewall*) Policy was installed on the Firewall engine.
ipsPolicyInstall	ipsSecurityPolicy	(*IPS*) Policy was installed on the IPS engine.
nodeBoot	-	Node bootup complete.
nodeHwmon	nodeHwmonEvent	Hardware monitoring system has detected problems.
nodeOffline	nodeOperState	Node changed to offline or standby state.
nodeOnline	nodeOperState	Node changed to online state.
nodeShutdown	-	Node is shutting down.
nodeTestFailure	nodeTestIdentity	Test subsystem reported a test failure on the node.
nodeFailedUserLogin	nodeLastLogin	(*Firewall and Layer 2 Firewall*) Login failed on the firewall engine’s console or through SSH.
nodeUserLogin	nodeLastLogin	Login initiated on the engine’s console or through SSH.
nodeUserLogout	nodeLastLogin	(*Firewall and Layer 2 Firewall*) Logout on the firewall engine’s console or through SSH.

The STONESOFT-SMI-MIB defines the top-level enterprise registrations for the Stonesoft's products in the .iso.org.dod.internet.private.enterprises.stonesoft branch (OID `.1.3.6.1.4.1.1369`). The Stonesoft-specific MIB files can be downloaded from the Stonesoft website at http://www.stonesoft.com/

The Stonesoft-specific MIBs are:

- STONEGATE-FIREWALL-MIB: see Table E.2.
- STONEGATE-IPS-MIB: see Table E.3.
- STONEGATE-NETNODE-MIB: see Table E.4.

Security Engines in the Firewall/VPN and Layer 2 Firewall roles support objects in STONEGATE-FIREWALL-MIB. Security Engines in the IPS role support objects in STONEGATE-IPS-MIB. Security Engines in all roles support objects in STONEGATE-NETNODE-MIB.

Security Engines in the Firewall/VPN role also support objects in the following standard MIBs:

- IF-MIB (RFC 2863 and RFC 2233): see Table E.5.
- IP-MIB (RFC 2011): see Table E.6.
- SNMP-USER-BASED-SM-MIB (RFC 3414): see Table E.7.
- SNMPv2 MIB (RFC 3418): see Table E.8.

Table E.2 STONEGATE-FIREWALL-MIB Objects

Object Name	Object Description in MIB
fwPolicyTime	The time when the security policy was installed to the Firewall or Layer 2 Firewall
fwSecurityPolicy	Name of the current security policy on the Firewall or Layer 2 Firewall
fwSoftwareVersion	Version string of the Firewall or Layer 2 Firewall software
fwConnNumber	Number of current connections
fwAccepted	Number of accepted packets
fwDropped	Number of dropped packets
fwLogged	Number of logged packets
fwAccounted	Number of accounted packets
fwRejected	Number of rejected packets
fwIfTable	This table contains an entry for each interface in system
fwIfStatsEntry	Row for a interface
fwIfStatsIndex	A unique value, greater than zero, for each interface or interface sub-layer in the managed system
fwIfName	Name of interface
fwIfAcceptedPkts	Number of accepted packets by Firewall or Layer 2 Firewall rules
fwIfDroppedPkts	Number of dropped packets by Firewall or Layer 2 Firewall rules
fwIfForwardedPkts	Number of forwarded packets by Firewall or Layer 2 Firewall rules

Table E.2 STONEGATE-FIREWALL-MIB Objects (Continued)

Object Name	Object Description in MIB
fwIfLoggedPkts	Number of logged packets by Firewall or Layer 2 Firewall rules
fwIfRejectedPkts	Number of rejected packets by Firewall or Layer 2 Firewall rules
fwIfAccountedPkts	Number of accounted packets by Firewall or Layer 2 Firewall rules
fwIfAcceptedBytes	Number of accepted bytes by Firewall or Layer 2 Firewall rules
fwIfForwardedBytes	Number of forwarded bytes by Firewall or Layer 2 Firewall rules
fwIfDroppedBytes	Number of dropped bytes by Firewall or Layer 2 Firewall rules
fwIfLoggedBytes	Number of logged bytes by Firewall or Layer 2 Firewall rules
fwIfRejectedBytes	Number of rejected bytes by Firewall or Layer 2 Firewall rules
fwIfAccountedBytes	Number of accounted bytes by Firewall or Layer 2 Firewall rules
fwCpuTable	This table contains an entry for each CPU in a system and total usage of all CPUs
fwCpuStats	Row with information about CPU usage
fwCpuStatsId	A unique value, greater than zero, for each CPU in the managed system. First element with Id '0' is designed for total values
fwCpuName	Name of data current line concern
fwCpuTotal	The total CPU load percentage
fwCpuUser	The percentage of time the CPU has spent running users' processes that are not niced
fwCpuSystem	The percentage of time the CPU has spent running the kernel and its processes
fwCpuNice	The percentage of time the CPU has spent running user’s processes that have been niced
fwCpuIdle	The percentage of time the CPU was idle
fwCpuIoWait	The percentage of time the CPU has been waiting for I/O to complete
fwCpuHwIrq	The percentage of time the CPU has been servicing hardware interrupts
fwCpuSoftIrq	The percentage of time the CPU has been servicing software interrupts
fwSwapBytesTotal	Total swap space
fwSwapBytesUsed	Used space of swap
fwSwapBytesUnused	Amount of unused space of swap
fwMemBytesTotal	Number of available bytes of physical memory
fwMemBytesUsed	Amount of memory being in use
fwMemBytesUnused	Amount of unused bytes of physical memory
fwMemBytesBuffers	Amount of memory used as buffers

Table E.2 STONEGATE-FIREWALL-MIB Objects (Continued)

Object Name	Object Description in MIB
fwMemBytesCached	Amount of memory used as cache
fwDiskSpaceUsageTable	Table contains an entry for each partition mounted in a system
fwDiskStats	Row of information concerning one partition
fwPartitionIndex	A unique value, greater than zero, for each partition
fwPartitionDevName	A unique name of a device
fwMountPointName	Name of a mount point
fwPartitionSize	Total size of the partition
fwPartitionUsed	Amount of used space of the partition (in kilobytes)
fwPartitionAvail	Information about amount of free space on partition (in kilobytes)
adslModulation	Modulation protocol
adslChannel	Channel type
adslConnStatus	The status of the DSL link or communication status with DSL modem in case of communication error
adslConnUptime	Uptime of current ADSL connection
adslLineStatus	Current status of DSL line
adslInOctets	Number of bytes received by ADSL interface
adslOutOctets	Number of bytes transmitted by ADSL interface
adslSynchroSpeedUp	The actual rate at which data is flowing upstream
adslSynchroSpeedDown	The actual rate at which data is flowing downstream
adslAttenuationUp	An estimate of the average loop attenuation upstream
adslAttenuationDown	An estimate of the average loop attenuation downstream
adslNoiseMarginUp	This is a signal-to-noise ratio (SNR) margin for traffic going upstream
adslNoiseMarginDown	This is a signal-to-noise ratio (SNR) margin for traffic going downstream
adslHecErrorsUp	The total number of header error checksum errors upstream
adslHecErrorsDown	The total number of header error checksum errors downstream
adslOcdErrorsUp	The number of out-of-cell delineation errors upstream
adslOcdErrorsDown	The number of out-of-cell delineation errors downstream
adslLcdErrorsUp	The total of lost-cell-delineation errors upstream
adslLcdErrorsDown	The total of lost-cell-delineation errors downstream

Table E.2 STONEGATE-FIREWALL-MIB Objects (Continued)

Object Name	Object Description in MIB
adslBitErrorsUp	The number of bit errors upstream
adslBitErrorsDown	The number of bit errors downstream

Table E.3 STONEGATE IPS-MIB Objects

Object Name	Object Description in MIB
ipsPolicyTime	The time when the security policy was installed to the IPS engine
ipsSecurityPolicy	Name of the current security policy on the IPS engine
ipsSoftwareVersion	Version string of the IPS software

Table E.4 STONEGATE-NETNODE-MIB Objects

Object Name	Object Description in MIB
nodeClusterId	The identification number of the cluster this node belongs to
nodeCPULoad	The CPU load percentage on the node
nodeHwmonEvent	Reason for the hardware monitoring event
nodeLastLogin	The most recent login event on the node
nodeLastLoginTime	Timestamp of the most recent login event on the node
nodeMemberId	Node's member identification within the cluster
nodeOperState	The operative (clustering) state of the node
nodeTestIdentity	Identification string of a nodeTest
nodeTestResult	The most recent result of the nodeTest
nodeTestResultTime	The timestamp of the most recent result of the nodeTest

Table E.5 IF-MIB Supported Objects

Object Name	Object Description in MIB
ifAdminStatus	The desired state of the interface. The testing(3) state indicates that no operational packets can be passed. When a managed system initializes, all interfaces start with ifAdminStatus in the down(2) state. As a result of either explicit management action or per configuration information retained by the managed system, ifAdminStatus is then changed to either the up(1) or testing(3) states (or remains in the down(2) state).
ifAlias	This object is an 'alias' name for the interface as specified by a network manager, and provides a non-volatile 'handle' for the interface. On the first instantiation of an interface, the value of ifAlias associated with that interface is the zero-length string. As and when a value is written into an instance of ifAlias through a network management set operation, then the agent must retain the supplied value in the ifAlias instance associated with the same interface for as long as that interface remains instantiated, including across all re- initializations/reboots of the network management system, including those which result in a change of the interface's ifIndex value. An example of the value which a network manager might store in this object for a WAN interface is the (Telco's) circuit number/identifier of the interface. Some agents may support write-access only for interfaces having particular values of ifType. An agent which supports write access to this object is required to keep the value in non-volatile storage, but it may limit the length of new values depending on how much storage is already occupied by the current values for other interfaces.
ifDescr	A textual string containing information about the interface. This string includes the name of the manufacturer, the product name and the version of the interface hardware/software.
ifHCInMulticastPkts	The 64-bit wide number of packets, delivered by this sub-layer to a higher (sub-)layer, which were addressed to a multicast address at this sub-layer. For a MAC layer protocol, this includes both Group and Functional addresses. This object is a 64-bit version of ifInMulticastPkts. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of ifCounterDiscontinuityTime. The 32-bit ifInMulticastPkts reports the low 32-bits of this counter’s value.
ifHCInOctets	The 64-bit wide total number of octets received on the interface, including framing characters. This object is a 64-bit version of ifInOctets. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of ifCounterDiscontinuityTime. The 32-bit ifInOctets reports the low 32-bits of this counter’s value.
ifHCInUcastPkts	The 64-bit wide number of packets, delivered by this sub-layer to a higher (sub-)layer, which were not addressed to a multicast or broadcast address at this sub-layer. This object is a 64-bit version of ifInUcastPkts. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of ifCounterDiscontinuityTime. The 32-bit ifInUcastPkts reports the low 32-bits of this counter’s value.

Table E.5 IF-MIB Supported Objects (Continued)

Object Name	Object Description in MIB
ifHCOutOctets	The 64-bit wide total number of octets transmitted out of the interface, including framing characters. This object is a 64-bit version of ifOutOctets. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of ifCounterDiscontinuityTime. The 32-bit ifOutOctets reports the low 32-bits of this counter's value.
ifHCOutUcastPkts	The 64-bit wide total number of packets that higher-level protocols requested be transmitted, and which were not addressed to a multicast or broadcast address at this sub-layer, including those that were discarded or not sent. This object is a 64-bit version of ifOutUcastPkts. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of ifCounterDiscontinuityTime. The 32-bit ifOutUcastPkts reports the low 32-bits of this counter's value.
ifHighSpeed	An estimate of the interface's current bandwidth in units of 1,000,000 bits per second. If this object reports a value of 'n' then the speed of the interface is somewhere in the range of 'n-500,000' to 'n+499,999'. For interfaces which do not vary in bandwidth or for those where no accurate estimation can be made, this object contains the nominal bandwidth. For a sub-layer which has no concept of bandwidth, this object must be zero.
ifIndex	A unique value, greater than zero, for each interface. It is recommended that values are assigned contiguously starting from 1. The value for each interface sub-layer must remain constant at least from one re-initialization of the entity's network management system to the next re- initialization.
ifInDiscards	The number of inbound packets which were chosen to be discarded even though no errors had been detected to prevent their being deliverable to a higher-layer protocol. One possible reason for discarding such a packet could be to free up buffer space. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of ifCounterDiscontinuityTime.
ifInErrors	For packet-oriented interfaces, the number of inbound packets that contained errors preventing them from being deliverable to a higher-layer protocol. For character-oriented or fixed-length interfaces, the number of inbound transmission units that contained errors preventing them from being deliverable to a higher-layer protocol. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of ifCounterDiscontinuityTime.
ifInMulticastPkts	The 32-bit wide number of packets, delivered by this sub-layer to a higher (sub-)layer, which were addressed to a multicast address at this sub-layer. For a MAC layer protocol, this includes both Group and Functional addresses. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of ifCounterDiscontinuityTime. This object reports the low 32-bits of the 64-bit ifHCInMulticastPkts counter's value.

Table E.5 IF-MIB Supported Objects (Continued)

Object Name	Object Description in MIB
ifInOctets	The 32-bit wide total number of octets received on the interface, including framing characters. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of ifCounterDiscontinuityTime. This object reports the low 32-bits of the 64-bit ifHCInOctets counter's value.
ifInUcastPkts	The 32-bit wide number of packets, delivered by this sub-layer to a higher (sub-)layer, which were not addressed to a multicast or broadcast address at this sub-layer. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of ifCounterDiscontinuityTime. This object reports the low 32-bits of the 64-bit ifHCInUcastPkts counter's value.
ifLastChange	The value of sysUpTime at the time the interface entered its current operational state. If the current state was entered prior to the last re-initialization of the local network management subsystem, then this object contains a zero value.
ifLinkUpDownTrapEnable	Indicates whether linkUp/linkDown traps are generated for this interface. By default, this object must have the value enabled(1) for interfaces which do not operate on 'top' of any other interface (as defined in the ifStackTable), and disabled(2) otherwise.
ifMtu	The size of the largest packet which can be sent/received on the interface, specified in octets. For interfaces that are used for transmitting network datagrams, this is the size of the largest network datagram that can be sent on the interface.
ifName	The textual name of the interface. The value of this object must be the name of the interface as assigned by the local device and must be suitable for use in commands entered at the device's ‘console'. This might be a text name, such as ‘le0' or a simple port number, such as ‘1', depending on the interface naming syntax of the device. If several entries in the ifTable together represent a single interface as named by the device, then each will have the same value of ifName. Note that for an agent which responds to SNMP queries concerning an interface on some other (proxied) device, then the value of ifName for such an interface is the proxied device's local name for it. If there is no local name, or this object is otherwise not applicable, then this object contains a zero-length string.
ifNumber	The number of network interfaces (regardless of their current state) present on this system.
ifOperStatus	The current operational state of the interface. The testing(3) state indicates that no operational packets can be passed. If ifAdminStatus is down(2) then ifOperStatus is down(2). If ifAdminStatus is changed to up(1) then ifOperStatus changes to up(1) if the interface is ready to transmit and receive network traffic; it changes to dormant(5) if the interface is waiting for external actions (such as a serial line waiting for an incoming connection); it remains in the down(2) state if and only if there is a fault that prevents it from going to the up(1) state; it remains in the notPresent(6) state if the interface has missing (typically, hardware) components.

Table E.5 IF-MIB Supported Objects (Continued)

Object Name	Object Description in MIB
ifOutDiscards	The number of outbound packets which were chosen to be discarded even though no errors had been detected to prevent their being transmitted. One possible reason for discarding such a packet could be to free up buffer space. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of ifCounterDiscontinuityTime.
ifOutErrors	For packet-oriented interfaces, the number of outbound packets that could not be transmitted because of errors. For character-oriented or fixed-length interfaces, the number of outbound transmission units that could not be transmitted because of errors. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of ifCounterDiscontinuityTime.
ifOutOctets	The 32-bit wide total number of octets transmitted out of the interface, including framing characters. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of ifCounterDiscontinuityTime. This object reports the low 32-bits of the 64-bit ifHCOutOctets counter’s value.
ifOutUcastPkts	The 32-bit wide total number of packets that higher-level protocols requested be transmitted, and which were not addressed to a multicast or broadcast address at this sub-layer, including those that were discarded or not sent. Discontinuities in the value of this counter can occur at re-initialization of the management system, and at other times as indicated by the value of ifCounterDiscontinuityTime. This object reports the low 32-bits of the 64-bit ifHCOutUcastPkts counter’s value.
ifPhysAddress	The interface's address at its protocol sub-layer. For example, for an 802.x interface, this object normally contains a MAC address. The interface's media-specific MIB must define the bit and byte ordering and the format of the value of this object. For interfaces that do not have such an address (for example, a serial line), this object must contain an octet string of zero length.
ifPromiscuousMode	This object has a value of false(2) if this interface only accepts packets/frames that are addressed to this station. This object has a value of true(1) when the station accepts all packets/frames transmitted on the media. The value true(1) is only legal on certain types of media. If legal, setting this object to a value of true(1) may require the interface to be reset before becoming effective. The value of ifPromiscuousMode does not affect the reception of broadcast and multicast packets/frames by the interface.
ifSpeed	An estimate of the interface's current bandwidth in bits per second. For interfaces which do not vary in bandwidth or for those where no accurate estimation can be made, this object must contain the nominal bandwidth. If the bandwidth of the interface is greater than the maximum value reportable by this object then this object must report its maximum value (4,294,967,295) and ifHighSpeed must be used to report the interface’s speed. For a sub-layer which has no concept of bandwidth, this object must be zero.
ifType	The type of interface. Additional values for ifType are assigned by the Internet Assigned Numbers Authority (IANA), through updating the syntax of the IANAifType textual convention.

Table E.6 IP-MIB Supported Objects

Object Name	Object Description in MIB
icmpInAddrMaskReps	The number of ICMP Address Mask Reply messages received.
icmpInAddrMasks	The number of ICMP Address Mask Request messages received.
icmpInDestUnreachs	The number of ICMP Destination Unreachable messages received.
icmpInEchoReps	The number of ICMP Echo Reply messages received.
icmpInEchos	The number of ICMP Echo (request) messages received.
icmpInErrors	The number of ICMP messages which the entity received but determined as having ICMP-specific errors (bad ICMP checksums, bad length, etc.).
icmpInMsgs	The total number of ICMP messages which the entity received. Note that this counter includes all those counted by icmpInErrors.
icmpInParmProbs	The number of ICMP Parameter Problem messages received.
icmpInRedirects	The number of ICMP Redirect messages received.
icmpInSrcQuenchs	The number of ICMP Source Quench messages received.
icmpInTimeExcds	The number of ICMP Time Exceeded messages received.
icmpInTimestampReps	The number of ICMP Timestamp Reply messages received.
icmpInTimestamps	The number of ICMP Timestamp (request) messages received.
icmpOutAddrMaskReps	The number of ICMP Address Mask Reply messages sent.
icmpOutAddrMasks	The number of ICMP Address Mask Request messages sent.
icmpOutDestUnreachs	The number of ICMP Destination Unreachable messages sent.
icmpOutEchoReps	The number of ICMP Echo Reply messages sent.
icmpOutEchos	The number of ICMP Echo (request) messages sent.
icmpOutErrors	The number of ICMP messages which this entity did not send due to problems discovered within ICMP such as a lack of buffers. This value must not include errors discovered outside the ICMP layer such as the inability of IP to route the resultant datagram. In some implementations there may be no types of error which contribute to this counter's value.
icmpOutMsgs	The total number of ICMP messages which this entity attempted to send. Note that this counter includes all those counted by icmpOutErrors.
icmpOutParmProbs	The number of ICMP Parameter Problem messages sent.
icmpOutRedirects	The number of ICMP Redirect messages sent. For a host, this object will always be zero, since hosts do not send redirects.
icmpOutSrcQuenchs	The number of ICMP Source Quench messages sent.
icmpOutTimeExcds	The number of ICMP Time Exceeded messages sent.

Table E.6 IP-MIB Supported Objects (Continued)

Object Name	Object Description in MIB
icmpOutTimestampReps	The number of ICMP Timestamp Reply messages sent.
icmpOutTimestamps	The number of ICMP Timestamp (request) messages sent.
ipAdEntAddr	The IP address to which this entry's addressing information pertains.
ipAdEntBcastAddr	The value of the least-significant bit in the IP broadcast address used for sending datagrams on the (logical) interface associated with the IP address of this entry. For example, when the Internet standard all-ones broadcast address is used, the value will be 1. This value applies to both the subnet and network broadcasts addresses used by the entity on this (logical) interface.
ipAdEntIfIndex	The index value which uniquely identifies the interface to which this entry is applicable. The interface identified by a particular value of this index is the same interface as identified by the same value of RFC 1573's ifIndex.
ipAdEntNetMask	The subnet mask associated with the IP address of this entry. The value of the mask is an IP address with all the network bits set to 1 and all the hosts bits set to 0.
ipAdEntReasmMaxSize	The size of the largest IP datagram which this entity can re-assemble from incoming IP fragmented datagrams received on this interface.
ipDefaultTTL	The default value inserted into the Time-To-Live field of the IP header of datagrams originated at this entity, whenever a TTL value is not supplied by the transport layer protocol.
ipForwarding	The indication of whether this entity is acting as an IP router in respect to the forwarding of datagrams received by, but not addressed to, this entity. IP routers forward datagrams. IP hosts do not (except those source-routed via the host).
ipForwDatagrams	The number of input datagrams for which this entity was not their final IP destination, as a result of which an attempt was made to find a route to forward them to that final destination. In entities which do not act as IP routers, this counter will include only those packets which were Source-Routed via this entity, and the Source-Route option processing was successful.
ipFragCreates	The number of IP datagram fragments that have been generated as a result of fragmentation at this entity.
ipFragFails	The number of IP datagrams that have been discarded because they needed to be fragmented at this entity but could not be, for example because their Don't Fragment flag was set.
ipFragOKs	The number of IP datagrams that have been successfully fragmented at this entity.
ipInAddrErrors	The number of input datagrams discarded because the IP address in their IP header's destination field was not a valid address to be received at this entity. This count includes invalid addresses (for example, 0.0.0.0) and addresses of unsupported Classes (for example, Class E). For entities which are not IP routers and therefore do not forward datagrams, this counter includes datagrams discarded because the destination address was not a local address.

Table E.6 IP-MIB Supported Objects (Continued)

Object Name	Object Description in MIB
ipInDelivers	The total number of input datagrams successfully delivered to IP user-protocols (including ICMP).
ipInDiscards	The number of input IP datagrams for which no problems were encountered to prevent their continued processing, but which were discarded (for example, for lack of buffer space). Note that this counter does not include any datagrams discarded while awaiting re-assembly.
ipInHdrErrors	The number of input datagrams discarded due to errors in their IP headers, including bad checksums, version number mismatch, other format errors, time-to-live exceeded, errors discovered in processing their IP options, etc.
ipInReceives	The total number of input datagrams received from interfaces, including those received in error.
ipInUnknownProtos	The number of locally addressed datagrams received successfully but discarded because of an unknown or unsupported protocol.
ipNetToMediaIfIndex	The interface on which this entry's equivalence is effective. The interface identified by a particular value of this index is the same interface as identified by the same value of RFC 1573's ifIndex.
ipNetToMediaNetAddress	The IpAddress corresponding to the media-dependent ‘physical' address.
ipNetToMediaPhysAddress	The media-dependent ‘physical' address.
ipNetToMediaType	The type of mapping. Setting this object to the value invalid(2) has the effect of invalidating the corresponding entry in the ipNetToMediaTable. That is, it effectively disassociates the interface identified with said entry from the mapping identified with said entry. It is an implementation- specific matter as to whether the agent removes an invalidated entry from the table. Accordingly, management stations must be prepared to receive tabular information from agents that corresponds to entries not currently in use. Proper interpretation of such entries requires examination of the relevant ipNetToMediaType object.
ipOutDiscards	The number of output IP datagrams for which no problem was encountered to prevent their transmission to their destination, but which were discarded (for example, for lack of buffer space). Note that this counter would include datagrams counted in ipForwDatagrams if any such packets met this (discretionary) discard criterion.
ipOutNoRoutes	The number of IP datagrams discarded because no route could be found to transmit them to their destination. Note that this counter includes any packets counted in ipForwDatagrams which meet this ‘no-route' criterion. Note that this includes any datagrams which a host cannot route because all of its default routers are down.
ipOutRequests	The total number of IP datagrams which local IP user- protocols (including ICMP) supplied to IP in requests for transmission. Note that this counter does not include any datagrams counted in ipForwDatagrams.
ipReasmFails	The number of failures detected by the IP re-assembly algorithm (for whatever reason: timed out, errors, etc.). Note that this is not necessarily a count of discarded IP fragments since some algorithms (notably the algorithm in RFC 815) can lose track of the number of fragments by combining them as they are received.

Table E.6 IP-MIB Supported Objects (Continued)

Object Name	Object Description in MIB
ipReasmOKs	The number of IP datagrams successfully re-assembled.
ipReasmReqds	The number of IP fragments received which needed to be reassembled at this entity.
ipReasmTimeout	The maximum number of seconds which received fragments are held while they are awaiting reassembly at this entity.

Table E.7 SNMP-USER-BASED-SM-MIB Objects

Object Name	Object Description in MIB
usmStatsDecryptionErrors	The total number of packets received by the SNMP engine which were dropped because they could not be decrypted.
usmStatsNotInTimeWindows	The total number of packets received by the SNMP engine which were dropped because they appeared outside of the authoritative SNMP engine's window.
usmStatsUnknownEngineIDs	The total number of packets received by the SNMP engine which were dropped because they referenced an snmpEngineID that was not known to the SNMP engine.
usmStatsUnknownUserNames	The total number of packets received by the SNMP engine which were dropped because they referenced a user that was not known to the SNMP engine.
usmStatsUnsupportedSecLevels	The total number of packets received by the SNMP engine which were dropped because they requested a security Level that was unknown to the SNMP engine or otherwise unavailable.
usmStatsWrongDigests	The total number of packets received by the SNMP engine which were dropped because they didn't contain the expected digest value.
usmUserSpinLock	An advisory lock used to allow several cooperating Command Generator Applications to coordinate their use of facilities to alter secrets in the usmUserTable.
usmUserStatus	The status of this conceptual row. Until instances of all corresponding columns are appropriately configured, the value of the corresponding instance of the usmUserStatus column is 'notReady'. In particular, a newly created row for a user who employs authentication, cannot be made active until the corresponding usmUserCloneFrom and usmUserAuthKeyChange have been set. Further, a newly created row for a user who also employs privacy, cannot be made active until the usmUserPrivKeyChange has been set. The RowStatus TC [RFC2579] requires that this DESCRIPTION clause states under which circumstances other objects in this row can be modified: The value of this object has no effect on whether other objects in this conceptual row can be modified, except for usmUserOwnAuthKeyChange and usmUserOwnPrivKeyChange. For these 2 objects, the value of usmUserStatus MUST be active.

Table E.8 SNMPv2-MIB Supported Objects

Object Name	Object Description in MIB
snmpEnableAuthenTraps	Indicates whether the SNMP entity is permitted to generate authenticationFailure traps. The value of this object overrides any configuration information; as such, it provides a means whereby all authenticationFailure traps may be disabled. Note that it is strongly recommended that this object be stored in non-volatile memory so that it remains constant across re-initializations of the network management system.
snmpInASNParseErrs	The total number of ASN.1 or BER errors encountered by the SNMP entity when decoding received SNMP messages.
snmpInBadCommunityNames	The total number of SNMP messages delivered to the SNMP entity which used a SNMP community name not known to said entity.
snmpInBadCommunityUses	The total number of SNMP messages delivered to the SNMP entity which represented an SNMP operation which was not allowed by the SNMP community named in the message.
snmpInBadVersions	The total number of SNMP messages which were delivered to the SNMP entity and were for an unsupported SNMP version.
snmpInPkts	The total number of messages delivered to the SNMP entity from the transport service.
snmpProxyDrops	The total number of GetRequest-PDUs, GetNextRequest-PDUs, GetBulkRequest-PDUs, SetRequest-PDUs, and InformRequest-PDUs delivered to the SNMP entity which were silently dropped because the transmission of the (possibly translated) message to a proxy target failed in a manner (other than a time-out) such that no Response-PDU could be returned.
snmpSetSerialNo	An advisory lock used to allow several cooperating SNMPv2 entities, all acting in a manager role, to coordinate their use of the SNMPv2 set operation. This object is used for coarse-grain coordination. To achieve fine-grain coordination, one or more similar objects might be defined within each MIB group, as appropriate.
snmpSilentDrops	The total number of GetRequest-PDUs, GetNextRequest-PDUs, GetBulkRequest-PDUs, SetRequest-PDUs, and InformRequest-PDUs delivered to the SNMP entity which were silently dropped because the size of a reply containing an alternate Response-PDU with an empty variable-bindings field was greater than either a local constraint or the maximum message size associated with the originator of the request.
sysContact	The textual identification of the contact person for this managed node, together with information on how to contact this person. If no contact information is known, the value is the zero-length string.
sysDescr	A textual description of the entity. This value must include the full name and version identification of the system's hardware type, software operating-system, and networking software.
sysLocation	The physical location of this node (for example, ‘telephone closet, 3rd floor'). If the location is unknown, the value is the zero-length string.

Table E.8 SNMPv2-MIB Supported Objects (Continued)

Object Name	Object Description in MIB
sysName	An administratively assigned name for this managed node. By convention, this is the node's fully qualified domain name. If the name is unknown, the value is the zero-length string.
sysObjectID	The vendor's authoritative identification of the network management subsystem contained in the entity. This value is allocated within the SMI enterprises subtree (1.3.6.1.4.1) and provides an easy and unambiguous means for determining `what kind of box' is being managed. For example, if vendor `Flintstones, Inc.' was assigned the subtree 1.3.6.1.4.1.4242, it could assign the identifier 1.3.6.1.4.1.4242.1.1 to its `Fred Router'.
sysServices	A value which indicates the set of services that this entity may potentially offers. The value is a sum. This sum initially takes the value zero, Then, for each layer, L, in the range 1 through 7, that this node performs transactions for, 2 raised to (L - 1) is added to the sum. For example, a node which performs only routing functions would have a value of 4 (2^(3-1)). In contrast, a node which is a host offering application services would have a value of 72 (2^(4-1) + 2^(7-1)). Note that in the context of the Internet suite of protocols, values must be calculated accordingly: layer functionality 1 physical (for example, repeaters) 2 datalink/subnetwork (for example, bridges) 3 Internet (for example, supports IP) 4 end-to-end (for example, supports TCP) 7 applications (for example, supports SMTP) For systems including OSI protocols, layers 5 and 6 may also be counted.
sysUpTime	The time (in hundredths of a second) since the network management portion of the system was last re-initialized.

APPENDIX F

SCHEMA UPDATES FOR EXTERNAL LDAP SERVERS

This section lists the Stonesoft-specific LDAP classes and attributes that you add to the schema of external LDAP servers.

The Stonesoft-specific attribute and class names start with “sg”. The classes are listed in the table below.

Table F.1 Stonesoft Specific LDAP Classes

Class	Description
sggroup	Stonesoft user group
sguser	Stonesoft user account

The Stonesoft-specific attributes are listed in the table below.

Table F.2 Stonesoft Specific LDAP Attributes

Attribute	Related Classes	Description
sgactivation	sguser	Activation date for the user account.
sgauth	sggroup, sguser	Authentication service for the user or group.
sgdelay	sggroup, sguser	Number of days the user account is valid after the activation.
sgexpiration	sguser	Last day when the user account is valid and the user can log in.
sggrouptype	sggroup	Indicates the type of the group: a subtree or discrete group.
sgmember	sggroup	The Distinguished Name (DN) for the user member of this group.

Table F.2 Stonesoft Specific LDAP Attributes (Continued)

Attribute	Related Classes	Description
sgpassword	sguser	MD5 message digest hash of the user password.
sgpresharedkey	sguser	IPsec PreSharedKey for the user account.
sgsubjectaltnames	sguser	IPsec certificate SubjectAltNames for the user account.
sgvirtualip	sggroup, sguser	Virtual IP allocation allowed for the user.

Example schema updates are provided in the Management Servers’ `<`*`installation directory`*`>/samples/LDAPSamples/LDAP/` directory:

- `SG_AD.ldif` is an example schema update for Windows Server 2003 and 2008 Active Directory. For additional considerations and instructions on extending the schema, consult Microsoft’s Documentation:
  - For Windows 2003, see http://technet.microsoft.com/en-us/library/cc759633%28WS.10%29.aspx.
  - For Windows 2008, see http://technet.microsoft.com/en-us/library/cc771796%28WS.10%29.aspx.
- `SG-v3.schema` is an example schema update in the LDAPv3 format (RFC 2252) used by OpenLDAP v.2.0.x and later, for example.
- `SG-schema.conf` is an example schema update in slapd.conf format, used by Netscape Directory server and OpenLDAP version 1.2.11, for example.

In addition to updating the directory schema, there may be some server-specific requirements. For the Netscape and the OpenLDAP version 1.2.11 servers, you must configure the following lines to the LDAP server’s `slapd.conf` configuration file after stopping the LDAP service:

Illustration F.1 Additional Configuration for OpenLDAP v1.2.11 and Netscape Server

```
include /etc/openldap/slapd.at.conf
include /etc/openldap/slapd.oc.conf
include /etc/openldap/sg-schema.conf
schemacheck on
```

For the OpenLDAP server versions 2.0 and later, you must configure the following lines to the LDAP server’s `slapd.conf` configuration file after stopping the LDAP service:

Illustration F.2 Additional Configuration for OpenLDAP version 2.0 or later

```
include /etc/openldap/schema/core.schema
include /etc/openldap/schema/cosine.schema
include /etc/openldap/schema/inetorgperson.schema
include /etc/openldap/schema/sg-v3.schema
```

Appendix G

Log Fields

The following sections are included:

Log Entry Fields

The following tables list the fields of the log entry table and the corresponding XML fields exported to syslog for exportable log entry fields. The rights of the administrator who views the logs and the log type(s) that the administrator has selected for viewing determine which fields are displayed.

- Non-exportable Log Entry Fields (page 1212).
- Exportable Alert Log Entry Fields (page 1216).
- Exportable Alert Trace Log Entry Fields (page 1216).
- Exportable Audit Log Entry Fields (page 1217).
- Exportable Firewall and Layer 2 Firewall Log Entry Fields (page 1218).
- Exportable IPS Log Entry Fields (page 1220).
- Exportable IPS Recording Log Entry Fields (page 1232).
- Exportable SSL VPN Log Entry Fields (page 1233).

Non-exportable Log Entry Fields

The following log entry fields can be displayed in the log table, but cannot be exported to syslog.

Table G.1 Non-exportable Log Entry Fields

Field	Description
Additional Situation	Identifier of an additional situation that was detected simultaneously with the situation that triggered the log event.
Blacklist response.Blacklist duration	Duration of blacklisting in seconds.
Blacklist response.Blacklist executor	Firewall or sensor that blacklisted the traffic that triggered the log event.
Blacklist response.Endpoint1 addr	Blacklisted IP addresses for Endpoint1.
Blacklist response.Endpoint1 mask	Netmask for blacklisted Endpoint1 IP address (32 = host address).
Blacklist response.Endpoint1 port	Blacklisted Endpoint1 port (empty = all ports).
Blacklist response.Endpoint1 port range	Blacklisted Endpoint1 port range.
Blacklist response.Endpoint2 addr	Blacklisted IP addresses for Endpoint2.
Blacklist response.Endpoint2 mask	Netmask for blacklisted Endpoint2 IP address (32 = host address).
Blacklist response.Endpoint2 port	Blacklisted Endpoint2 port (empty = all ports).

Table G.1 Non-exportable Log Entry Fields (Continued)

Field	Description
Blacklist response.Endpoint2 port range	Blacklisted Endpoint2 port range.
Blacklist response.Firewall ID	ID number of firewall node for which the blacklist request is assigned (this must match the Firewall ID given to the blacklist Analyzer module).
Blacklist response.IP Protocol	IP protocol of the blacklist response.
Blacklist response.Value missing in	Blacklist Response field for which value resolving failed.
Certificate verify error	TLS/SSL Certificate verify error code related to this event.
Correlation base component ID	The policy used to decide a response after successful correlation. Usually the value of this field is the same as “Component ID”, and the field is omitted.
Data type	Data type of the log.
Element Domain	Administrative Domain of the element associated with the event.
Endpoint	The VPN Endpoint through which the traffic that triggered the log event was sent or received.
Ethernet main type	Ethernet frame main type (Ethernet 2, IPX, LLC, SNAP).
Event type	Description of the event triggered the log creation.
GRE protocol	Protocol number of the GRE payload packet.
GRE version	Version of the GRE header.
IP frag conflict range.IP frag different bytes	Total number of conflicting bytes.
IP frag conflict range.IP frag different bytes first	First conflicting byte in the IP fragment.
IP frag conflict range.IP frag different bytes last	Last conflicting byte in the IP fragment.
IP frag conflict range.IP frag different new first	Value of the first conflicting byte in the latest fragment.
IP frag conflict range.IP frag different new last	Value of the last conflicting byte in the latest fragment.
IP frag conflict range.IP frag different old first	Value of the first conflicting byte in an earlier fragment.
IP frag conflict range.IP frag different old last	Value of the last conflicting byte in an earlier fragment.

Table G.1 Non-exportable Log Entry Fields (Continued)

Field	Description
IPv6 extension header type	IPv6 extension header type as indicated by the next header value of the preceding header.
IPv6 extension header’s length	IPv6 extension header length as indicated by the value of the hdr_ext_len field in the extension header.
IPv6 hop limit	Hop limit field in the IPv6 header.
IPv6 option data length	IPv6 option data length.
IPv6 option offset	IPv6 option offset from the beginning of the IPv6 extension header.
IPv6 option type	IPv6 option type.
IPv6 routing final destination	Final destination address in the IPv6 routing header.
IPv6 routing header type	IPv6 routing header type.
IPv6 routing segments left	Segments left value in the IPv6 routing header.
LLC DSAP	Logical Link Control Destination Service Access Point.
LLC SSAP	Logical Link Control Source Service Access Point.
Login Domain	The administrative Domain in which the action that triggered the log event was taken.
Normalized	URI normalization was used to find the match.
Overview	Observed overview.
Overview Name	Name of the observed overview.
Overview Section	Summary of the observed section definition.
Reference event ID.Ref Comp Id	Sender identifier of the referred event.
Reference event ID.Ref Creation Time	Creation time of the referred event.
Reference event ID.Ref Event ID	Identifier of the referred event.
Roles	Roles of the Administrator who triggered the event.
Security Gateway	The VPN Security Gateway through which the traffic that triggered the log event was sent or received.
Sender Domain	Administrative Domain from which the log entry was sent.
Sender module version.Sender build	Build number of the engine that generated the event.
Sender module version.Sender module major	Major version of the engine module that generated the event.

Table G.1 Non-exportable Log Entry Fields (Continued)

Field	Description
Sender module version.Sender module minor	Minor version of the engine module that generated the event.
Sender module version.Sender module pl	Patch version of the engine module that generated the event.
SNAP Organization Code	Subnetwork Access Protocol Organization Code.
SSL/TLS Domain	Domain name field in SSL/TLS certificate related to the event.
State	Connection state in connection monitoring.
Subexpression Count	The number of concurrent independent subexpressions.
TCP urgent pointer	Urgent pointer value in the TCP header.
TCP window size	TCP receive window size.
TCP window shrinkage	The amount by which the TCP window shrunk.
Threshold Check Time	Threshold measurement end time.
Threshold Description	Description of threshold limitation.
Threshold Measured Value	Value exceeding the threshold.
TLS Alert Description	TLS/SSL alert message description.
TLS Alert Level	TLS/SSL alert message alert level.
TLS cipher suite	TLS/SSL cipher suite.
TLS compression method	TLS/SSL compression method.
TLS Protocol Version	TLS/SSL protocol version.
Tunneling level	Number of tunneling protocol layers encapsulating this protocol layer.
User and Group Information	User and Group Information related to the event.
Virus Identifier	Virus Identifier.
VPN	The VPN through which the traffic that triggered the log event was sent or received.

Exportable Alert Log Entry Fields

Table G.2 Alert Log Entry Fields

Field	Syslog Export Field	Description
Acknowledged	ACK	Acknowledged alert.
Alert Type	ALERT	Type of alert.
Component ID	COMP_ID	The identifier of the creator of the log entry.
Creation Time	TIMESTAMP	Log entry creation time.
Data Identifier	LOG_ID	Data Identifier of the log entry.
Description	LONG_MSG	Long field description of the alert.
Dst Addr	DST	Packet destination IP address.
Event ID	EVENT_ID	Event identifier, unique within one sender.
Information message	INFO_MSG	A description of the log event that further explains the entry.
Protocol	PROTOCOL	Connection IP protocol.
Reception time	RECEPTION_TIME	Time when the entry was received by the Log Server.
Reference event ID	REF_EVENT	Reference to a related event.
Rule Tag	RULE_ID	Rule tag of the rule that triggered the log event.
Sender	NODE_ID	IP address of the engine or server that sent the log entry.
Sender type	SENDER_TYPE	The type of engine or server that sent the log entry.
Severity	ALERT_SEVERITY	Severity of the situation related to the alert event.
Situation	SITUATION	The identifier of the situation that triggered the log event.
Situation Type	SIT_CATEGORY	The type of the situation that triggered the log event.

Exportable Alert Trace Log Entry Fields

Table G.3 Alert Trace Log Entry Fields

Field	Syslog Export Field	Description
Address	EVENT_ADDRESS	Destination for the alert notification.
Alert Event	EVENT_TYPE	Type of alert event.
Alert Identifier	EVENT_LOG_ID	Data Identifier of the alert.
Alert Time	EVENT_TIME	Timestamp of the alert.

Table G.3 Alert Trace Log Entry Fields (Continued)

Field	Syslog Export Field	Description
Component ID	COMP_ID	The identifier of the creator of the log entry.
Creation Time	TIMESTAMP	Log entry creation time.
Data Identifier	LOG_ID	Data Identifier of the log entry.
Event description	EVENT_INFO	Description of the alert event.
Storage Server	STORAGE_SERVER_ID	Server where the alert is stored.
User	EVENT_USER	User who executed the action that produced the alert.

Exportable Audit Log Entry Fields

Table G.4 Audit Log Entry Fields

Field	Syslog Export Field	Description
Administrator	USER_ORIGINATOR	Administrator who triggered the audit event.
Client IP address	CLIENT_IP_ADDRESS	Address of the client that triggered the audit event.
Component ID	COMP_ID	The identifier of the creator of the log entry.
Creation Time	TIMESTAMP	Log entry creation time.
Elements	OBJECT_NAME	Elements being manipulated in the audit event.
Event ID	EVENT_ID	Event identifier, unique within one sender.
Incident case	INCIDENT_CASE	The Incident case to which the logs and/or audit events are related.
Information message	INFO_MSG	A description of the log event that further explains the entry.
Operation type	TYPE_DESCRIPTION	Type of action that triggered the audit entry.
Origin name	ORIGIN_NAME	Name of the component that triggered the audit event.
Result	RESULT	Result state after the audited event.
Sender	NODE_ID	IP address of the engine or server that sent the log entry.
Sender type	SENDER_TYPE	The type of engine or server that sent the log entry.

Exportable Firewall and Layer 2 Firewall Log Entry Fields

Table G.5 Firewall Log Entry Fields

Field	Syslog Export Field	Description
Acknowledged	ACK	Acknowledged alert.
Action	ACTION	Action of the rule that triggered the log event. The action values are Allow, Discard, Refuse, Terminate, Wait for further actions, and Wait for authentication. For more information on action values, see the table Table G.12.
Alert Type	ALERT	Type of alert.
Auth. Rule Tag	AUTH_RULE_ID	Rule number of the rule that triggered the log event.
Auth. User	AUTH_NAME	Username of the authorized user related to this event.
Bytes Rcvd	ACC_RX_BYTES	Number of bytes received during the connection.
Bytes Sent	ACC_TX_BYTES	Number of bytes sent during the connection. The number of bytes sent is counted when accounting entries are created.
Component ID	COMP_ID	The identifier of the creator of the log entry.
Creation Time	TIMESTAMP	Log entry creation time.
Data Identifier	LOG_ID	Data Identifier of the log entry.
DSCP Mark	DSCP_MARK	The DSCP Mark associated with the traffic that triggered the log event.
Dst Addr	DST	Packet destination IP address.
Dst Port	DPORT	TCP or UDP destination port in the packet header.
Elapsed Time	ACC_ELAPSED	Elapsed time of the connection in seconds. The elapsed time is recorded when accounting entries are created at the time of connection closing.
Event	EVENT	The event that triggered the log creation, for example, New connection, Connection closed, Connection discarded. For more information on event values, see the table Table G.12.
Event ID	EVENT_ID	Event identifier, unique within one sender.
Facility	FACILITY	Firewall subsystem that generated the log event. For more information on facility values, see the table Table G.9.
FP situation	FP_SITUATION	Situation identifier of a matching fingerprint.
ICMP code	ICMP_CODE	ICMP code field. ICMP code provides further information about message type (for example, network unreachable). For more information, refer to *RFC 792* and *RFC 950.*

Table G.5 Firewall Log Entry Fields (Continued)

Field	Syslog Export Field	Description
ICMP ID	ICMP_ID	The ICMP identifier recorded by the engine when ICMP packets pass through the firewall. The ICMP identifier may be used by the echo sender to aid in matching the replies with the echo requests. For example, the identifier might be used like a port in TCP or UDP to identify a session. For more information on ICMP ID and the ICMP protocol, refer to *RFC 792* and *RFC 950.*
IKE Cookie	IKE_COOKIE	IKE Cookie used in the VPN negotiation.
Information message	INFO_MSG	A description of the log event that further explains the entry.
IPsec SPI	IPSEC_SPI	The IPsec Security Parameter Index (SPI) is the connection identifier of the IPsec connection. The IPsec SPI value is displayed as a hexadecimal number.
NAT Dst	NAT_DST	Translated packet destination IP address.
NAT Dst Port	NAT_DPORT	Translated packet destination protocol port.
Nat Rule Tag	NAT_RULE_ID	The rule number of the NAT rule that triggered the log event.
NAT Src	NAT_SRC	Translated packet source IP address.
NAT Src Port	NAT_SPORT	Translated packet source protocol port.
Priority	QOS_PRIORITY	The priority assigned to the traffic according to the QoS policy.
Protocol	PROTOCOL	Connection IP protocol.
Protocol Agent	SRVHELPER_ID	Protocol Agent numerical ID code.
QoS Class	QOS_CLASS	The Quality of Service class assigned to the traffic according to the QoS policy.
Reception time	RECEPTION_TIME	Time when the entry was received by the Log Server.
Reference event ID	REF_EVENT	Reference to a related event.
Round trip	RTT	Round trip time for outbound Multi-Link link testing. Time indicated is from sending queries to the first reply. The unit is 0.01 seconds.
Rule Tag	RULE_ID	Rule tag of the rule that triggered the log event.
Sender	NODE_ID	IP address of the engine or server that sent the log entry.
Sender type	SENDER_TYPE	The type of engine or server that sent the log entry.
Service	SERVICE	Special field for filtering logs using the defined services. Does not appear in the log entry table.

Table G.5 Firewall Log Entry Fields (Continued)

Field	Syslog Export Field	Description
Severity	ALERT_SEVERITY	Severity of the situation related to the log event.
Situation	SITUATION	The identifier of the situation that triggered the log event.
Situation Type	SIT_CATEGORY	The type of the situation that triggered the log event.
Src Addr	SRC	Packet source IP address.
Src IF	Srcif	Defined source interface number for the firewall cluster.
Src Port	SPORT	TCP or UDP source port in the packet header.
Src VLAN	SRC_VLAN	The source VLAN ID number (up to 4095).
Syslog	SYSLOG_TYPE	Syslog is a system service used in some operating systems, for example, UNIX- and software packages. For more information on syslog and syslog types, refer to *RFC 3164.*
Type	TYPE	Log entry severity type. For more information on type values, see the table Table G.10.

Exportable IPS Log Entry Fields

Table G.6 IPS Log Entry Fields

Field	Syslog Export Field	Description
Agent mem usage	AGENT_MEMUSAGE	Memory usage of each IPS agent.
Alert Type	ALERT	Type of alert.
Attacker IP	IP_ATTACKER	IPv4 address of the attacking host.
Blacklist executor	FIREWALL_ID	Firewall that blacklisted the traffic that triggered the log event.
Blacklist response	BLACKLIST_RESPONSE	Firewall blacklist response that triggered the log event.
Cluster ID	CLUSTER_ID	The identifier of the cluster to which the node that created the log entry belongs.
Component ID	COMP_ID	The identifier of the creator of the log entry.
Connection analysis end	CONNECTION_ANALYSIS_END	The application could not continue analyzing the traffic stream after this event.
Connection dropped	DROP_CONNECTION	The connection was dropped by a Drop Response in the rule.
Content type of message body	SIP_CONTENT_TYPE	Content type of the SIP message body.

Table G.6 IPS Log Entry Fields (Continued)

Field	Syslog Export Field	Description
Correlation base component ID	CORRELATION_COMP_ID	The policy that decides the response after successful correlation.
Correlation begin time	TIME_FRAME_BEGIN	NTP stamp of the beginning of the time frame for a match to a correlation situation.
Correlation end time	TIME_FRAME_END	NTP stamp of the end of the time frame for a match to a correlation situation.
Creation Time	TIMESTAMP	Log entry creation time.
Data Identifier	LOG_ID	Data Identifier of the log entry.
Datagram dropped	DROP_DATAGRAM	The datagram was dropped by a Drop Response in the rule.
Description	LONG_MSG	Long field description of the alert.
Destination port	PORT_DEST	TCP or UDP destination port in the packet header. Included only for backwards compatibility with legacy IPS engines. For other cases, use Dst Port.
DNS class	DNS_CLASS	DNS resource record class.
DNS hdr ancount	DNS_HDR_ANCOUNT	DNS answers count.
DNS hdr arcount	DNS_HDR_ARCOUNT	DNS additional section count.
DNS hdr flag tc	DNS_HDR_FLAG_TC	DNS header flag TC.
DNS hdr id	DNS_HDR_ID	DNS message ID.
DNS hdr is request	DNS_HDR_IS_REQUEST	DNS message is a request.
DNS hdr nscount	DNS_HDR_NSCOUNT	DNS authority section count.
DNS hdr opcode	DNS_HDR_OPCODE	DNS operation code.
DNS hdr qdcount	DNS_HDR_QDCOUNT	DNS questions count.
DNS hdr rcode	DNS_HDR_RCODE	DNS return code.
DNS name length	DNS_NAME_LENGTH	Length of DNS name in a message.
DNS offset	DNS_OFFSET	DNS message offset where the situation occurs.
DNS pointer	DNS_POINTER	Name pointer in a DNS message.
DNS qclass	DNS_QCLASS	Query resource record class in a DNS message.
DNS qname	DNS_QNAME	First queried name in a DNS message.
DNS qtype	DNS_QTYPE	Query type in a DNS message.
DNS section	DNS_SECTION	Section name in a DNS message.

Table G.6 IPS Log Entry Fields (Continued)

Field	Syslog Export Field	Description
DNS type	DNS_TYPE	DNS resource record type.
DNS UDP payload	DNS_UDP_PAYLOAD	UDP payload size of a DNS message.
DNS UDP payload by opt	DNS_UDP_PAYLOAD_BY_OPT	UDP payload advertised in a DNS OPT record.
Dst Addr	DST	Packet destination IP address.
Dst Port	DPORT	TCP or UDP destination port in the packet header.
Error id	ERROR_ID	Identifier of the error that triggered the log event.
Eth frame length	ETH_FRAME_LENGTH	Length of the Ethernet frame.
Eth min frame length	ETH_MIN_FRAME_LENGTH	Minimum length for Ethernet frame.
Ethernet type	ETH_TYPE	Type field in Ethernet frame.
Event count	EVENT_COUNT	Event count in the defined time frame.
Event ID	EVENT_ID	Event identifier, unique within one sender.
Event update	EVENT_UPDATE	Event ID for which this event is an update.
Excerpt data	EXCERPT	Short recording of the application level data stream of the attack.
Excerpt position	EXCERPT_POS	Position in the attached short recording.
Failed response cnt	FAILED_RESP_CNT	Number of failed response attempts.
Fields updatable	FIELDS_UPDATABLE	Map of updatable log fields.
FP situation	FP_SITUATION	Situation identifier of a matching fingerprint.
Frame dropped	DROP_FRAME	The frame was dropped by a Drop Response in the rule.
From address	SIP_FROM	SIP From address.
FTP account len	FTP_ACCOUNT_LEN	Length of the FTP account string.
FTP adat argument len	FTP_ADAT_ARG_LEN	Length of ADAT command argument.
FTP allocate size	FTP_ALLOCATE_SIZE	Size of FTP allocate.
FTP arg len	FTP_ARG_LEN	Length of the FTP command argument.
FTP auth arg len	FTP_AUTH_ARG_LEN	Length of the AUTH argument.
FTP client state name	FTP_CLIENT_STATE_NAME	The detected FTP client state.

Table G.6 IPS Log Entry Fields (Continued)

Field	Syslog Export Field	Description
FTP clnt arg len	FTP_CLNT_ARG_LEN	Length of the FTP CLNT argument.
FTP command	FTP_COMMAND	Name of the FTP command.
FTP conf arg len	FTP_CONF_ARG_LEN	Length of the CONF command argument.
FTP enc arg len	FTP_ENC_ARG_LEN	Length of the ENC command argument.
FTP eprt arg len	FTP_EPRT_ARG_LEN	Length of the EPRT command argument.
FTP estp arg len	FTP_ESTP_ARG_LEN	Length of the ESTP command argument.
FTP help arg len	FTP_HELP_ARG_LEN	Length of the HELP command argument.
FTP lang arg len	FTP_LANG_ARG_LEN	Length of the LANG command argument.
FTP lprt arg len	FTP_LPRT_ARG_LEN	Length of the LPRT command argument.
FTP marker len	FTP_MARKER_LEN	Length of the REST command argument.
FTP mic arg len	FTP_MIC_ARG_LEN	Length of the MIC command argument.
FTP opts arg len	FTP_OPTS_ARG_LEN	Length of the OPTS command argument.
FTP password len	FTP_PASSWORD_LEN	Length of the detected FTP password.
FTP pathname len	FTP_PATHNAME_LEN	Length of the detected FTP pathname.
FTP protection buffer size	FTP_PROTECTION_BUFFER_SIZE	Size of the detected PBSZ protection buffer.
FTP reply	FTP_REPLY	The detected FTP server reply.
FTP reply code	FTP_REPLY_CODE	The detected FTP server reply code.
FTP reply len	FTP_REPLY_LEN	Length of an FTP server reply that is too long.
FTP reply line len	FTP_REPLY_LINE_LEN	Length of an FTP server reply line that is too long.
FTP server action	FTP_SERVER_ACTION	FTP server action after a suspicious client command.
FTP server banner	FTP_SERVER_BANNER	The detected FTP server banner.
FTP server state name	FTP_SERVER_STATE_NAME	The detected FTP server state.
FTP site arg len	FTP_SITE_ARG_LEN	Length of the SITE command argument.
FTP state name	FTP_STATE_NAME	The detected FTP session state.
FTP username len	FTP_USERNAME_LEN	Length of the detected FTP username.
HTTP content length	HTTP_CONTENT_LENGTH	HTTP content length.
HTTP content type	HTTP_CONTENT_TYPE	HTTP content type.

Table G.6 IPS Log Entry Fields (Continued)

Field	Syslog Export Field	Description
HTTP header	HTTP_HEADER	The detected HTTP header field.
HTTP header name	HTTP_HEADER_NAME	The detected HTTP header field name.
HTTP no request	HTTP_NO_REQUEST	The detected HTTP response could not be associated to any request.
HTTP request host	HTTP_REQUEST_HOST	HTTP request host.
HTTP request line	HTTP_REQUEST_LINE	The detected HTTP request line.
HTTP request message field name length	HTTP_REQUEST_MESSAGE_ FIELD_NAME_LENGTH	Length of the HTTP request header field name.
HTTP request message field value length	HTTP_REQUEST_MESSAGE_ FIELD_VALUE_LENGTH	Length of the HTTP request header field value.
HTTP request method	HTTP_REQUEST_METHOD	The detected HTTP request method.
HTTP request URI	HTTP_REQUEST_URI	The detected HTTP request URI.
HTTP request version	HTTP_REQUEST_VERSION	The detected HTTP request version.
HTTP requests not stored	HTTP_REQUESTS_NOT_ STORED	Number of requests not stored due to HTTP pipeline overflow.
HTTP response code	HTTP_RESPONSE_CODE	The detected HTTP response code.
HTTP response message field name length	HTTP_RESPONSE_MESSAGE_ FIELD_NAME_LENGTH	Length of the HTTP response header field name.
HTTP response message field value length	HTTP_RESPONSE_MESSAGE_ FIELD_VALUE_LENGTH	Length of the HTTP response header field value.
HTTP URI length	HTTP_URI_LENGTH	Length of HTTP request URI
ICMP code	ICMP_CODE	ICMP code field. ICMP code provides further information about message type (for example, network unreachable). For more information, refer to *RFC 792* and *RFC 950.*
ICMP expected message length	ICMP_EXPECTED_MESSAGE_ LENGTH	Expected length of the ICMP message.
ICMP field addr entry size	ICMP_FIELD_ADDR_ENTRY_SI ZE	Value of the detected ICMP address entry size field.

Table G.6 IPS Log Entry Fields (Continued)

Field	Syslog Export Field	Description
ICMP field address mask	ICMP_FIELD_ADDRESS_MASK	Value of detected ICMP address mask field.
ICMP field domain name	ICMP_FIELD_DOMAIN_NAME	Value of the detected ICMP domain name field.
ICMP field gateway IP addr	ICMP_FIELD_GATEWAY_IP_ADDR	Value of the detected ICMP gateway address field.
ICMP field lifetime	ICMP_FIELD_LIFETIME	Value of the ICMP lifetime field.
ICMP field num addrs	ICMP_FIELD_NUM_ADDRS	Value of the ICMP number of addresses field.
ICMP field originate timestamp	ICMP_FIELD_ORIGINATE_TIMESTAMP	Value of the ICMP originate timestamp field.
ICMP field outbound hop count	ICMP_FIELD_OUTBOUND_HOP_COUNT	Value of the ICMP outbound hop count field.
ICMP field output link mtu	ICMP_FIELD_OUTPUT_LINK_MTU	Value of the ICMP output link MTU field.
ICMP field output link speed	ICMP_FIELD_OUTPUT_LINK_SPEED	Value of the ICMP output link speed field.
ICMP field pointer	ICMP_FIELD_POINTER	The offset in the related datagram where the situation occurred.
ICMP field preference level	ICMP_FIELD_PREFERENCE_LEVEL	Value of the ICMP preference level field.
ICMP field receive timestamp	ICMP_FIELD_RECEIVE_TIMESTAMP	Value of the ICMP receive timestamp field.
ICMP field return hop count	ICMP_FIELD_RETURN_HOP_COUNT	Value of the ICMP return hop count field.
ICMP field router addr	ICMP_FIELD_ROUTER_ADDRESS	Value of the ICMP router address field.
ICMP field sequence num	ICMP_FIELD_SEQUENCE_NUMBER	Value of the ICMP sequence number field.
ICMP field traceroute id	ICMP_FIELD_TRACEROUTE_ID	Value of the ICMP traceroute ID field.
ICMP field transmit timestamp	ICMP_FIELD_TRANSMIT_TIMESTAMP	Value of the ICMP transmit timestamp field.

Table G.6 IPS Log Entry Fields (Continued)

Field	Syslog Export Field	Description
ICMP ID	ICMP_ID	The ICMP identifier recorded by the engine when ICMP packets pass through the firewall. The ICMP identifier may be used by the echo sender to aid in matching the replies with the echo requests. For example, the identifier might be used like a port in TCP or UDP to identify a session. For more information on ICMP ID and the ICMP protocol, refer to *RFC 792* and *RFC 950.*
ICMP message length	ICMP_MESSAGE_LENGTH	Length of the ICMP message.
ICMP referenced destination IP addr	ICMP_REFERENCED_ DESTINATION_IP_ADDR	Destination IP address of the datagram related to the ICMP message.
ICMP referenced destination port	ICMP_REFERENCED_ DESTINATION_PORT	Destination port of the datagram related to the ICMP message.
ICMP referenced IP proto	ICMP_REFERENCED_IP_ PROTO	IP Protocol field of the datagram related to the ICMP message.
ICMP referenced source IP addr	ICMP_REFERENCED_SOURCE _IP_ADDR	Source IP address of the datagram related to the ICMP message.
ICMP referenced source port	ICMP_REFERENCED_SOURCE _PORT	Source port of IP datagram related to the ICMP message.
ICMP Type	ICMP_TYPE	The Internet Control Message Protocol is an extension to the Internet Protocol (IP) that supports packets containing error, control and informational messages. ICMP messages are sent using the basic IP header. The first octet of the data portion of the datagram is an ICMP *type* field. For more information, refer to *RFC 792* and *RFC 950.*
Imf encoded word	IMF_ENCODED_WORD	Encoded word token related to this event.
Imf header field	IMF_HEADER_FIELD	Contents (possibly partial) of the mail header field related to this event.
Imf header field name	IMF_HEADER_FIELD_NAME	Name of the mail header field related to this event.
Imf header field position	IMF_HEADER_FIELD_ POSITION	Number of characters processed in this header field when this event was generated.
Imf token	IMF_TOKEN	Syntactical token in the mail body related to this event.
Imf token length	IMF_TOKEN_LENGTH	Length of the syntactical token in the mail body related to this event.
Information message	INFO_MSG	A description of the log event that further explains the entry.

Table G.6 IPS Log Entry Fields (Continued)

Field	Syslog Export Field	Description
Inspection check parameter	MODPAR_VA	List of agent parameters and the defined values.
IP checksum	IP_CHECKSUM	Value of the IP header checksum.
IP datagram length	IP_DATAGRAM_LENGTH	Length of the IP datagram.
IP datagram new length	IP_DATAGRAM_NEW_LENGTH	The new suggested length for the IP datagram.
IP destination	IP_DEST	Destination IP address in the packet header. Included only for backwards compatibility for legacy IPS. For other cases, use Dst Addr.
IP frag conflict range	IP_FRAG_CONFLICT_RANGE AGMENT_OFFSET	Conflicting byte range in a fragment.
IP fragment offset	IP_FRAGMENT_OFFSET	Fragment offset in the IP header.
IP header length	IP_HEADER_LENGTH	Length of the IP header.
IP identification	IP_IDENTIFICATION	Identification field in the IP header.
IP offset	IP_OFFSET	Start IP offset from the beginning of the Ethernet frame.
IP option length	IP_OPTION_LENGTH	Length of the IP option that triggered the response.
IP option number	IP_OPTION_NUMBER	IP option number that triggered the response.
IP protocol	PROTOCOL	IP protocol of the traffic that generated the log event.
IP source	IP_SOURCE	Source IP address in the packet header. Included for backwards compatibility with legacy IPS. For other cases, use Src Addr.
IP total length	IP_TOTAL_LENGTH	Total length of the IP datagram.
IP version	IP_VERSION	Version field value in the IP header.
Length of message body	SIP_CONTENT_LENGTH	Length of the SIP message body.
Logical interface	IF_LOGICAL	Logical interface for a packet.
MAC destination	MAC_DEST	Destination MAC address in the packet header.
MAC source	MAC_SOURCE	Source MAC address in the packet header.
Module	SENDER_MODULE_ID	Sender module identification.
Module mem usage	MODULE_MEMUSAGE	Memory usage of each module.
Node configuration	NODE_CONFIGURATION	Current configuration of the node that sent the log entry.

Table G.6 IPS Log Entry Fields (Continued)

Field	Syslog Export Field	Description
Node dynup	NODE_DYNUP	Dynamic update package level of the node that sent the log entry.
Node version	NODE_VERSION	Node version of the node that sent the log entry.
Not final value	NOT_FINAL_VALUE	Entry is not final.
One LAN	ONE_LAN	The “View interface as one LAN” option was enabled on the logical interface through which the packet was received.
Orig config id	ORIG_CONFIG_ID	Configuration identifier related to the Situation in the referred event.
Orig sender module version	ORIG_SENDER_MODULE_ VERSION	Module version in the referred event.
Orig sender os ver	ORIG_SENDER_OS_VER	The operating system version of the sender of the referred event.
Original Alert Type	ORIG_ALERT	Type of alert in the referred event.
Original correlation begin time	ORIG_TIME_FRAME_BEGIN	NTP stamp of the beginning of the time frame in the referred event.
Original correlation end time	ORIG_TIME_FRAME_END	NTP stamp of the end of the time frame in the referred event.
Original event count	ORIG_EVENT_COUNT	Number of events in the time frame of the referred event.
Original module	ORIG_SENDER_MODULE_ID	Sender module identification in the referred event.
Original severity	ORIG_ALERT_SEVERITY	Severity of the referred event.
Original situation	ORIG_SITUATION	Identifier of the situation that triggered the referred event.
Original time	ORIG_TIMESTAMP	Creation time of the referred event.
Packet analysis end	PACKET_ANALYSIS_END	Module could not continue analyzing packet or datagram after this event.
Packet not seen	PACKET_NOT_SEEN	Flag indicating that the related packet was not seen.
Physical interface	IF_PHYSICAL	Physical interface for a packet.
Protocol	PROTOCOL	Connection IP protocol.
Protocol Agent	SRVHELPER_ID	Protocol Agent numerical ID code.
Reception time	RECEPTION_TIME	Time when the entry was received by the Log Server.
Record ID	RECORD_ID	Identifier of the traffic recording.

Table G.6 IPS Log Entry Fields (Continued)

Field	Syslog Export Field	Description
Reference event ID	REF_EVENT	Reference to a related event.
Rule Tag	RULE_ID	Rule tag of the rule that triggered the log event.
Scan ICMP echo no reply cnt	SCAN_ICMP_ECHO_NO_RESPONSE_COUNTER	Number of distinct ICMP Echo Request (ping) destinations that did not reply to a request.
Scan ICMP echo request cnt	SCAN_ICMP_ECHO_REQUEST_COUNTER	Number of ICMP Echo Request destinations detected.
Scan ICMP echo targets	SCAN_ICMP_ECHO_TARGETS	List of the detected ICMP Echo Request destinations.
Scan ICMP mask no reply cnt	SCAN_ICMP_NETMASK_NO_RESPONSE_COUNTER	Number of ICMP Netmask Request destinations with no reply.
Scan ICMP mask request cnt	SCAN_ICMP_NETMASK_REQUEST_COUNTER	Number of distinct ICMP Netmask Request destinations detected.
Scan ICMP mask targets	SCAN_ICMP_NETMASK_TARGETS	List of the detected ICMP Netmask Request destinations.
Scan ICMP no reply cnt	SCAN_ICMP_NO_RESPONSE_COUNTER	Number of ICMP Echo, Timestamp, and Netmask Request destinations with no reply.
Scan ICMP request cnt	SCAN_ICMP_REQUEST_COUNTER	Number of ICMP Echo, Timestamp, and Netmask Request destinations.
Scan ICMP time no reply cnt	SCAN_ICMP_TIMESTAMP_NO_RESPONSE_COUNTER	Number of ICMP Timestamp Request destinations with no reply.
Scan ICMP time request cnt	SCAN_ICMP_TIMESTAMP_REQUEST_COUNTER	Number of distinct ICMP Timestamp Request destinations detected.
Scan ICMP time targets	SCAN_ICMP_TIMESTAMP_TARGETS	List of detected ICMP Timestamp Request destinations.
Scan start time	SCAN_START_TIME	Detected starting time of the port scanning activity that triggered the log event.
Scan TCP negative cnt	SCAN_TCP_NEGATIVE_COUNTER	Number of TCP destinations that replied with TCP Reset.
Scan TCP no ack cnt	SCAN_TCP_NO_ACK_COUNTER	Number of TCP destinations targeted for illegal TCP segments.
Scan TCP no ack targets	SCAN_TCP_NO_ACK_TARGETS	List of TCP destinations targeted for illegal TCP segments.
Scan TCP no reply cnt	SCAN_TCP_NO_RESPONSE_COUNTER	Number of TCP destinations with no reply to connection attempts.

Table G.6 IPS Log Entry Fields (Continued)

Field	Syslog Export Field	Description
Scan TCP normal cnt	SCAN_TCP_NORMAL_COUNTER	Number of TCP destinations with handshake and two-directional data transfer.
Scan TCP positive cnt	SCAN_TCP_POSITIVE_COUNTER	Number of TCP destinations with handshake but no data sent by client.
Scan TCP targets	SCAN_TCP_TARGETS	List of detected TCP port scan destinations in the traffic that triggered the log event.
Scan UDP negative cnt	SCAN_UDP_NEGATIVE_COUNTER	Number of destinations that replied with ICMP Port Unreachable.
Scan UDP positive cnt	SCAN_UDP_POSITIVE_COUNTER	Number of two-directional UDP conversations detected.
Scan UDP probe cnt	SCAN_UDP_PROBE_COUNTER	Number of destinations that did not reply using UDP.
Scan UDP target cnt	SCAN_UDP_TARGET_COUNTER	Total number of UDP destinations detected.
Scan UDP targets	SCAN_UDP_TARGETS	List of detected UDP destinations.
Sender	NODE_ID	IP address of the engine or server that sent the log entry.
Sender module version	SENDER_MODULE_VERSION	Version of the engine module that generated the event.
Sender type	SENDER_TYPE	The type of engine or server that sent the log entry.
Service	SERVICE	Special field for filtering logs using the defined services. Does not appear in the log entry table.
Severity	ALERT_SEVERITY	Severity of the situation related to the alert event.
SIP call ID	SIP_CALL_ID	SIP call ID.
SIP contact address	SIP_CONTACT	SIP contact address.
SIP header field contents	SIP_HEADER	SIP header field contents.
SIP header field name	SIP_HEADER_NAME	SIP header field name.
SIP request method	SIP_REQUEST_METHOD	Method of the SIP request.
SIP request URI	SIP_REQUEST_URI	URI of the SIP request.
SIP request version	SIP_REQUEST_VERSION	Version of the SIP request.

Table G.6 IPS Log Entry Fields (Continued)

Field	Syslog Export Field	Description
SIP response reason-phrase	SIP_RESPONSE_REASON_ PHRASE	SIP response reason-phrase.
SIP response status code	SIP_RESPONSE_STATUS_ CODE	Status code of the SIP response.
SIP VIA address	SIP_VIA	SIP VIA address.
Situation	SITUATION	The identifier of the situation that triggered the log event.
Situation Type	SIT_CATEGORY	The type of the situation that triggered the log event.
SMTP command	SMTP_COMMAND	Suspicious SMTP command sent by the client.
SMTP mail stats	SMTP_MAIL_STATS	Statistics on e-mail messages.
SMTP misplaced command	SMTP_MISPLACED_ COMMAND	Command given in the wrong place in the command sequence.
SMTP recipient	SMTP_RECIPIENT	Recipient forward path in RCPT command parameter.
SMTP reply	SMTP_REPLY	Suspicious SMTP reply message sent by the server.
SMTP reverse path	SMTP_REVERSE_PATH	SMTP reverse path in MAIL FROM command parameter.
SMTP server action	SMTP_SERVER_ACTION	Suspicious server action after a suspicious client command.
SMTP server banner	SMTP_SERVER_BANNER	Banner sent by the SMTP server at the beginning of the connection.
SMTP transaction state	SMTP_TRANSACTION_STATE	Session state of the SMTP transaction.
Source file	SOURCE_FILE	Name of the source file.
Source file line	SOURCE_FILE_LINE)	Line number in the source file.
Source port	PORT_SOURCE	TCP or UDP source port in the packet header. Included for backwards compatibility with legacy IPS. For other cases, see Src Port.
Src Addr	SRC	Packet source IP address.
Src Port	SPOR	TCP or UDP source port in the packet header.
SSH calc client crypto bit ratio	SSH_CALC_CLIENT_CRYPTO_ BIT_RATIO	Calculated SSH client crypto bit ratio.
SSH calc server crypto bit ratio	SSH_CALC_SERVER_CRYPTO _BIT_RATIO	Calculated SSH server crypto bit ratio.
SSH1 host key bits	SSH1_HOST_KEY_BITS	Bit length of the SSHv1 host key.

Table G.6 IPS Log Entry Fields (Continued)

Field	Syslog Export Field	Description
SSH1 server key bits	SSH1_SERVER_KEY_BITS	Bit length of the SSHv1 server key.
Syslog	SYSLOG_TYPE	Syslog is a system service used in some operating systems, for example, UNIX- and software packages. For more information on syslog and syslog types, refer to *RFC 3164*.
Target IP	IP_TARGET	IPv4 address of the target host in a detected attack.
TCP connection start time	TCP_CONNECTION_START_TIME	Start time of the TCP connection.
TCP handshake seen	TCP_HANDSHAKE_SEEN	Initial handshake of the TCP connection detected.
TCP option kind	TCP_OPTION_KIND	Type of the TCP option.
TCP option length	TCP_OPTION_LENGTH	Length of the TCP option that caused the response.
To address	SIP_TO	SIP To address.
UDP datagram size	UDP_DATAGRAM_SIZE	Size of the UDP datagram.
Vulnerability References	VULNERABILITY_REFERENCES	References to known vulnerabilities in a vulnerability database. Generated from situation and original situation.
Whole session seen	WHOLE_SESSION_SEEN	True if no data of this session has been missed up to this point.

Exportable IPS Recording Log Entry Fields

Table G.7 IPS Recording Log Entry Fields

Field	Syslog Export Field	Description
Component ID	COMP_ID	The identifier of the creator of the log entry.
Creation Time	TIMESTAMP	Log entry creation time.
Packet data	PACKET_DATA	Recorded packet data.
Record frame cached	RECORD_FRAME_CACHED	Marker showing that this frame was received before the recording was started. The frame included in the recording was taken from a memory cache.
Record ID	RECORD_ID (*IPS and IPS recording only*)	Identifier of the traffic recording.
Sender	NODE_ID	IP address of the engine or server that sent the log entry.

Exportable SSL VPN Log Entry Fields

Table G.8 SSL VPN Log Entry Fields

Field	Syslog Export Field	Description
Application	APPLICATION	The SSL VPN service that generated the log.
Application Detail	APPLICATION_DETAIL	The type of log sent by the SSL VPN. Possible values: SYSTEM, AUDIT, HTTP, DEBUG, BILLING, RADIUS, EVENT.
Creation Time	TIMESTAMP	Log entry creation time.
Information message	INFO_MSG	A description of the log event that further explains the entry.
Reception time	RECEPTION_TIME	Time when the entry was received by the Log Server.
Sender	NODE_ID	IP address of the engine or server that sent the log entry.
Username	USERNAME	Username of the user to which this log event is related.
Message Id	MESSAGE_ID	SSL VPN-internal identifier of the log entry.
Session Id	SESSION_ID	ID of the User Session.
Log Severity	LOG_SEVERITY	The severity assigned to the log entry by the SSL VPN-internal logging system. Possible values: Debug, Fatal, Info, Unknown, Warning.

Facility Field Values

The following table lists the possible values for the Facility field in the log table.

Table G.9 Facility Field Values

Value
Accounting
Authentication
Blacklisting
Cluster Daemon
Cluster Protocol
Connection Tracking
Data Synchronization
DHCP Client
DHCP Relay
Invalid

Table G.9 Facility Field Values (Continued)

Value
IPsec
License
Load balancing filter
Log Server
Logging System
Management
Monitoring
NetLink Incoming HA
Network Address Translation
Packet Filter
Protocol Agent
Server Pool
SNMP Monitoring
State Synchronization
Syslog
System
Tester
Undefined
User Defined

Type Field Values

The following table lists the possible values for the Type field in the log table.

Table G.10 Type Field Values

Value
Critical Error
Debug high
Debug low
Debug mid
Diagnostic
Emergency - System Unusable
Error
Informational
Internal max
Max
Notification
System Alert
Undefined
Warning

Action Field Values

The following table show the most common log occurrences for the Action field.

Table G.11 Action Field Values

Action	Description
Allow	A connection was allowed through the engine. This can be a new connection, a related connection (for example, an FTP data connection), a related packet (for example ICMP error messages related to an earlier TCP connection), or a new connection through an existing VPN tunnel.
Discard	A connection or packet was discarded by the engine.
Permit	A connection was allowed through according to the Inspection Rules on the engine.
Refuse	A connection was refused by the engine.
Terminate	A connection was terminated by the engine.
Terminate (failed)	An attempt to terminate a connection failed.
Terminate (passive)	A connection matched a rule with the passive Terminate action, and a log entry indicating that the connection would have been terminated was produced.
Terminate (reset)	A connection was terminated by the engine and TCP resets were sent to both communicating hosts.
Wait for Authentication	A connection was waiting for successful user authentication before it could continue.
Wait for Further Actions	A connection was waiting for some other action before it could continue.
Wait for RPC Reply	A connection was waiting for an RPC reply before it could continue.

Event Field Values

The following table show the most common log occurrences for the Event field.

Table G.12 Event Field Values

Event	Description
Allowed a connection from blacklister	A connection from a blacklister was allowed.
Application protocol version is not supported	The application protocol version used in the traffic is not supported.
Application protocol version not recognized	The application protocol version used in the traffic was not recognized.
Authentication error	There was an error in the user authentication process.

Table G.12 Event Field Values (Continued)

Event	Description
Authentication failed	A user did not successfully authenticate.
Authentication Server does not respond	There is no response from the Authentication Server.
Authentication succeeded	A user successfully authenticated.
Automatic online transition	An engine automatically went online.
Automatic standby transition	An engine automatically went to standby.
Blacklister not allowed	The component that attempted to send a blacklist request is not on the list of Allowed Blacklisters.
Blacklisting connection closed	A connection from a blacklister was closed.
Blacklisting entries flushed	All entries were removed from the engine’s blacklist.
Blacklisting entry deleted	An entry was removed from the engine’s blacklist.
Blacklisting entry expired	A blacklisting entry reached the end of its Duration time.
Can't connect to log server	The engine is unable to connect to the Log Server.
Configuration changed	The engine’s configuration changed.
Configuration information for this connection	The engine’s configuration at the time the connection was logged.
Connection cannot be redirected to CIS due to absence of source NAT rule	Redirection to a Content Inspection Server failed because there was no NAT rule to redirect the connection.
Connection closed	A connection was closed.
Connection Discarded	A connection was discarded by the engine.
Connection Queued	A connection was queued according to the QoS rules.
Connection redirected to Content Inspection Server	A connection was redirected to an external Content Inspection Server.
Connection Refused	A connection was refused by the engine.
Connection Terminated	A connection was terminated by the engine.
Data connection cannot be redirected to CIS due to absence of source NAT rule	Redirection to a Content Inspection Server failed because there was no NAT rule to redirect the data connection.
Data connection redirected to content inspection server	A data connection was redirected to an external Content Inspection Server.
DHCP message received	A DHCP message was received.

Table G.12 Event Field Values (Continued)

Event	Description
DHCP Relay address not configured, reply discarded	A DHCP reply was discarded because no DCHP address is configured for the engine.
DHCP Relay address spoofed, request discarded	A DHCP request was discarded because the DHCP relay address was considered spoofed.
DHCP reply received	A DHCP reply was received.
DHCP reply sent	A DHCP reply was sent.
DHCP request forwarded	A DHCP request was forwarded.
DHCP request received	A DHCP request was received.
DHCP request sent	A DHCP request was sent.
Dropped AH packet	An IPsec AH packet was dropped.
Dropped ESP packet	An IPsec ESP packet was dropped.
Error in receiving a new configuration	There was an error when trying to transfer a new configuration to the engine.
Error with Content Inspection Server	There was an error when attempting to redirect a connection to an external Content Inspection Server.
Failed to allow a related connection to open	The engine failed to open a related connection for a connection that had already been allowed.
Force offline by test failure	The engine was forced offline as the result of an automated test failing.
Going locked offline by command	An administrator commanded the engine to go to the locked offline state.
Going locked online by command	An administrator commanded the engine to go to the locked online state.
Going offline by command	An administrator commanded the engine to go offline.
Going offline by test failure	The engine went offline as the result of an automated test failing.
Going online by command	An administrator commanded the engine to go online.
Going standby by command	An administrator commanded the engine to go to standby.
Hybrid authentication done	Hybrid authentication successfully completed.
Hybrid authentication failed	Hybrid authentication failed.
Incomplete connection closed	A connection for which the TCP handshake did not complete was closed.

Table G.12 Event Field Values (Continued)

Event	Description
Internal engine error	An internal error occurred on the engine.
Internal error	An internal error occurred.
Invalid license	The engine has an invalid license.
Invalid properties of custom Protocol Agent	Invalid options have been configured for a custom Protocol Agent.
IPsec authentication error	An error occurred in IPsec authentication.
IPsec client cfg download done	The configuration for an IPsec VPN Client has finished downloading.
IPsec client cfg download failed	An attempt to download the configuration for an IPsec VPN Client failed.
IPsec client cfg download from	The configuration for an IPsec VPN Client was downloaded by the client at the source address.
IPsec IKE error	There was an error in the IKE negotiation for an IPsec VPN.
LDAP Server does not respond	An LDAP Server is not responding.
License exceeded	A throughput based license was exceeded.
Log spool corrupted	The data in the engine’s log spool partition has become corrupted.
Log spool is becoming full	The engine’s log spool partition is becoming full.
New blacklisting entry	A new entry was added to the engine’s blacklist.
New configuration successfully installed	A new configuration was installed on the engine.
New connection	A new connection was allowed through the engine.
New VPN connection	A new connection through an existing VPN tunnel was allowed.
No space left on device	The engine’s hard drive is full.
No suitable NAT rule found	No NAT rule matched a connection.
No suitable NAT rule found for related connection	No NAT rule matched a related connection.
Node booted	An engine node booted up.
Node down	An engine node is down.
Node up	An engine node is up.
Oversized DHCP message discarded	An excessively large DHCP message was discarded.

Table G.12 Event Field Values (Continued)

Event	Description
Packet Discarded	A packet was discarded by the engine.
Packet too long	A packet was too long.
Packet too short	A packet was too short.
Receive ICMP echo	An ICMP echo (ping) was received.
Related Connection	A related connection was allowed through the engine. For example, an FTP data connection.
Related Packet	A related packet was allowed through the engine. For example, ICMP error messages related to an earlier TCP connection.
Requested NAT cannot be done	There was an error when applying NAT to the traffic.
Security Policy reload	New security policy is loaded on the engine.
Send ICMP echo	An ICMP echo (ping) was sent.
Sending DHCP reply failed	The engine failed to send a DHCP reply.
Sending DHCP request failed	The engine failed to send a DHCP request.
Sending sync messages	The engine is sending synchronization messages.
Server pool member went offline	A Server Pool member went offline.
Server pool member went online	A Server Pool member went online.
SSL Handshake failed	An SSL handshake failed.
Starting hybrid authentication	Hybrid authentication started.
Starting IKE initiator negotiation	IKE initiator negotiation started.
Starting IKE responder negotiation	IKE responder negotiation started.
State sync communication failure	State synchronization communication between cluster nodes failed.
State sync configuration changed	The configuration of the synchronization communication between cluster nodes changed.
Unknown DCHP Relay error	An unknown error occurred in DHCP relay.
Unrecognized protocol	A protocol in the logged traffic was not recognized.
Went locked offline	The engine went to the locked offline state.
Went locked online	The engine went to the locked online state.
Went offline	The engine went offline.
Went online	The engine went online.

Table G.12 Event Field Values (Continued)

Event	Description
Went standby	The engine went to standby.

A successful engine login causes an event that is displayed in the Logs view with the following type of message in the Info Message field: `date time login[id]:USERNAME LOGIN on 'device'`. A failed login causes an info message of the following type: `date time login[id]:FAILED LOGIN (#) on 'device' FOR 'UNKNOWN'`.

IPsec VPN Log Messages

The tables in this section list the most common IPsec VPN log messages (Facility=IPsec). Some messages can only be seen when the VPN diagnostics are enabled during the VPN negotiations. The messages listed appear in the Information Message fields of logs as information or error messages. The Situation field in some of the logs contains similar messages.

- VPN Notifications (page 1241)
- VPN Errors (page 1243)
- VPN Error Codes (page 1245)

VPN Notifications

The table below lists messages that are seen in the logs as part of normal IPsec VPN operation.

Table G.13 Common IPsec VPN Messages in Normal Operation

Information Message	Description
SA traffic selectors local: [...]	This message is visible only when IPsec diagnostics are enabled. The first message generated when new VPN negotiations are triggered. Negotiation of a new VPN tunnel follows.
IKE SA proposal [...]	This message is visible only when IPsec diagnostics are enabled. Shows the proposal that the initiator in the negotiations sent to the responder (displayed in both roles).
Starting IKE main mode initiator negotiation Starting IKE main mode responder negotiation	The beginning of IKE negotiations (in main mode). Which message is displayed depends on whether the gateway is the initiator or the responder in the negotiation. Repeated negotiations for the same connection are normal in a Multi-Link environment.
IKE Phase-1 initiator done [...] IKE Phase-1 responder done [...]	IKE Phase-1 negotiations were successfully completed, Phase-2 negotiations will begin. Which message is displayed depends on whether the gateway is the initiator or the responder in the negotiation.

Table G.13 Common IPsec VPN Messages in Normal Operation (Continued)

Information Message	Description
IKE Phase-2 initiator done [...] IKE Phase-2 responder done [...]	IKE Phase-2 negotiations were successfully completed. The VPN tunnel is now established and ESP or AH message(s) should appear shortly. Which message is displayed depends on whether the gateway is the initiator or the responder in the negotiation.
Starting Hybrid Authentication	Hybrid authentication is started for a Stonesoft IPsec VPN client user.
Hybrid Authentication Done	Hybrid authentication succeeded for a Stonesoft IPsec VPN client user.
IKE SA import succeeded IPsec SA import succeeded	This message is visible only when IPsec diagnostics are enabled. Synchronization of Phase 1 (IKE) and Phase 2 (IPsec) information between clustered firewall engines was successful.
ESP [...] AH [...]	Encrypted traffic going through the VPN tunnel. When you enable IPsec diagnostics you may see more of these messages.
Unknown IKE cookie	This message is visible only when IPsec diagnostics are enabled. The other gateway identified an SA that does not exist on this node. If this is a cluster, this message is normal when the SA has been negotiated with a different node and the correct SA is then queried from the other nodes, allowing the connection to continue. This message can also appear if the SA has been deleted, for example, because of a timeout or dead peer detection (DPD).
Sending delete notification [...] Delete notification received [...]	This message is visible only when IPsec diagnostics are enabled. Messages between the gateways forming the tunnel informing the other party that the gateway has removed the settings indicated in the message. As a result, the other gateway also clears the settings, allowing for renegotiations if the tunnel is still needed.
Sending IKE SA delete sync Receiving IKE SA expire/delete sync	This message is visible only when IPsec diagnostics are enabled. Synchronization of SA deletion information between clustered firewall engines.
Initial contact notification received	The gateway at the other end of the tunnel has sent an Initial-Contact message (indicating that it has no knowledge of previous negotiations). If there are old SAs with the gateway, they are deleted at this point (recently negotiated SAs are not, as may be indicated by a further log message). If SAs exist, the notification may indicate that the other end has been cleared, for example, in a reboot.

VPN Errors

The table below lists common errors that indicate problems in an IPsec VPN tunnel. The log messages inform you about the stage of negotiations and then give the actual error message, for example, “IKE Phase-2 error: No proposal chosen”. The table lists only the actual message part without additional variable details such as IP addresses or identifiers.

Table G.14 Common IPsec VPN Errors

Error Message	Description
Access group mismatch	The connecting VPN client is not authorized.
Authentication failed	One of the parties rejected the authentication credentials or something went wrong during the authentication process. If the problem is not apparent in the available logs, activate diagnostics to generate more verbose logs that give you more information about the next negotiations.
Authentication method mismatch	The authentication method used by the other gateway is not allowed in the configuration of this gateway. Check the settings in the VPN Profile that is selected for this VPN.
Can not get policy [...] No matching connection	May indicate that the gateway has no valid VPN certificate.
Can not get QM policy [...]	Indicates that there is a mismatch in granularity settings between the negotiating gateways. In Stonesoft, granularity is controlled with the Security Association Granularity setting on the IPsec Settings tab of the VPN Profile.
Could not allocate inbound SPI	Indications that the gateway has run out of memory. The reasons for this may include inappropriate configuration settings (such as using the “SA per host” setting with a very large number of hosts) in addition to other considerations (such as hardware specifications).
Could not create outbound IPsec rule	
Could not register outbound SPI	
Old outbound SPI entry not found	
Out of memory	
SA install failed	
Session attaching failed	
Transform creation failed	
Dead peer detection failed IKE peer was found dead [...]	Dead peer detection checks the other gateway periodically when the VPN is established. If no response is received, the VPN tunnel is closed. Indicates that the other gateway is down, unreachable, or considers the VPN tunnel already closed.
Encapsulation mode mismatch	Encapsulation modes (AH and/or ESP) did not match between gateways.

Table G.14 Common IPsec VPN Errors (Continued)

Error Message	Description
IKE error notify received: [...]	This message is visible only when IPsec diagnostics are enabled. The other gateway has sent the error notification that is shown in this message.
IKE negotiation rate-limit reached, discard connection	This message is visible only when IPsec diagnostics are enabled. There is an excessive number of new VPN connection attempts within a short period of time. This mechanism is meant to protect the firewall from certain types of denial-of-service attacks.
Invalid argument	Generic error. Check the other log messages for more useful information. If the problem is not apparent in the available logs, activate diagnostics to generate more verbose logs that give you more information about the next negotiations.
Invalid syntax	
IPsec SA proposal not accepted	This message is visible only when IPsec diagnostics are enabled. The VPN gateway at the other end of the tunnel sent a proposal that the Stonesoft gateway could not accept. This message includes information about the rejected proposal and a further log message should contain information on Stonesoft's local proposal.
NAT-T is not allowed for this peer	This message is visible only when IPsec diagnostics are enabled. NAT-T was requested by the other gateway but it is not allowed in the configuration of the gateway that sends this message.
No proposal chosen	IKE negotiations failed. If the problem is not apparent in the available logs, activate diagnostics to generate more verbose logs that give you more information about the next negotiations.
Payload malformed [...]	Most likely due to a mismatch in preshared keys between the initiator and the responder. May also be due to corruption of packets in transit.
Peer IP address mismatch	The IP address of the other gateway uses is not configured as a VPN gateway end-point on this gateway.
Proposal did not match policy	There is a mismatch in the configurations of the two negotiating parties.
Remote address not allowed	A VPN client is trying to use an IP address that is out of the allowed address range. Make sure all valid IP addresses are actually included in the range of allowed addresses in the Internal VPN Gateway properties and check the DHCP server configuration.
Remote ID mismatch	The IKE Phase 1 ID defined for the external security gateway in Stonesoft is different from the ID with which the gateway actually identified itself. The ID and its type are set for each tunnel End-Point in the properties of the external Gateway. Note that if an IP address is used as identity, the IP address used as the identity may be different from the IP address used for communications.
Remote identity [...] used in IKE negotiation doesn't match to policy [...]	
SA unusable	Usually means that an SA is being deleted when some new traffic arrives to use the tunnel.

Table G.14 Common IPsec VPN Errors (Continued)

Error Message	Description
Sending error notify: [...]	This message is visible only when IPsec diagnostics are enabled. Negotiations have failed and Stonesoft is sending the error notification that is shown in this message to the other gateway.
SPD doesn't allow connection [...]	Most likely indicates that the Site definitions do not match the IP addresses used. Check the addresses included under the Sites for both Gateways, and also that the translated addresses are included under the Site, if NAT is used for communications inside the VPN.
Timed out	Indicates connection problems or that the other end has deleted the SA that Stonesoft is using in the negotiation. Check the logs at the other end to see if the connection makes it through.
Traffic selector mismatch	There is a mismatch in the configurations of the two negotiating parties. You must define a matching pair for all settings; double-check all settings at both ends.
Tunnel policy mismatch [...]	This message is visible only when IPsec diagnostics are enabled. Usually indicates IKE negotiations failed because of a mismatch in the configurations of the two negotiating parties.
Tunnel selection failed	An Access rule matched this connection, but the traffic could not be sent across the VPN. Most likely, this is due to the (possibly NATed) source or destination IP address not being included in the local or remote gateway's Site as required. This message also appears if a connection that is not intended for the VPN matches the VPN rule (note that inbound cleartext traffic can be allowed from the same addresses as tunneled traffic with the Apply action in the VPN rule).
Tunnel type mismatch [...]	This message is visible only when IPsec diagnostics are enabled. Only gateway-to-gateway VPN or client-to-gateway VPN is configured, but the connecting device is of the other type. For example, a VPN client tried to connect, but VPN client access is not configured (correctly) on the gateway.

VPN Error Codes

Under some conditions, multiple IPsec VPN errors may be detected simultaneously and combined in a single log message. The most significant error is shown as text, and the other detected errors are indicated using a combined (with bitwise OR) hexadecimal error code.

Example IKE Phase-1 Initiator error: Proposal did not match policy (100002).
Here, the hexadecimal codes
00100000 for "Proposal did not match policy" and
00000002 for "Peer IP address mismatch") produces the code
00100002 = 100002.

The table below lists codes that are valid for engine software versions 5.0 and above.

Table G.15 Hexadecimal Error Codes in VPN Log Messages

Hex Code	Error Message
00000020	Access group mismatch
00008000	Authentication method mismatch
00020000	Encapsulation mode mismatch
00000002	Peer IP address mismatch
00100000	Proposal did not match policy
00400000	Remote address not allowed
00000040	Traffic selector mismatch (local)
00000080	Traffic selector mismatch (remote)
00200000	Tunnel type mismatch
00000200	Remote ID mismatch
00000100 00000004 00000001	Internal configuration-related problems. See the other messages to troubleshoot.

Audit Entry Types

The following table explains the audit entry types.

Table G.16 Audit Entry Types

Type	Definition
audit.info	Internal messages of the audit system.
audit.start	Start of an audit.
audit.stop	End of an audit.
stonesoft.admin.changeIp.mgtserver	Audited when management server IP address is changed.
stonesoft.admin.changeMgtIp.logserver	Audited when log server management IP address is changed.
stonesoft.admin.comment.change	Audited when a comment is changed.
stonesoft.admin.create	Creation of an administrator.
stonesoft.admin.delete	Deletion of an administrator.
stonesoft.admin.login	Audited when the administrator logs in to the management server.

Table G.16 Audit Entry Types (Continued)

Type	Definition
stonesoft.admin.logout	Audited when the administrator logs out from the management server.
stonesoft.admin.name.change	Change of administrator name.
stonesoft.admin.password.change	Change of password for an administrator.
stonesoft.admin.permission.change	Change of permissions for an administrator.
stonesoft.admin.session	Audits administrator sessions.
stonesoft.alert	Audited when management system sends an alert.
stonesoft.alert.policy.upload	Uploading a policy to an alert server - success or failure.
stonesoft.audit.archive.create	Audited when audit data archive is created.
stonesoft.audit.archive.delete	Audited when audit data archive is deleted.
stonesoft.audit.archive.restore	Audited when audit data archive is restored.
stonesoft.backup.create	Audited when a backup is created in the origin server.
stonesoft.backup.delete	Audited when a backup is deleted in the origin server.
stonesoft.backup.restore	Audited when a backup is restored in the origin server.
stonesoft.database.migrate	Audited when the server database is migrated.
stonesoft.database.password.change	Audited when database password is changed.
stonesoft.directarchive.start	Audited when the direct archive option is set to ON.
stonesoft.directarchive.stop	Audited when the direct archive option is set to OFF.
stonesoft.export.start	Audited when an export operation is started.
stonesoft.firewall.connections.terminate	Audited when a connection is terminated.
stonesoft.firewall.diagnostic	Diagnostic mode selected for a firewall.
stonesoft.firewall.disable.userdatabase	Audited when user database is disabled.
stonesoft.firewall.enable.userdatabase	Audited when user database is enabled.
stonesoft.firewall.initial.contact	Firewall performed initial contact to management server.
stonesoft.firewall.initial.generate	Initial configuration generated for a firewall.
stonesoft.firewall.monitor.off	A firewall monitoring change by an administrator to deactivated.
stonesoft.firewall.monitor.on	A firewall monitoring change by an administrator to activated.
stonesoft.firewall.policy.upload	Uploading a policy to a single firewall - success or failure.

Table G.16 Audit Entry Types (Continued)

Type	Definition
stonesoft.firewall.reboot	A firewall reboot by an administrator through the management system.
stonesoft.firewall.reset.database	Audited when the user database is reset.
stonesoft.firewall.state.lockoffline	A firewall state change by an administrator to locked offline.
stonesoft.firewall.state.lockonline	A firewall state change by an administrator to locked online.
stonesoft.firewall.state.offline	A firewall state change by an administrator to offline.
stonesoft.firewall.state.online	A firewall state change by an administrator to online.
stonesoft.firewall.state.standby	A firewall state change by an administrator to standby.
stonesoft.firewall.time.adjust	Firewall node time adjustment.
stonesoft.firewall.upgrade.end	Firewall node upgrade end through management system.
stonesoft.firewall.upgrade.start	Firewall node upgrade start through management system.
stonesoft.import.start	Audited when an import operation is started.
stonesoft.ips.analyzer.diagnostic	Diagnostic mode selected for an analyzer.
stonesoft.ips.analyzer.monitor.off	Monitoring mode offline for a sensor.
stonesoft.ips.analyzer.monitor.on	Monitoring mode online for a sensor.
stonesoft.ips.analyzer.policy.upload	Uploading a policy to an analyzer - single analyzer cluster success or failure.
stonesoft.ips.analyzer.reboot	Analyzer reboot through the management system.
stonesoft.ips.analyzer.state.lockoffline	Analyzer state changed to locked offline.
stonesoft.ips.analyzer.state.lockonline	Analyzer state changed to locked online.
stonesoft.ips.analyzer.state.offline	Analyzer state changed to offline.
stonesoft.ips.analyzer.state.online	Analyzer state changed to online.
stonesoft.ips.analyzer.state.standby	Sensor state changed to standby.
stonesoft.ips.analyzer.time.adjust	Analyzer node time adjusted.
stonesoft.ips.analyzer.upgrade.end	Analyzer node upgrade through management system ends.
stonesoft.ips.analyzer.upgrade.start	Analyzer node upgrade through management system begins.
stonesoft.ips.sensor.diagnostic	Diagnostic mode selected for a sensor.

Table G.16 Audit Entry Types (Continued)

Type	Definition
stonesoft.ips.sensor.monitor.off	Monitoring mode offline for a sensor.
stonesoft.ips.sensor.monitor.on	Monitoring mode online for a sensor.
stonesoft.ips.sensor.policy.upload	Uploading a policy to a sensor - single sensor success or failure.
stonesoft.ips.sensor.reboot	Sensor rebooted through the management system.
stonesoft.ips.sensor.state.lockoffline	Sensor state changed to locked offline.
stonesoft.ips.sensor.state.lockonline	Sensor state changed to locked online.
stonesoft.ips.sensor.state.offline	Sensor state changed to offline.
stonesoft.ips.sensor.state.online	Sensor state change by an administrator to online.
stonesoft.ips.sensor.state.standby	Sensor state changed to standby.
stonesoft.ips.sensor.time.adjust	Sensor node time adjusted.
stonesoft.ips.sensor.upgrade.end	Sensor node upgrade through management system ends.
stonesoft.ips.sensor.upgrade.start	Sensor node upgrade through management system begins.
stonesoft.layer2firewall.connections.terminate	Audited when a connection is terminated.
stonesoft.layer2firewall.diagnostic	Diagnostic mode selected for a Layer 2 Firewall.
stonesoft.layer2firewall.disable.userdatabase	Audited when user database is disabled.
stonesoft.layer2firewall.enable.userdatabase	Audited when user database is enabled.
stonesoft.layer2firewall.initial.contact	Layer 2 Firewall performed initial contact to management server.
stonesoft.layer2firewall.initial.generate	Initial configuration generated for a Layer 2 Firewall.
stonesoft.layer2firewall.monitor.off	A Layer 2 Firewall monitoring change by an administrator to deactivated.
stonesoft.layer2firewall.monitor.on	A Layer 2 Firewall monitoring change by an administrator to activated.
stonesoft.layer2firewall.policy.upload	Uploading a policy to a single Layer 2 Firewall - success or failure.
stonesoft.layer2firewall.reboot	A Layer 2 Firewall reboot by an administrator through the management system.
stonesoft.layer2firewall.reset.database	Audited when the user database is reset.

Table G.16 Audit Entry Types (Continued)

Type	Definition
stonesoft.layer2firewall.state.lockoffline	A Layer 2 Firewall state change by an administrator to locked offline.
stonesoft.layer2firewall.state.lockonline	A Layer 2 Firewall state change by an administrator to locked online.
stonesoft.layer2firewall.state.offline	A Layer 2 Firewall state change by an administrator to offline.
stonesoft.layer2firewall.state.online	A Layer 2 Firewall state change by an administrator to online.
stonesoft.layer2firewall.state.standby	A Layer 2 Firewall state change by an administrator to standby.
stonesoft.layer2firewall.time.adjust	Layer 2 Firewall node time adjustment.
stonesoft.layer2firewall.upgrade.end	Layer 2 Firewall node upgrade end through management system.
stonesoft.layer2firewall.upgrade.start	Layer 2 Firewall node upgrade start through management system.
stonesoft.license.activate	Audited when a license file or a license component is activated.
stonesoft.license.delete	Audited when a license component is deleted.
stonesoft.license.import	Audited when a license file is imported.
stonesoft.license.inactivate	Audited when a license is deactivated.
stonesoft.logdatamanager.abort	Audited when a scheduled task is aborted in the log server.
stonesoft.logdatamanager.complete	Audited when a scheduled task is completed in the log server.
stonesoft.logdatamanager.create	Audited when a scheduled task is created in the log server.
stonesoft.logdatamanager.delete	Audited when a scheduled task is deleted in the log server.
stonesoft.logdatamanager.modify	Audited when a scheduled task is modified in the log server.
stonesoft.logdatamanager.start	Audited when the user manually starts a task.
stonesoft.logpruningfilter.apply	Audited when a pruning filter is applied to the log server.
stonesoft.logpruningfilter.delete	Audited when a pruning filter is deleted from the log server.

Table G.16 Audit Entry Types (Continued)

Type	Definition
stonesoft.logpruningfilter.refresh	Audited when, following to a log server re-logging to the management, all the pruning filters are retrieved at the management and re-applied.
stonesoft.logreception.start	Log reception process begins.
stonesoft.logreception.stop	Log reception process ends.
stonesoft.logserver.certify	Audited when the log server is certified.
stonesoft.mgtserver.certify	Audited when the management server is certified.
stonesoft.object.delete	Audited when an object is deleted.
stonesoft.object.insert	Audited when a new object is added.
stonesoft.object.update	Audited when an object is updated.
stonesoft.policy.display	Generate a policy for display.
stonesoft.policy.upload.end	Uploading a policy ends.
stonesoft.policy.upload.start	Uploading a policy starts.
stonesoft.server.diskfull	Audited when the log server disk gets full.
stonesoft.server.start	Audited when the log server is started.
stonesoft.server.stop	Audited when the log server is stopped.
stonesoft.vpn.certificate.download	Audited when client downloaded a VPN certificate.
stonesoft.vpn.certificate.request	Audited when a VPN certificate is requested.
stonesoft.vpn.certificate.sign	Audited when a VPN certificate is signed.
stonesoft.vpn.gateway.remove	Audited when a VPN gateway is removed.
stonesoft.vpn.site.remove	Audited when a VPN site is removed.
stonesoft.vpn.validity.check	Audited when the VPN validity is checked.

Syslog Entries

The following table presents the categories for messages that appear in log entries sent to an external syslog server.

Table G.17 Syslog Entries

Value
Clock daemon for BSD systems
Clock daemon for System V systems
File transfer protocol
Kernel messages
Line printer subsystem
Mail system
Messages generated internally by syslogd
Network news subsystem
Network time protocol
Random user-level messages
Security/authorization messages
Security/authorization messages (private)
System daemons
UUCP subsystem

Log Fields Controlled by the Additional Payload Option

The following table presents the log fields that may be logged when the Additional Payload option is selected in an Inspection rule's Logging options.

Table G.18 Additional Payload Log Fields

Value
DNS qname
FTP command
FTP reply
FTP server banner
HTTP header
HTTP header name

Table G.18 Additional Payload Log Fields (Continued)

Value
HTTP request URI
HTTP request method
HTTP request version
ICMP field datagram reference
Imf encoded word
Imf header field
Imf token
SMTP command
SMTP misplaced command
SMTP recipient
SMTP reply
SMTP reverse path
SMTP server banner

Connection States

The following states are used both in the **State** column in the Connections view and (in part) in the Logs view in conjunction with info messages or logs on the closing of connections. They reflect the standard states regarding the initiation and termination of TCP connections as seen by the firewall in the transmissions. The table below lists the possible states.

Table G.19 Connection States

State	Description
CP established	Stonesoft cluster protocol packet is recognized.
ICMP echo	Ping reply is expected.
ICMP reply wait	Other ICMP request/reply types.
Invalid	The communication has violated the protocol.
IPsec established	IPsec tunnel packet is recognized.
New	New connection is being opened.
Related	New connection related to an existing one is expected soon.
Remove	Connection cannot be physically removed yet.

Table G.19 Connection States (Continued)

State	Description
Remove soon	Expecting to still see some packets (multiple reset packet), so delaying the removal for a few seconds. Eliminates unnecessary packet filtering and possible logging of dropped packets.
TCP close wait	One end of the connection waits for the FIN packet (passive close).
TCP close wait ack	Waiting for ACK for the FIN before going to close wait status (passive close).
TCP closing	Closing packet (FIN) sent by one end of the connection (simultaneous).
TCP closing ack	Waiting for ACK for the FIN before going to closing status (active close).
TCP established	Normal status of TCP connections for data transfer.
TCP fin wait 1	One end of the connection waits for sending the FIN packet (active close).
TCP fin wait 2	One end of the connection waits for receiving ACK packet.
TCP last ack	One end of the connection sent a FIN packet (passive close).
TCP last ack wait	Waiting for the FIN packet to be acknowledged.
TCP syn ack seen	Second phase of the TCP three-way handshake, the server has replied to client sent SYN with SYN+ACK, next status will be established.
TCP syn fin seen	T/TCP (Transactional TCP) connection, RFC 1644.
TCP syn return	Received simultaneous SYN from the other end (simultaneous open).
TCP syn seen	Very first packet sent by one end of the connection.
TCP time wait	One end of the connection acknowledged closing packet (FIN).
TCP time wait ack	Waiting for ACK for the FIN status before going to time wait status (active close).
UDP established	UDP connection is recognized.
Unknown established	Connection from other transport level protocol.

Appendix H

Keyboard Shortcuts

This chapter explains the shortcut keys that you can use in the Stonesoft Management Client. Most of these shortcuts are also shown in the menus of the Management Client next to each menu item that has a shortcut.

The following sections are included:

General Shortcuts

The following table lists general shortcuts that are available in most views in the Management Client.

Table H.1 General Shortcuts

Action	Shortcut
Bookmark all tabs	Ctrl+Shift+D
Bookmark current view	Ctrl+D
Cancel	Escape
Close tab	Ctrl+W
Close view	Ctrl+Q
Close window	ALT+F4
Copy	Ctrl+C
Cut	Ctrl+X
Delete, Remove	Delete
Go up/down/left/right	Up/down/left/right arrow key
Manage bookmarks	Ctrl+B
Maximize/Restore	Ctrl+M
Move to home level in element tree	F12, Alt+Up arrow key
Move to parent folder in element tree	Backspace
Open action in new tab	Ctrl+left mouse click
Open action in new window	Shift+left mouse click
Open a new tab	Ctrl+T
Open a new window	Ctrl+N
Open current view in new tab	Ctrl+Shift+T
Open current view in new window	Ctrl+Shift+N
Open help topics	F1
Open Info panel	Shift+F1
Open list of view-specific tools	Alt+T
Open Logs view	Ctrl+1
Open properties of the selected element	Ctrl+R

Table H.1 General Shortcuts (Continued)

Action	Shortcut
Open sidebar	F4
Paste	Ctrl+V
Print	Ctrl+P
Refresh	F5
Save	Ctrl+S
Search	Ctrl+F
Search references	Ctrl+G
Select all	Ctrl+A
Switch between open tabs	Ctrl+Tab

Shortcuts for Browsing Logs and Alerts

The following table explains the shortcuts available for browsing different types of logs in the Audit, Blacklist, Connections, and Logs views.

Table H.2 Keyboard Shortcuts in Log Browsing Views

Context	Action	Shortcut
Audit View, Logs View	Abort query	Escape
Audit View, Logs View	Acknowledge selected alert	Space
Audit View, Logs View	Apply query	Ctrl+Enter
Audit View, Logs View	Cancel query	Shift+Escape
Audit View, Logs View	Select columns	Ctrl+L
Audit View, Logs View	Current Events mode on/off	F11
Audit View, Logs View	Decrease text size	Ctrl+Shift+Minus
Audit View, Logs View	Enforce time range	F9
Audit View, Logs View	Go to the first log record	Ctrl+Home
Audit View, Logs View	Go to last log record	Ctrl+End
Audit View, Logs View	Increase text size	Ctrl+Shift+Plus
Audit View, Logs View	Jump backward in the timeline	Ctrl+ Page Up
Audit View, Logs View	Jump forward in the timeline	Ctrl+ Page Down
Audit View, Logs View	Refresh statistics	F6

Table H.2 Keyboard Shortcuts in Log Browsing Views (Continued)

Context	Action	Shortcut
Audit View, Logs View	Scroll horizontally	Left/Right arrow key
Audit View, Logs View	Scroll one page up/down	Page Up/Page Down
Audit View, Logs View	Scroll Up and Down	Up/Down arrow key
Audit View, Logs View	Normal Text Size	Ctrl+Shift+Zero
Audit View, Logs View	Zoom in on timeline	Ctrl+Plus
Audit View, Logs View	Zoom out on timeline	Ctrl+Minus
Audit View, Blacklist View, Connections View, Logs View	Increase/decrease text size	Ctrl+mouse wheel
Blacklist View, Connections View	Pause	F7
Blacklist View, Connections View	Play	F11
Blacklist View	Remove entry	Delete
Connections View	Terminate connection	Delete
Query Panel	Apply query	Ctrl+Enter
Query Panel	Cancel query	Shift+Escape

Other View-Specific Shortcuts

The table below lists other useful shortcuts that are available in different views in the Management Client.

Table H.3 Other Keyboard Shortcuts

Context	Action	Shortcut
Diagram Editor	Zoom In	Ctrl+Plus
Diagram Editor	Zoom Out	Ctrl+Minus
Diagram Editor, Security Policy Editor	Redo	Ctrl+Y
Diagram Editor, Security Policy Editor	Undo	Ctrl+Z
Diagram Editor, Security Policy Editor, VPN Editor	Close editor	Shitf+Escape
Diagram Editor, Security Policy Editor, VPN Editor	Start editing	Ctrl+E
Info Panel	Examine	Ctrl+Space
Overviews View, Reports View	Revert to last saved version	Shitf+Escape
Security Policy Editor	Add rule after	Ctrl+Insert
Security Policy Editor	Clear cell	Delete
Security Policy Editor	Collapse Comment Rule section/ sub-policy section	Left arrow key
Security Policy Editor	Delete rule	Ctrl+Delete
Security Policy Editor	Edit cell	F2
Security Policy Editor	Expand Comment Rule section/ sub-policy section	Right arrow key
Security Policy Editor	Move rule down	Alt+Down arrow key
Security Policy Editor	Move rule up	Alt+Up arrow key
Security Policy Editor	Search for next matching rule	F3
Security Policy Editor	Search for previous matching rule	Shift+F3

GLOSSARY

A

Access Control List

A list of Elements that can be used to define the Elements that an administrators with restricted permissions can access. See also Administrator Role and Granted Element.

Action

What the engine should do with a packet that matches the criteria for a particular rule in the security policy.

Action Option

Additional action-specific selections that affect how the traffic is handled set in the Action cell in rules.

Active Management Server

The Management Server that currently has control of all Domains in a system that has at least one Additional Management Server.

Additional Log Server

A Log Server defined as a backup channel for components that primarily send their logs to some other Log Server.

Additional Management Server

A redundant Management Server that replicates the configuration data from the Active Management Server under normal conditions so that the services offered by the Management Server can be used without interruption if components fail or are otherwise unavailable.

Address Range

A Network Element that defines a range of IP addresses. Use to avoid having to repeatedly type in the same IP addresses when defining address ranges that do not correspond to whole networks.

Address Resolution Protocol (ARP)

An Internet standard (RFC 826) protocol used to associate IP addresses with the media hardware address of a network interface card on a local area network (LAN).

Administrator

An Element that defines the details of a single person that is allowed to log on to the SMC using the Management Client. If used as a general term, Web Portal Users are also considered as administrators.

Administrator Role

An Element that defines which actions an Administrator with restricted permissions is allowed to take. See also Granted Element and Permission Level.

Aggressive Mode

The authentication of two IPsec end-points with only three messages, as opposed to Main Mode's six. Aggressive mode also does not provide PFS support, and SA negotiation is limited. See Main Mode (page 1277). See also Security Association (SA) (page 1284).

AH (Authentication Header)

See Authentication Header (AH) (page 1264).

Alert Chain

A list of rules defining which Alert Channels are used, and in which order, when an alert entry is directed to the Alert Chain from an Alert Policy to be escalated out from the Management Center. See also Alert Escalation.

Alert Channel

A method of sending alerts out from the Log Server. You can send alerts via SMTP (e-mail), SNMP, SMS text messages, or some other action you define in a custom script. Alert Channels are defined in the Log Server's properties, after which they can be used in Alert Chains.

Alert Element

An Element that gives the name and description to an Alert Event. The Alert element can be used as a matching criteria in the rules of an Alert Policy.

Alert Entry

A log message with an alert status that has been raised based on some Situation (which you can see in the Logs View). Alert entries trigger Alert Escalation.

Alert Escalation

Sending alerts out from the Management Center to administrators through Alert Channels (such as e-mail) according to a predefined Alert Chain until the original Alert Entry is acknowledged by some administrator in the Logs View.

Alert Event

A pattern in traffic or a problem in the system's operation that matches to some Situation used in a policy or internally in the system, and thus triggers an Alert Entry.

Alert Policy

A list of rules defining if an Alert Entry is escalated and which Alert Chain is used for escalating which type of alert entries. See also Alert Escalation.

Alert Server

A Stonesoft Management Center component that handles receiving and handling of Alerts. The Alert Server cannot be installed separately, it is integrated in the Log Server installation.

Alias

An Element that can be used to represent other network elements in configurations. It differs from a group element in that it does not represent all the elements at once: the value it takes in a configuration can be different on each engine where it is used.

Allow Action

An Action parameter that allows a connection matching that rule to pass through the Firewall to its destination.

Analyzer

1) A legacy device in the Stonesoft IPS system that analyzes the log information from Sensors according to its policy to find patterns, so that separate log entries can be combined together. See also Log Server, Security Engine.

2) The legacy Element that represents an Analyzer device in the Management Center.

Antispoofing

Technique used to protect against malicious packages whose IP header information has been altered. See also IP Spoofing (page 1274).

Application

A category of Tags for Situations. Meant for grouping Situations that detect known vulnerabilities in a particular software application.

Application Layer Gateway; Application Level Firewall

A firewall system, or gateway, in which packets are examined based on the application protocol being used (e.g., telnet, FTP, SMTP). Proxies for each application-level service are installed on the gateway, and are often configured to relay a conversation between two systems. That is, a packet's destination is the gateway, which then establishes a separate connection to the other system to complete the connection.

Apply VPN Action

A Firewall Action parameter that directs traffic from protected local networks into the Virtual Private Network (VPN) tunnel and allows traffic that arrives through a VPN, but does not match non-VPN traffic from outside networks into the protected networks. See also Enforce VPN Action (page 1270).

ARP (Address Resolution Protocol)

See Address Resolution Protocol (ARP) (page 1261).

Asymmetric Encryption

A cryptographic technology that uses a pair of keys. The message is encrypted with the public half of a pair and can then be decrypted only with the matching private half of the key pair. Public key technology can be used to create digital signatures and deal with key management issues. Also referred to as public key encryption. See also Symmetric Encryption (page 1287) and Public-key Cryptography (page 1282).

Auditing

A Management Center feature that logs administrators' actions and allows administrators with unrestricted permissions to view and manage these logs to keep track of system changes.

Authentication

The process of proving that someone or something is who or what they claim to be. For example, typing a simple username-password combination is a form of authentication.

Authentication Header (AH)

A security protocol supported by the IPsec protocol to enhance traffic security. It enables the authentication and integrity of data against packet corruption or tampering. AH protocol can use SHA-1 or MD5 to generate a hash signature based on a secret component from the SA, the packet payload and some parts of the packet header. See also Security Association (SA) (page 1284).

Authentication Server

A component of the Management Center that provides authentication services for end-user and Administrator logins.

Authentication Token/Authenticator

A portable device for authenticating a user. Authentication tokens typically operate by challenge/response, time-based code sequences, or other techniques. One of the most commonly used tokens is the RSA SecurID card.

Authorization

The process of giving someone/something permission to do or have something. Usually related to authentication; once a user has authenticated (proved who they are), they are authorized (given permission) to perform certain actions.

B

Balancing Mode

A Stonesoft cluster mode that attempts to divide the traffic as equally as possible between the online engines participating in the cluster. Confer to Standby Mode (page 1286).

Bandwidth Management

The process of determining and enforcing bandwidth limits and guarantees for different types of traffic either together with Traffic Prioritization or on its own. Also see QoS Class (page 1282) and QoS Policy (page 1282).

Blacklisting

1) The process of blocking unwanted network traffic either manually or automatically.

2) Persistently blocking access to certain URLs manually.

Bookmark

A stored link to a view or layout in the Management Client.

Bookmark Folder

A folder in the toolbar of the Management Client for storing and sharing Bookmarks.

Boot Recovery

A Stonesoft setting that brings the engines automatically back online after boot-up.

Border Routing

Routing of connections between different autonomous systems.

BrightCloud

A Web Filtering categorization service that provides categories for malicious sites as well as several categories for different types of non-malicious content that may be considered objectionable.

Buffer Overflow

When a program's data in the memory of a computer exceeds the space reserved for it (the buffer), data may in some circumstances be written on other parts of the memory area. Attackers may use buffer overflows to execute harmful program code on a remote system.

Bugtraq

A mailing list for discussing network security related issues, such as vulnerabilities.

Bulk Encryption Algorithm

Describes symmetric encryption algorithms which operate on fixed-size blocks of plaintext and generates a block of ciphertext for each.

C

CA

See Certificate Authority (CA) (page 1266).

CAN

A candidate for a CVE entry.

Capture Interface

An IPS Engine interface that can listen to traffic passing in the network, but which is not used for routing traffic through the engine. See also Inline Interface.

Category

A way of organizing elements and policies to display a specific subset at a time when configuring a large Stonesoft system in the Management Client to make it easier to find the relevant elements when configuring the system. For example, a Managed Service Provider (MSP) who manages networks of several different customers can add a customer-specific category to each element and policy to be able to view one customer's elements and policies at a time.

Certificate

Electronic identification of a user or device. Certificates prove the user or device is who/what they claim to be. This is done through using public/private key pairs and digital signatures. Certificates are used in Stonesoft for authenticating communications between the system components and for Virtual Private Network (VPN) authentication. Digital certificates are granted and verified by a Certificate Authority (CA), such as the internal CA included in the Management Server.

Certificate Authority (CA)

A trusted third-party organization or company that issues digital certificates, used to create digital signatures and public-private key pairs. The role of the CA in this process is to guarantee that the individual granted the unique certificate is, in fact, who he or she claims to be.

Challenge/Response

An authentication technique whereby a server sends an unpredictable challenge to the user, who computes a response using some form of authentication token, which can be an authenticator, or pre-shared keys used to encrypt random data.

Checksum

A one-way function applied to a file to produce a unique “fingerprint” of the file for later reference. File tampering can then be discovered by verifying the checksum value in the future.

CIS

See Content Inspection Server (CIS) (page 1267).

Client

In a client-server architecture, a client is usually an application running on a computer or a workstation that uses services provided by a Server.

Client Protection Certificate Authority

Contains the credentials that the engine uses to sign replacement server-side certificates the engine creates and presents to clients when inspecting the clients’ HTTPS connections with external servers. Also see Server Protection Credentials (page 1285).

Client-to-Gateway VPN

A Virtual Private Network (VPN) between a software client and a Security Gateway (SGW). Allows connecting mobile and home office workers safely to corporate resources using a secure (authenticated and encrypted) connection through insecure networks.

Cluster

A group of devices, or nodes, that share a given work load. In Stonesoft, you can cluster Firewalls, IPS engines, and Layer 2 Firewalls to share the load and provide redundancy, allowing, for example, scheduled maintenance that takes one node out of service without interrupting services to the users.

Cluster Mode

Determines if all members of a cluster participate to traffic processing at all times (Balancing Mode) or if other members remain inactive until a traffic-processing member stops processing traffic (Standby Mode).

Cluster Virtual IP Address (CVI)

An IP and MAC address shared by all nodes in a cluster, which are used by every node in a cluster for communication. These interfaces give the cluster a single identity on the network, reducing the complexity of routing and network design. CVIs handle the traffic directed to the Firewall for inspection in Firewall Clusters.

Combined Sensor-Analyzer

1) A legacy IPS device that has both Sensor and Analyzer engines running simultaneously on the same hardware.

2) The legacy Element that represents a Combined Sensor-Analyzer device in the Management Center.

See also IPS Engine.

Connection Tracking

The set of data maintained for a connection. Used for relating incoming packets to existing connections. Connection tracking information also includes information to necessary for NAT (Network Address Translation), Load Balanced Routing and Protocol Agents. May also contain accounting information.

Contact Address

The IP address that is needed to contact a device performing a function in the Stonesoft Management Center when there is NAT (Network Address Translation) being performed in between the two devices and thus the actual IP address assigned to the network interface cannot be used directly.

Content Inspection Server (CIS)

A server that performs detailed examination of a connection’s data and assists in the determination to allow or discard packets. Common examples include virus scanning or filtering of web URLs. Also known as *content screening*.

Continue Action

A policy parameter that sets default values to those used in the rule. The defaults are used in all subsequent rules except where specifically overridden until some other rule with the Continue action changes the values or the policy ends.

Context

An Element that is added to a Situation to define what the Situation should match. Provides a framework for defining parameters, which are most entered as a regular expression, or through a set of fields and options that the administrators adjust.

Correlation Situation

A Situation that defines the patterns that the Analyzer looks for when it examines event data produced by Sensors.

CRL Server

A server that maintains a Certificate Revocation List (CRL), which can be used in Authentication to check if the certificate has been cancelled.

Custom Alert

An Alert Element that is defined by a Stonesoft administrator, as opposed to a ready-made Default Element created by Stonesoft.

CVE

A dictionary that provides common names for publicly known information security vulnerabilities and exposures and thus a standardized description for each vulnerability that links the vulnerability information of different tools and databases.

CVI

See Cluster Virtual IP Address (CVI) (page 1266).

D

Default Element

An Element that is present in the system at installation, or is added to the system during an upgrade or from a Dynamic Update (Package). Default elements cannot be modified or deleted by administrators, but they may be modified or deleted by dynamic update packages or upgrades.

Defragmentation

The process by which a large block of data is broken up into smaller pieces (datagrams), so that it can be packaged and transmitted by the underlying network technology (Fragmentation). Once the smaller pieces arrive at their destination, the datagrams are reassembled into the larger block of data (defragmentation).

DHCP (Dynamic Host Configuration Protocol)

A protocol for dynamically assigning IP addresses and other network information to an interface, based on BOOTP. A device on a network with no network information can broadcast a request for an IP address, subnet mask, default gateway and other information from a DHCP server on that same network. DHCP is defined in RFC 2131.

Diagram

An Element that contains one or more network diagrams created using the Diagram Editor.

Digital Certificate

See Certificate (page 1265).

Discard Action

An Action parameter that stops all connections matching to the rule without sending any notification to the connecting host. Confer to Refuse Action (page 1282).

Dispatch Clustering

See Packet Dispatch (page 1280).

DMZ Network

A DMZ (DeMilitarized Zone Network) is a network separate from both internal and external networks, and connected through a gateway. Often used for isolating bastion hosts or publicly available machines, e.g., mail and HTTP servers are typically located on a DMZ network. Sometimes also referred to as a *screened subnetwork*.

DNS Spoofing

An attack method whereby the DNS name of a system is assumed by a malicious system, either by corrupting the name service cache of a victim, or by compromising a domain name server for a valid domain. The victim system is then directed to the malicious system instead of the original server.

Domain

Domains are administrative boundaries that allow you to separate the configuration details and other information in the system for the purpose of limiting administrator access.

DoS Attack (Denial of Service)

An attack with the objective of causing enough disruption in a computer system that its usability to legitimate users suffers. For example, and attacker may target a website so that it becomes overloaded, and slows down so much that it becomes unusable for people wishing to view it.

DSCP (DiffServ Code Point)

The Differentiated Services (DiffServ) Type of Service (ToS Flag) field added to packets in the network.

DSCP Mark

A field in QoS Policy rules that writes a particular DSCP (DiffServ Code Point) marker to the packets, if the QoS Policy is applied on the interface the packets use to exit the Firewall.

DSCP Match

A field in QoS Policy rules that assigns the QoS Class specified in the rule to incoming packets that have a specific DSCP (DiffServ Code Point) marker set, if the QoS Policy is applied on the interface the packets use to enter of the Firewall.

Dynamic IP address

An IP address that is assigned by using the DHCP (Dynamic Host Configuration Protocol).

Dynamic NAT

A way to translate network addresses, where for each original address, a translated address and possibly a port are selected dynamically from a predefined pool.

Dynamic Update (Package)

A file supplied by Stonesoft that provides updates to Default Elements and policies, most importantly to the Situation and Vulnerability information that is used for traffic inspection in Inspection Rules.

Element

A Stonesoft object representing the equipment in your physical networks or some area or concept of configuration. Elements may, for example, represent a single device such as a server, a range of IP addresses, or some configuration aid in the Management Center, such as a Category. Also see Network Element (page 1279).

Encryption

Used for data security, encryption translates any data into a secret code. Public-key encryption and symmetric encryption are the main types of encryption. Decrypting ciphertext (encrypted data) into plaintext requires access to a secret key.

Encryption Domain

Networks that are defined to be behind a certain VPN gateway in a Virtual Private Network (VPN) configuration.

Encryption Key

The data that is used to convert plaintext to ciphertext. In symmetric algorithms, the same key is the decryption key as well. In public key algorithms, a different, but related key is used to convert the ciphertext back into plaintext.

Encryption Policy

Settings that define which encryption and authentication methods are used to establish a Virtual Private Network (VPN).

Enforce VPN Action

A Firewall Action parameter that directs traffic from protected local networks into the Virtual Private Network (VPN) tunnel and allows traffic that arrives through a VPN, and drops any non-VPN traffic from external networks to the local network that matches the rule. See also Apply VPN Action (page 1263).

Ethernet Rules

A set of rules in the IPS Policy that define which Ethernet traffic is allowed or discarded by a Sensor in Transparent Access Control Mode.

Expression

An Element that can be used to accurately define a whole complex set of elements by including and excluding elements using logical expressions.

External Gateway

Any Security Gateway (SGW) that is managed by a different Management Server than the one on which the Virtual Private Network (VPN) is being configured.

F

Filter

A description of log fields and their values combined together using operators for the purpose of sorting in or out log, alert, and audit entries. Used, for example, to filter out logs from the display in the Logs View so that those entries that are interesting at the moment can be found more easily.

Firewall

1) An Element that represents the firewall device in the Management Center. Either a Single Firewall or a Firewall Cluster.

2) The device running the Stonesoft Firewall software.

Firewall Cluster

A Group of two or more *Firewall Engines* that work together as if they were a single unit.

Firewall Engine

The device that runs the Stonesoft Firewall software; a standard server, an engine installed on a virtualization platform, or a Stonesoft appliance. Represented by the Firewall Node in the Management Client.

Firewall Node

An individual Firewall Engine in the Management Client, representing a device that runs Stonesoft Firewall software as part of a Firewall Cluster or a Single Firewall.

Forward Action

A Firewall Action parameter that directs traffic from protected local networks or from a Virtual Private Network (VPN) tunnel into another VPN tunnel.

Fragmentation

The process by which a large block of data is broken up into smaller pieces (datagrams), so that it can be packaged and transmitted by the underlying network technology (fragmentation). Once the smaller pieces arrive at their destination, the datagrams are reassembled into the larger block of data (Defragmentation).

G

Gateway

A device that provides VPN access for other devices.

Gateway Certificate

A Certificate used for authenticating a Gateway to other Gateways and VPN Clients in a VPN.

Gateway Profile

An element that defines a set of VPN-related capabilities that a VPN Gateway supports.

Gateway Settings

An element that contains general settings for Stonesoft Firewall/VPN engines related to VPN performance.

Gateway-to-Gateway VPN

In Stonesoft, a Virtual Private Network (VPN) element which is set up so that the VPN is established between two gateway devices providing connectivity to networks behind the gateways.

Geolocation

Elements that define a geographical location of an IP address. Used for illustrating networks and network traffic on a map and other informative purposes in the Management Client.

Granted Element

An Element or Security Policy that an administrator has been given permission to edit and install when their Administrator Role would otherwise prevent them from doing so.

Group

A Network Element that includes other elements and represents them all at once in policies and other parts of the configuration. For example, you can define a Group of several WWW-servers, and then use the Group element in policies when you need to make a rule that concerns all of the WWW-servers.

H

Hardware

A category of Tags for Situations. Meant for grouping Situations that detect known vulnerabilities in applications that run on a particular hardware platform.

Hash Signature

A cryptography-related concept that refers to a digital fingerprint associated with a given message and computed with one-way algorithms. Hash signatures are used to secure the integrity of encrypted data, ensuring that no tampering has taken place during transmission. See also Client-to-Gateway VPN (page 1266), and SHA-1 (page 1285).

Heartbeat

A protocol that the nodes of a Firewall Cluster or Sensor Cluster use to monitor each other and for other tasks that are needed for collaboration between each Node.

High Availability

The implementation of clustering technology, hot standby technology, or general redundancy in a system to increase the availability of an application, service, or network beyond what a single system is capable of providing. Increased availability is achieved by eliminating all single points of failure, with clustering technology providing the highest level of availability.

Host

1) A Network Element that represents any single device that has an IP address.

2) Any device connected to a TCP/IP network, including the Internet, with one or more IP addresses. Hosts are distinguishable from gateways or routers, in that they do not forward, or route, packets to other networks.

Hot Standby

A solution where one node handles the work load with the support of a back-up node, which takes over connections in case of failure in the first node.

Hybrid Authentication

A system using both Asymmetric Encryption and Symmetric Encryption. Asymmetric techniques are used for key management and digital signatures. The symmetric algorithms are used to encrypt the bulk of data with reduced strain on resources.

I

IKE Proposal

The suggested encryption algorithms, authentication methods, hash algorithms, and Diffie-Hellman information in the Security Association (SA) component of an IPsec VPN. The initiator of an IPsec tunnel can make multiple proposals, but the responder only sends one proposal in return. See also Internet Key Exchange (IKE) (page 1274) and Security Association (SA) (page 1284).

Incident Case

An Element that administrators can use to gather together all the data, actions, system configuration information, and files related to a specific incident of suspicious activity.

Incident History

A collection of all the logs and audit entries that track actions performed in a particular Incident Case window.

Info Panel

A tab in Management Client windows that shows information on the selected element or other object. The Info view shows, for example, the nodes belonging to a selected cluster.

Inherited Rule

A rule either hidden or shown on a grey background in a Security Policy or Template Policy which has been added in a template higher up in the policy hierarchy so that it has been passed down to the security policy or template policy. Inherited rules are enforced just as any other rules, but they can be edited only in the template where the rule was originally added.

Inline Interface

An IPS Engine or Layer 2 Firewall interface that combines together two physical interfaces, enabling the traffic to be routed through as if the engine were an extension of the network cable, but allowing the engine to actively monitor packets and connections and stop them according to its Actions and Inspection Rules.

Insert Point

The place in a Security Policy or Template Policy where new rules can be inserted when no rules have been inserted in that place yet (shown as a green row) or the place in a template policy where rules can be inserted in inheriting policies and template policies (shown as an orange row).

Inspection Rule

The definitions on the Inspection tab in a Firewall or IPS policy that defines options for deeper inspection and reactions to traffic accepted in Actions. The matching in Inspection rules is done based on matching information provided by Situation elements. Confer to Action (page 1261).

Internal Gateway

A Stonesoft Firewall/VPN engine that are managed by the same Management Server on which the Virtual Private Network (VPN) is being configured.

Internal Network

The networks and network resources that Stonesoft is protecting. In Stonesoft, there is no concept of internal and external networks in the system.

Internet Key Exchange (IKE)

A protocol defined by the IPsec (IP Security) standard for securely exchanging key-related information between connecting hosts when establishing a Virtual Private Network (VPN).

Internet Service Provider (ISP)

A company that provides Internet connectivity to subscribers.

Intrusion Detection System (IDS)

A system that monitors network traffic for determining, and making administrators aware of data security exploits or attempts by providing logs or other network information. Confer to Intrusion Prevention System (IPS).

Intrusion Prevention System (IPS)

A system that monitors network traffic (like an Intrusion Detection System (IDS)) and has the capability of actively stopping traffic if it is deemed malicious or otherwise unwanted.

IP Address Bound License

A License file for the engines that includes the information on the IP address of the component it licenses. If you need to change the IP address of the component, you must request an IP address change at the Stonesoft Licensing website. On engines, an alternative to a Management Bound License (page 1277).

IPComp (IP Payload Compression Protocol)

A protocol used to reduce the size of IP datagrams. Increases the overall communication performance between a pair of communicating gateways by compressing the datagrams, provided the nodes have sufficient computation power, and the communication is over slow or congested links. IPComp is defined in RFC 2393.

IP Splicing (or Hijacking)

An attack performed by intercepting and using an active, established session. Often occurs after the authentication phase of the connection is complete, giving the attacker the permissions of the original, authenticated user. Encryption at the session or network layer is typically the best defense from such an attack.

IP Spoofing

A technique used to obtain unauthorized access to computers by sending connection requests with tampered headers, simulating a trusted source.

IPsec (IP Security)

A set of protocols supporting secure exchange of packets. Used for the implementation of Virtual Private Network (VPN) solutions when high performance and/or support for a wide variety of protocols are needed. IPsec provides transport and tunnel encryption modes. IPsec is defined in RFC 2401.

IPsec Proposal

Suggested encryption algorithms, hash algorithms, authentication methods, etc. to be used for an IPsec (IP Security) tunnel. See also IKE Proposal (page 1273).

IPS Cluster

Group of two or more IPS engine nodes that work together as if they were a single IPS.

IPS Engine

1)A Stonesoft IPS component that captures all the traffic from a physical network link, inspects it according to its policy, and if installed inline, selects which connections are allowed to continue.

2) The device that runs IPS software; a standard server, an engine installed on a virtualization platform, or a Stonesoft appliance.

IPS Policy

The Security Policy for IPS Engines that contains the Action and Inspection Rule definitions that determine how traffic is inspected and how the system reacts when a match is found.

IPv4 Access Rule

A row in a Firewall or IPS policy that defines how one type of IPv4 connection is handled by providing matching criteria based on the source, destination, and protocol information. Confer to IPv6 Access Rule (page 1275).

IPv6 Access Rule

A row in an IPS policy that defines how one type of IPv6 connection is handled by providing matching criteria based on the source, destination, and protocol information. Confer to Action (page 1261).

ISAKMP (Internet Security Association Key Management Protocol)

An open-ended encoding protocol necessary for IKE negotiation when establishing Security Associations. See also Security Association (SA) (page 1284).

ISP (Internet Service Provider)

See Internet Service Provider (ISP) (page 1274).

J

Journal

A tool in the Incident Case window that allows administrators to create a permanent record of their actions while investigating an incident.

Jump Action

A Security Policy parameter that directs the inspection to a Sub-Policy, against which connections matching the rule with the Jump action are checked. Can be used to speed up traffic processing, as connections that do not match the Jump rules are not checked against rules in the sub-policies.

L

Layer 2 Firewall

A basic Stonesoft component that provides access control and deep inspection of traffic.

License

Files you import to the system to tell the Management Server that the components you have installed have been legally purchased. You generate the Licenses at the Stonesoft Licensing website and import them to the Management Server using the Management Client.

Lifetime

The interval at which the IPsec participants should begin to negotiate a replacement Security Association (SA) (soft lifetime) or the interval at which the current SA for an IPsec tunnel is no longer valid (hard lifetime) in a Virtual Private Network (VPN).

Load Balancing

A process for distributing work evenly across multiple, available devices to avoid overwhelming any single system.

Load Balancing Filter

A software component that determines which network connections should be handled by a particular node in a cluster, based on address information, current load, performance of individual machines, and other factors.

Load Balanced Routing

A method for choosing routes to destinations based on determining the fastest response time through multiple gateways. The application of Multi-Link technology to determine which network link provides the best round trip time.

Load Sharing

The distribution of work between multiple devices. Similar to Load Balancing, but not as effective, since the techniques used do not ensure an *equal* distribution of the work load. Load sharing is typically a static method of distributing a load, whereas load balancing is often a dynamic method.

Location

An Element that groups together system components that are on the same side of a device doing NAT (Network Address Translation). Used to define Contact Addresses for components that communicate within the Stonesoft Management Center.

Logging Options

A selection available in all rules in policies that determines if and how a record is created when the rule matches.

Logging Profile

Defines how the Log Server converts Syslog data received from a particular type of third-party component into Stonesoft log entries.

Log Server

A component of the Management Center responsible for storing and managing log (and alert) data, and analyzing and correlating events detected by multiple Security Engines.

Log Spool

A temporary storage area in an engine node for log data before it is sent to a Log Server.

Logical Interface

An IPS Element used in the IPS policies to represent one or more physical network interfaces as defined in the Sensor properties.

Logs View

A tool that allows browsing logs, alerts, audit data, and connections each in an adapted version of the same user interface.

M

Main Mode

An IKE negotiation mode, which exchanges six messages between the end-points of an IPsec tunnel to complete the negotiation of authentication and keys for a Virtual Private Network (VPN). Optionally, Perfect Forward Secrecy (PFS) can be applied to protect further negotiations. See also Aggressive Mode (page 1262) and Perfect Forward Secrecy (PFS) (page 1280).

Malware

Malicious software designed to infiltrate or damage a computer system.

Management Bound License

A License file for Stonesoft engines that is based on information on the Management Server’s Proof of License (POL) code. An alternative to an IP Address Bound License (page 1274).

Management Center

The system consisting of a Management Server, one or more Log Servers and none to several Web Portal Servers that is used to manage the Firewall Engines, and to store and manage traffic and system related data.

Management Client

A graphical user interface component that provides the tools for configuring, managing, and monitoring the Security Engines, and other components in the Stonesoft system. The Management Client connects to the Management Server to provide these services based on the Administrator information that you use when launching the Management Client software.

Management Network

The network used for communication between firewalls, Management Servers, Log Servers and the Management Client.

Management Server

A system component that stores all information about the configurations of all Security Engines, and other Stonesoft components in the system, monitors their state, and provides access for Management Clients when administrators want to change the configurations or command the engines. The most important component in the system.

Master Engine

A physical engine device that provides resources for Virtual Security Engines.

Maximum Transmission Unit (MTU)

The largest physical size of a datagram that can be transmitted over a network without fragmentation. Often expressed in bytes, it can apply to frames, packets, cells or other media, depending on the underlying topology.

Modem Interface

A Firewall interface that defines the settings of a 3G modem that provides a wireless outbound link for a Single Firewall.

Monitored Element

A Stonesoft server or engine component that is actively polled by the Management Server, so that administrators can keep track of whether it is working or not. All Stonesoft system components are monitored by default.

Monitoring Agent

A software component that can be installed on servers in a Server Pool to monitor the server's operation for the purposes of Traffic Management.

Multicast

A technique by which a set of packets are sent to a group of machines sharing a common address. Unlike broadcast, it does not include all machines, and unlike unicast, it usually has more than one member of the group.

Multi-Layer Inspection

A hybrid firewall technology that incorporates the best elements of application level and network level firewalls, with additional technology to enable the secure handling of many connection types.

Multi-Link

Patented Stonesoft technology to connect one site to another, or to the Internet, using more than one network link. Applications of Multi-Link technology include inbound and outbound traffic management for unencrypted as well as VPN traffic. See also Outbound Multi-link (page 1280).

N

NAT (Network Address Translation)

A mechanism for assigning local networks a set of IP addresses for internal traffic and another for external traffic. It increases security by hiding internal IP addresses and enables hosts with "invalid" (non-routable) addresses to communicate on the Internet.

NDI

See Node Dedicated IP Address (NDI) (page 1279).

NetLink

An Element used for implementing routing of Stonesoft's Multi-Link features. NetLinks can represent any IP-based network links (such as ISP routers, xDSL, leased lines, dial-up modems). NetLinks are combined together into an Outbound Multi-link.

Network Element

1) All Elements that represent one or more components that have an IP address, that is, a general category (‘Network Elements’) for those elements that represent physical devices and networks in Stonesoft.

2) The Network Element called ‘Network’ that represents a (sub)network of computers. Used for rules and configurations that are common for all hosts in a specific (sub)network.

Network Scan

A stage of an attack in which the attacker scans the target to enumerate or map the directly-connected network(s).

Node

The representation of an individual Security Engine in the Management Client.

Node Dedicated IP Address (NDI)

A unique IP address for each machine. The only interface type for Single Firewalls. Not used for operative traffic in Firewall Clusters, IPS engines, and Layer 2 Firewalls. Firewall Clusters use a second type of interface, Cluster Virtual IP Address (CVI), for operative traffic. IPS engines have two types of interfaces for traffic inspection: the Capture Interface and the Inline Interface. Layer 2 Firewalls only have Inline Interfaces for traffic inspection.

O

Operating System

A category of Tags for Situations. Meant for grouping Situations that detect known vulnerabilities in a particular operating system or applications that run on that operating system.

Outbound Multi-link

An Element used for combining NetLinks for load balancing outbound traffic. The NetLinks included in a Outbound Multi-link element are frequently tested to determine which is the fastest NetLink for new outbound connections.

P

Packet

A segment of data sent across a network that includes a header with information necessary for the transmission, such as the source and destination IP addresses.

Packet Dispatch

A Cluster Virtual IP Address (CVI) mode in which only one node in the cluster receives packets. This dispatcher node then forwards the packets to the correct node according to Load Balancing, as well as handles traffic as a normal node. The recommended cluster mode for new installations.

Packet Filtering

A method of controlling access to a network, or set of networks, by examining packets for source and destination address information, and permitting those packets to pass, or halting them based on defined rules.

Packet Sniffer

See Sniffer (page 1285).

Perfect Forward Secrecy (PFS)

A property of IKE transactions that enhances the secrecy of keys, but requires additional processing overhead. PFS ensures that the distribution of key-related information remains independent from previously existing key material. See also Internet Key Exchange (IKE) (page 1274).

Permission Level

The general level of rights that an Administrator has. Permissions are customized with Administrator Roles and Granted Elements.

Permit Action

An Inspection Rule action that stops the inspection of all traffic that matches to the rule that uses the Permit action and lets the traffic continue to its destination.

Phishing

A Social Engineering attack in which a malicious e-mail or web page attempts to solicit sensitive information such as usernames, passwords, and credit card details by masquerading as coming from a trustworthy entity.

Player

Any element or IP address that was involved in an incident that is being investigated using the Incident Case element.

Policy

A container for the Access rules, Inspection rules, and NAT rules.

Policy Routing

User-defined routing based on information that is not normally used in routing, such as the source IP address, port information, or service type.

Policy Snapshot

A record of policy configuration that shows the configuration in the form that it was installed or refreshed, including the rules of the policy, the elements included and their properties, as well as the time when the policy was uploaded, and which administrator performed the upload. Helps in keeping track of configuration changes.

Port Address Translation (PAT)

A process, similar to NAT (Network Address Translation), where the source or destination port is changed to a different port. PAT is often used to disguise, or masquerade a service in place of another. See also NAT (Network Address Translation) (page 1279).

Pre-shared Key

A string of characters that is stored on two (or more) systems and that is used for authenticating or encrypting communications between the systems.

Probing Profile

Settings that define how a Log Server monitors third-party components.

Proof of License (POL)

A code used for verifying the legitimate purchase of Stonesoft software products. Used for generating License files at the Stonesoft website.

Proof of Serial Number (POS)

Identification code attached to Stonesoft appliances.

Protocol

An element that is used inside Service elements to specific a Protocol Agent for the Firewall Actions and the protocol of the traffic for the Inspection Rules.

Protocol Agent

A process on the engines that assists the engine in handling a particular Protocol. Protocol Agents ensure that related connections for a service are properly grouped and evaluated by the engine, as well as assisting the engine with content filtering or network address translation tasks. See also Connection Tracking (page 1267).

Protocol Tag

A type for Protocol elements that are only used to define the protocol of traffic for inspection against the inspection rules. Confer to Protocol Agent.

Proxy ARP

Proxy ARP option on a device that does routing means that the device relays broadcast messages between two hosts that are in separate physical networks, but still have IP addresses from the same network. This proxy is needed for the ARP requests, as broadcast messages are not normally relayed from one network to another. See also Address Resolution Protocol (ARP) (page 1261).

Pruning

Deleting log entries according to Filters either as the logs arrive on the Log Server or before they are stored (after displaying them in the current view in the Logs view).

Public-key Cryptography

A cryptographic system that uses a pair of keys: a public key, used to encrypt a message, and a private (secret) key that can decrypt the message. This is also called asymmetric encryption.

Q

QoS Class

An Element that works as a link between a rule in a QoS Policy and one or more Firewall Actions. The traffic allowed in the access rule is assigned the QoS Class defined for the rule, and the QoS class is used as the matching criteria for applying QoS Policy rules.

QoS Policy

A set of rules for Bandwidth Management and Traffic Prioritization for traffic that has a particular QoS Class, or rules for assigning QoS Classes based on a DSCP Match found in the traffic.

R

Refragmentation

A technique to fragment outbound packets from the engine in the same manner in which they were fragmented when the engine received them. See also Virtual Defragmentation (page 1289).

Refuse Action

An Action parameter that blocks the packet that matches the rule and sends an error message to the originator of the packet. Confer to Discard Action (page 1268).

Regular Expression

A string that describes a set of strings. Used in many text editors and utilities to search for text patterns and, for example, replace them with some other string. In Stonesoft, regular expressions are used, for example, for defining patterns in traffic that you want a certain Situation to match when you give the Situation a Context that calls for a Regular Expression.

Related Connection

A connection that has a relationship to another connection defined by a Service. For example, the FTP protocol defines a relationship between a control connection, and one or more data connections at the application level. The engine may be required to allow a connection that would otherwise be discarded if it is related to an already allowed connection.

Request for Comments (RFC)

A document that outlines a proposed standard for a protocol. RFCs define how the protocol should function, and are developed by working groups of the Internet Engineering Task Force (IETF), and reviewed and approved by the Internet Engineering Steering Group (IESG). See `http://www.rfc-editor.org/`.

Retained License

A Management Bound License that has been used to install a policy on an engine and has then been unbound without relicensing or deleting the engine the license was bound to. Retained licenses cannot be bound to any engine before the engine the license was previously bound to is deleted or has a new policy refresh with a valid license.

RFC

See Request for Comments (RFC).

Rootkit

A set of tools that intruders to computer systems use for hiding their presence and the traces of their actions.

Route

The set of routers or gateways a packet travels through in order to reach its destination. In TCP/IP networks, individual packets for a connection may travel through different routes to reach the destination host.

Router

A Network Element representing a physical router in your network. Most often used to indicate next-hop routers in the Routing view and in Network Diagrams.

Routing Table

A database maintained on every router and gateway with information on paths to different networks. In Stonesoft, the routing table is represented graphically in the Routing view.

Rule

An expression used to define the eventual outcome of packets arriving at the engine, which match certain conditions (e.g., source and destination address, protocol, user).

Rules Tree

The main configuration tool for adjusting Inspection Rule definitions.

S

SA (Security Association)

See Security Association (SA) (page 1284).

Scan

See Network Scan (page 1279).

Secondary IP address

An IP address used for identifying an element with multiple addresses as a source or destination of traffic, defined in addition to a primary IP address.

Secret Key Cryptography

See Symmetric Encryption (page 1287).

Security Association (SA)

A unidirectional, logical connection established for securing Virtual Private Network (VPN) communications between two sites. A security association records the information required by one site to support one direction of the IPsec connection whether inbound or outbound. It uses transport mode for communications between two hosts and tunnel mode for communication between security gateways. See also Authentication Header (AH) (page 1264).

Security Engine

1) A type Element that represents a Security Engine device in the Management Center. See also Firewall, IPS Engine, and Layer 2 Firewall.

2)The device that runs Security Engine software in Firewall, IPS Engine, or Layer 2 Firewall mode. Can be a standard server, an engine installed on a virtualization platform, or a Stonesoft appliance.

Security Gateway (SGW)

A device, typically a firewall, that performs encryption/decryption on Virtual Private Network (VPN) packets sent between Sites through untrusted networks.

Security Parameter Index (SPI)

A value used by AH and ESP protocols to help the Firewall Cluster select the security association that will process an incoming packet. See also Authentication Header (AH) (page 1264).

Security Policy

The set of templates, policies, and sub-policies together or individually that define what traffic is acceptable and what traffic is unwanted. Policies are defined using the Management Client, stored on the Management Server and installed on Security Engines, which then use their installed version of the policies to determine the appropriate action to take regarding packets in the network.

Sensor

A legacy Stonesoft IPS component that captures all the traffic from a physical network link, inspects it according to its policy, and if installed inline, selects which connections are allowed to continue. Provides data for the Analyzer (see Analyzer (page 1263)).

Sensor Cluster

Group of two or more legacy IPS Sensor nodes that work together as if they were a single Sensor.

Server

1) A Network Element representing a physical server in your network. Generally, server elements are only defined to configure a specific server for use with the Management Center (such as a RADIUS server used for authenticating administrators), but generic Servers can be used in Network Diagrams instead of Host elements to better illustrate the network layout.

2) In a client-server architecture, a computer that is dedicated for running services used by Client computers. The services may include, for example, file storage, e-mail, or web pages.

Server Pool

A Network Element representing a group of Servers. Used for inbound traffic management.

Server Protection Credentials

An element that stores the private key and certificate of an internal HTTPS server. The private key and certificate allow the engine to present itself as the server to clients so that the engine can decrypt and inspect incoming HTTPS traffic. Also see Client Protection Certificate Authority (page 1266).

Service

An Element that is used for matching traffic to an application level protocol, for example, FTP, HTTP or SMTP. The TCP and UDP Services also determine the port number. Service elements are used in policies to make the rule match only a particular protocol, to enable Protocol Agents, and select traffic to be matched against Inspection Rules.

Session Stealing

See IP Splicing (or Hijacking) (page 1274).

SHA-1

A cryptographic algorithm used for hash functions. It generates a 160-bit signature from an input of any length. See also Hash Signature (page 1272).

Single Firewall

A firewall that has only one Firewall Engine.

Single Point of Failure

The point at which the failure of a single device or component of a system will lead to either the failure of the entire system, or the inability to use services normally provided by that system. Redundant systems, using high availability technologies, eliminate single points of failure.

Site

A set of resources protected by Stonesoft.

Situation

1) An Element that identifies and describes detected events in the traffic or in the operation of the system. Situations contain the Context information, i.e., a pattern that the system is to look for in the inspected traffic.

2) An Inspection Rule cell where Situation elements are inserted.

Situation Type

A category of Tags for Situations. Meant for indicating what kind of events the associated Situations detect (for example, Attacks, Suspicious Traffic).

Sniffer

A device or program that captures data traveling over a network. Sniffers are often used for troubleshooting network problems, as they can show the packet flow taking place. They can also be used maliciously to steal data off a network.

SNMP Agent

A software component that sends SNMP traps when specific events are encountered.

Social Engineering

An attack involving trickery or deception for the purpose of manipulating people into performing actions or divulging confidential information.

SPI (Security Parameter Index)

See Security Parameter Index (SPI) (page 1284).

SSH (Secure Shell)

A program to log into another computer over a network, to execute commands in a remote machine, and to move files from one machine to another. It provides strong authentication and secure communications over insecure channels. Often used as a replacement for insecure programs such as `telnet` or `rsh`. In Stonesoft, SSH can be used for remotely accessing the engine command line.

SSL VPN

A VPN technology that utilizes SSL encryption to secure users' remote access to specific applications. Allow authenticated users to establish secure connections to a limited number of specific internal services through a standard web browser ("clientless" access) or through a client application that allows a wider range of services.

Standby Mode

An operating state of a Stonesoft cluster that keeps one node online and the rest in standby, so that State Synchronization is done, but node does not process the traffic. If the online node is taken offline or fails, one of the standby nodes takes over the existing connections.

State Synchronization

The communication of connection tracking information between several Firewall nodes in a cluster. Can be either a full synchronization, where all connection tracking information is transferred to the other nodes of a cluster, or an incremental synchronization, where only the information on connections changed after the last synchronization are transferred. See also Connection Tracking (page 1267).

Static IP address

IP address that is typed in by a user or an administrator, and which does not change without their action.

Static NAT

NAT (Network Address Translation) where for each original address, there is a single, predefined translated address.

Static Routing

A form of routing that has permanent routes between networks programmed into every Routing Table.

Sub-Policy

A set of rules that are separated from the main policy, based on some common category, such as the service or the destination IP address. In this way, related rules can be grouped together to make the entire policy easier to understand. Because subrules are only processed if the general rule in the main policy matches, the overall processing time is improved.

Subtunnel

The actual tunnels that are combined logically within a multi-route VPN tunnel in a Stonesoft Multi-Link environment. They represent all possible routes that connect the end-points of the security gateways between which a Virtual Private Network (VPN) is formed. The individual subtunnels may connect the two gateways through different network links.

Symmetric Encryption

An Encryption mechanism that uses the same shared secret key for encrypting and decrypting messages. It is often referred to as symmetric bulk encryption since it processes large amounts of data rather quickly. Also known as conventional or secret key cryptography. There are two main types of symmetric encryption algorithms, bulk and stream encryption (also known as block ciphers and stream ciphers). Common symmetric algorithms are DES and 3DES. See also Asymmetric Encryption (page 1263).

Syslog

A standard protocol for exchanging logs between network components. Defined in RFC 5424.

System Summary

A panel in the System Status view that provides a general summary view of the current status of the monitored elements according to the component type.

T

Tag

An Element for organizing Situations. Tags can also be used in Inspection Rules, in the Situation cell, to represent all Situations marked with that Tag.

Takeover Period

The time interval during which the active nodes in a Security Engine cluster collaborate to redistribute the work load of a failed node.

Task

An Element that allows you to schedule commands to run automatically at a convenient time.

Template Policy

A combination of rules and Insert Points, which is used as a basis when creating policies or other template policies. Policies and template policies created from a particular template policy then inherit all the rules from that template policy and any of the template policies higher up in the inheritance hierarchy. The Inherited Rules cannot be edited within the inheriting policy. Used, for example, by high-privilege Administrators to restrict changes administrators with a lower Administrator Role can make to rules.

Temporary Filter

A log filter that is created from details of entries in the Logs View or the Connections view, and which is only available until the view is closed.

Terminate Action

An Inspection Rule parameter that stops or attempts to stop the connection matching to the rule according to the Action Option selected and the whether the Security Engine where the rule matching occurs is capable of stopping the connection.

Tester

A tool that can automatically run tests on Stonesoft engines to check system or network operation and take action based on the results of those tests.

Timeline

A tool in the Logs View that allows you to select and change the time range for the logs that are displayed.

ToS Flag

A data field in IP packet headers that provides a number representing the type of the service the packet is a part of. The ToS flag is used for Traffic Prioritization and is also know as DSCP (DiffServ Code Point).

Traffic Handler

The set of Network Elements used for inbound and outbound traffic management. Includes NetLinks, Outbound Multi-links, and Server Pools.

Traffic Management

The control, definition, and management of how packets or connections should flow through firewalls, routers, network links, VPNs or other gateway objects, based on load balancing, clusters, availability of links and more.

Traffic Prioritization

The process of assigning traffic a priority value, which is used to determine the order in which queued packets are sent forward, overriding the standard first-come-first-served operation of network devices. Used for assuring Quality of Service (QoS) for time-critical connections. Can be used together with Bandwidth Management or on its own. See also DSCP (DiffServ Code Point) (page 1269), QoS Class (page 1282) and QoS Policy (page 1282).

Transparent Access Control Mode

A Security Engine configuration in which the IPS Engine or Layer 2 Firewall examines Ethernet traffic according to the Ethernet Rules.

Transparent Proxy

A technique whereby a connection is routed to a proxy server, which then establishes a second connection to the original destination host, but the entire transaction takes place without notifying the user, or requiring the user to perform any additional actions.

Transport Protocol

Any protocol that communicates and functions on the transport layer of the TCP/IP protocol stack. These protocols function above the network layer, and are usually responsible for error correction, quality of service, and other characteristics not handled by the network layer. TCP, UDP, and IPsec are common examples of transport protocols.

Tunneling

A technology that enables one network to send its data through another, perhaps dissimilar, network. Tunneling works by encapsulating, or packaging, a network protocol within packets carried by the second network.

U

Use IPsec VPN Action

A Firewall Action parameter that directs traffic matching to the rule to a VPN. Can be either an Apply VPN Action or an Enforce VPN Action.

UDP Tracking

Information maintained by the Firewall engines to group together UDP requests and replies, handling them as a single virtual connection. See also Virtual Connection Tracking (page 1289).

User

An Element that defines an end-user in your network. Used for defining Authentication with or without Client-to-Gateway VPN access. Confer to Administrator (page 1261).

User Response

Defines additional notification actions for rule matches, such as redirecting access to a forbidden URL to a page on an internal web server instead.

UTM (Unified Threat Management)

A device that combines different types of traffic filtering in one physical appliance. The features offered in a UTM device vary greatly from vendor to vendor. The Stonesoft UTM comprises a Firewall, deep packet inspection (IDS), and anti-virus.

V

Virtual Adapter

A component of the Stonesoft IPsec VPN Client, or a third-party VPN client, that allows using a second, Virtual IP address for Virtual Private Network (VPN) traffic. Shown as a network adapter in the operating system.

Virtual Connection Tracking

A superset of UDP tracking, ICMP tracking, etc. A technology that is used by the Firewall engines for connectionless network protocols like UDP and ICMP. The Firewall engines keep track of virtual connections by grouping together packets that are related, based on information in the packet headers. See also Related Connection (page 1282).

Virtual Defragmentation

A procedure in which incoming packet fragments are collected. The packet is defragmented for processing by the engine, and refragmented before it is transmitted again. See also Fragmentation (page 1271).

Virtual IP address

A second IP address that is given to a VPN Client that has a Virtual Adapter enabled, and that is connecting to a security gateway using Client-to-Gateway VPN. A virtual IP address enables the use of certain services that require the client to have an IP address belonging to a specific address range, while enabling it to retain its primary IP address for maintaining other connections. The Virtual IP address for Stonesoft VPN Clients is always assigned by DHCP (Dynamic Host Configuration Protocol).

Virtual Local Area Network (VLAN)

A local area network which is defined through software in a switch or other networking device, rather than by the more traditional hardware division.

Virtual Private Network (VPN)

Refers to a confidential connection that is established through unsecured networks by the means of authentication, encryption, and integrity checking. The two major VPN technologies are IPsec (IP Security), which is better suited when a wide variety of network services and large traffic volumes are involved, and SSL VPN, which is used to provide access to a limited number of services to individual users without client-side device configuration.

Virtual Resource

An element that defines the set of resources on the Master Engine that are allocated to each Virtual Security Engine.

Virtual Security Engine

Logically-separate engines that run as virtual machines on a Master Engine.

VPN Client

Software that can be used to establish a Virtual Private Network (VPN) with a VPN gateway device to securely access remote resources over insecure networks.

VPN Profile

An element that defines the IPsec (IP Security)-related settings for one or more VPNs.

Vulnerability

An IPS element that contains information on a publicly known flaw that affects security of some system. Vulnerabilities are attached to Situations to provide you more information on what has happened when the Situation matches.

W

Web Filtering

A feature that compares the URLs that users attempt to open to a list of URLs to prevent users from intentionally or accidentally accessing most websites that are objectionable or potentially harmful.

Web Portal

Browser-based service that allows users to view logs, Policy Snapshots, and reports.

Whitelisting

The process of exempting specific traffic from being blocked by Blacklisting or Web Filtering.

Index

Stonesoft Guides

Administrator's Guides - step-by-step instructions for configuring and managing the system.

Installation Guides - step-by-step instructions for installing and upgrading the system.

Reference Guides - system and feature descriptions with overviews to configuration tasks.

User's Guides - step-by-step instructions for end-users.

For more documentation, visit
www.stonesoft.com/support/

Stonesoft Corporation

Itälahdenkatu 22 A
FI-00210 Helsinki
Finland

Tel. +358 9 476 711
Fax +358 9 4767 1349

Stonesoft Inc.

1050 Crown Pointe Parkway
Suite 900
Atlanta, GA 30338
USA

Tel. +1 770 668 1125
Fax +1 770 668 1131