



Sidewinder

Administration Guide

7.0.1.3H14

©1996–2016, Forcepoint LLC

All rights reserved.

10900-A Stonelake Blvd, Quarry Oaks 1, Suite 350, Austin TX 78759

Published 2016

The products and/or methods of use described in this document are covered by U.S. Patent Numbers 5,983,270; 6,606,659; 6,947,985; 7,185,015; 7,194,464 and RE40,187 and other patents pending.

This document may not, in whole or in part, be copied, photocopied, reproduced, translated, or reduced to any electronic medium or machine-readable form without prior consent in writing from Forcepoint LLC.

Every effort has been made to ensure the accuracy of this manual. However, Forcepoint LLC makes no warranties with respect to this documentation and disclaims any implied warranties of merchantability and fitness for a particular purpose. Forcepoint LLC shall not be liable for any error or for incidental or consequential damages in connection with the furnishing, performance, or use of this manual or the examples herein. The information in this documentation is subject to change without notice.

Trademarks

Forcepoint is a trademark of Forcepoint LLC. SureView, TRITON, ThreatSeeker, Sidewinder and Stonesoft are registered trademarks of Forcepoint LLC. Raytheon is a registered trademark of Raytheon Company. All other trademarks are the property of their respective owners.

Microsoft, Windows, Windows Server, and Active Directory are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Mozilla and Firefox are registered trademarks of the Mozilla Foundation in the United States and/or other countries.

eDirectory and Novell Directory Services are a registered trademarks of Novell, Inc., in the U.S and other countries.

Adobe, Acrobat, and Acrobat Reader are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States and/or other countries.

Red Hat is a registered trademark of Red Hat, Inc., in the United States and other countries. Linux is a trademark of Linus Torvalds, in the United States and other countries.

This product includes software distributed by the Apache Software Foundation (<http://www.apache.org>).

Copyright (c) 2000. The Apache Software Foundation. All rights reserved.

Other product names mentioned in this manual may be trademarks or registered trademarks of their respective companies and are the sole property of their respective manufacturers.

Contents

About this Guide	13
Who should read this guide	13
Where to find additional information	13
Online help	13
Reference materials	13
Typographical conventions	14

Introduction

1 Introduction to Sidewinder	17
About Forcepoint Sidewinder	17
The Type Enforcement environment	17
How Type Enforcement works	17
Type Enforcement's effects	19
Sidewinder operating characteristics	20
Burbs and network stack separation	20
Access control	21
Attack protection	22
Sidewinder deployment options	23
Routed mode	23
Transparent (bridged) mode	25
Hybrid mode	28
2 Administrator Basics	29
Managing your firewall	29
Understanding Sidewinder management	29
Admin Console basics	31
Using the Admin Console	34
Logging directly into the firewall	36
Configuring Admin Console access	37
Modifying the Admin Console rule	38
Configuring the Admin Console server	39
Restarting or shutting down the system	41
Rebooting or shutting down using the Admin Console	42
Rebooting or shutting down using a command line interface	43
Managing administrator accounts	44
Changing administrator passwords	47
Administering Sidewinder using Secure Shell	48
Configuring the Sidewinder as an SSH server	49
Configuring and using the Sidewinder as an SSH client	51
Tips on using SSH with the Sidewinder	53
Administering the Sidewinder using Telnet	54
Setting up an internal (trusted) Telnet server	54
Setting up an external Telnet server	54
Connecting to the Sidewinder using Telnet	54

Policy

3 Policy Configuration Overview	57
About policy configuration	57
A brief guide to planning your policy	59
Using groups to simplify policy management	59
Examining your policy using the Firewall Policy Report	61

About creating rules	61
4 Network Objects and Time Periods	63
Creating network objects	63
About the Network Objects: Domain window	65
About the Network Objects: Geo-Location window	66
About the Network Objects: Host window	67
About the Network Objects: IP Address window	68
About the Network Objects: IP Range window	69
About the Network Objects: Netmap window	70
About the Network Objects: Subnet window	72
About the Network Objects: Netgroup window	73
Creating time periods	76
5 Authentication	79
Understanding authentication	79
Who gets authenticated	79
Weak and strong authentication	80
Types of authentication methods	80
Alternate authentication methods	81
Authentication scenario	81
Configuring an authenticator	83
Setting up Passport authentication	85
Setting up standard password authentication	90
Setting up LDAP authentication	91
Setting up CAC authentication	95
Setting up Windows domain authentication	96
Setting up RADIUS authentication	97
Setting up SafeWord authentication	100
Telnet and FTP considerations	102
Setting up users to change their own passwords	103
Create a change password rule	103
How users can change their own password	103
Authenticating groups from an external group source	105
Authenticating groups from an internal group source	107
About the Users and User Groups tab	108
About the Create New User/Group window	110
About the Group Objects: Group Information tab	111
About the Group Objects: User Group Membership tab	112
About the User Objects: User Information tab	113
About the User Objects: User Password tab	114
6 Content Inspection	117
About content inspection	117
Configuring IPS inspection	117
Understanding signature-based IPS	118
Adding IPS inspection to rules	120
About signature file updates	121
Using IPS with other Sidewinder attack protection tools	121
Configuring a response mapping	122
Configuring a signature group	125
Managing signatures	130
Configuring IPS signature file updates	132
Configuring virus scanning services	134
Configuring virus scanning signature updates	135
Configuring the advanced virus scanning features	137
About TrustedSource	138
Using TrustedSource on a Sidewinder	138
Configuring TrustedSource	142
Updating the Geo-Location database	146
Configuring SmartFilter for Sidewinder	147

7	Services	149
	About services	149
	Using the main Services window	150
	Create and modify services	152
	Create and modify service groups	153
	Configuring proxy agents and services	154
	About proxy agents and services	154
	Configuring proxy agent properties	157
	Configuring proxy service properties	158
	Selecting the appropriate proxy agent	160
	Configuring packet filter agents and services	163
	About packet filter agents and services	163
	Selecting the appropriate packet filter agent	168
	Configuring the TCP/UDP packet filter agent properties	168
	Configuring packet filter service properties	170
	Configuring server agents	172
	About server agents	172
	Configuring server agent properties	172
	Selecting the appropriate server	177
	Configuring additional proxy agent properties	178
	Configuring URL translation on the HTTP proxy agent	178
	Using the SSH proxy agent	182
	Modifying the FTP proxy agent's accepted server responses	187
	Configuring the SMTP proxy agent to strip source routing	188
	Using the T.120 and H.323 proxy agents together	189
8	Application Defenses	195
	Understanding Application Defenses	195
	Creating HTTP or HTTPS Application Defenses	198
	Configuring the HTTP/HTTPS: Enforcements tab	198
	Configuring the HTTP/HTTPS: HTTP URL Control tab	202
	Configuring the HTTP: FTP URL Control tab	203
	Configuring the HTTP/HTTPS: HTTP Request tab	204
	Configuring the HTTP/HTTPS: HTTP Reply tab	205
	Configuring the HTTP/HTTPS: MIME/Virus/Spyware tab	206
	Configuring the HTTP/HTTPS: Content Control tab	209
	Configuring the HTTP/HTTPS: SmartFilter tab	211
	Configuring the HTTP/HTTPS: Connection tab	211
	Creating Mail (Sendmail) Application Defenses	212
	Configuring the Mail (Sendmail): Control tab	213
	Configuring the Mail (Sendmail): Size tab	213
	Configuring the Mail (Sendmail): Keyword Search tab	215
	Configuring the Mail (Sendmail): MIME/Virus/Spyware tab	217
	Creating Mail (SMTP proxy) Defenses	220
	Configuring the Mail (SMTP proxy): General tab	221
	Configuring the Mail (SMTP proxy): Commands tab	224
	Configuring the Mail (SMTP proxy): Header filters tab	225
	Creating Citrix Application Defenses	227
	Configuring the Citrix: Enforcements tab	227
	Configuring the Citrix: Filters tab	227
	Creating FTP Application Defenses	228
	Configuring the FTP: Enforcements tab	228
	Configuring the FTP: Command Filter tab	229
	Configuring the FTP: Virus/Spyware tab	230
	Creating IOP Application Defenses	234
	Configuring the IOP: Filter tab	234
	Creating T.120 Application Defenses	235
	Configuring the T.120: General tab	235
	Configuring the T.120: Filter tab	235
	Creating H.323 Application Defenses	236
	Configuring the H.323: General tab	236
	Configuring the H.323: Filter tab	236
	Creating Oracle Application Defenses	237

Configuring the Oracle: Enforcements tab	238
Configuring the Oracle: Service Name (SID) tab	239
Creating MS SQL Application Defenses	239
Creating SOCKS Application Defenses	240
Configuring the SOCKS: SOCKS 5 Filter tab	240
Configuring the SOCKS: Connection tab	241
Creating SNMP Application Defenses	242
Configuring the SNMP: Filter tab	242
Configuring the SNMP: v1 tab	243
Creating SIP Application Defenses	244
Configuring the SIP: General tab	245
Configuring the SIP: Media Filters tab	246
Creating SSH Application Defenses	247
Configuring the SSH: Channels tab	247
Configuring the SSH: Client Authentication tab	248
Configuring the SSH: Client Advanced tab	249
Configuring the SSH: Server Advanced tab	250
Creating Packet Filter Application Defenses	252
Configuring the Packet Filter: General tab	252
Configuring the Packet Filter: Advanced tab	253
Configuring Application Defense groups	254
9 Rules	257
About rules	257
Condition rule elements	258
Action rule elements	260
Example of a simple rule	260
Using NAT and redirection in rules	263
Understanding and configuring NAT	263
Understanding and configuring redirection	264
Viewing and ordering rules and rule groups	265
Ordering rules within your policy	265
About the default firewall policy	268
Creating an alternate policy	268
Using the main Rules window	269
Customizing the main Rules window view	271
Viewing and exporting your active policy	272
Creating, modifying, and duplicating rules	273
Creating and modifying rule groups	279
Viewing and modifying rule elements	280
Services	280
Time periods	281
Source burbs, endpoints, and NAT	281
Destination burbs, endpoints, and redirection	285
Application Defenses	287
IPS response mapping and signature groups	288
Authentication	289

Monitoring

10 The Dashboard	293
Monitoring Sidewinder status using the dashboard	293
Viewing device information	295
Viewing network traffic information	300
Viewing IPS attack and system event summaries	307
Understanding audit event severities	307
Viewing the summary statistics	308
11 Auditing	311
Understanding the Sidewinder audit process	311
Audit components	312
Audit file names	314

Understanding audit messages	314
Tools for viewing and customizing audit events	314
Supported log file formats	315
Viewing audit information	316
Filtering audit data	318
Viewing and transferring audit records	326
Managing log files	331
Creating or modifying an export entry	332
Signing export files	334
Exporting and rolling log files	334
Monitoring disk space using cron jobs	335
Identifying changes using change tickets	335
Exporting audit data to Firewall Reporter and syslog servers	336
12 Service Status	339
Understanding processes that control server status	339
daemon	339
Network Services Sentry (NSS)	340
Viewing service status	341
Viewing a service's process information	344
13 IPS Attack and System Event Responses	347
Understanding attack and system event responses	347
Creating IPS attack responses	348
Modifying an IPS attack response	349
Configuring the e-mail response settings	353
Creating system responses	354
Modifying a system response	355
Configuring the e-mail settings	358
Ignoring network probe attempts	359
Sidewinder SNMP traps	360
14 Network Defenses	363
Viewing Network Defense information	363
Configuring the TCP Network Defense	367
Configuring the IP Network Defense	369
Configuring the UDP Network Defense	371
Configuring the ICMP Network Defense	373
Configuring the ARP Network Defense	375
Configuring the IPsec Network Defense tab	377
Configuring the IPv6 Network Defense tab	379
15 The SNMP Agent	381
Understanding SNMP options	381
Overview of Sidewinder as a managed node	382
Communicating with an SNMP management station	382
About Sidewinder SNMP traps	383
About Sidewinder SNMP MIBs	384
About the management station	385
Setting up the SNMP agent on Sidewinder	386
Configuring the SNMP agent	386
Creating a rule to allow access to the SNMP agent	389
Sending SNMP traffic through Sidewinder	390

Networking

16 Burbs, Interfaces, and Quality of Service	393
Configuring burbs	393
Creating or modifying a burb	395
Creating or modifying a burb group	395
Configuring interfaces	397
About the Interfaces: Interface Configuration tab	399

About the Interfaces: NIC and NIC Group Configuration tab	401
Creating interfaces	403
Configuring Quality of Service	419
Example QoS scenarios	426
17 Routing	427
About routing on Sidewinder	427
Configuring static routes	428
Configure default routes	430
Check route status and reset the default route	434
Configure other static routes	435
RIP on Sidewinder	436
Configuring RIP (ripd)	437
Viewing and comparing ripd configurations	443
OSPF on Sidewinder	445
Configuring OSPF (ospfd)	446
Viewing and comparing OSPF configurations	449
OSPF IPv6 on Sidewinder	450
BGP on Sidewinder	453
Configuring BGP (bgpd)	453
Viewing and comparing BGP configurations	456
PIM-SM on Sidewinder	457
Configuring PIM-SM (pimd)	457
Viewing PIM-SM configurations	466
Dynamic routing in HA clusters	467
Troubleshooting dynamic routing issues	467
18 DNS (Domain Name System)	469
What is DNS?	469
About transparent DNS	470
About firewall-hosted DNS	471
Configuring transparent DNS	474
Configuring firewall-hosted DNS servers	475
Configuring the Server Configuration tab	477
Configuring the Zones tab	480
Configuring the Master Zone Attributes tab	482
Configuring the Master Zone Contents tab	487
Reconfiguring DNS	490
Reconfiguring transparent DNS	490
Reconfiguring single server hosted DNS	492
Reconfiguring split server hosted DNS	492
Manually editing DNS configuration files	495
DNS message logging	495
19 E-mail	497
Overview of mail on Sidewinder	497
About transparent mail (SMTP proxy)	497
About Sidewinder-hosted mail (sendmail)	498
Setting up and reconfiguring mail	498
Understanding sendmail on Sidewinder	500
Using sendmail on Sidewinder	500
Mail filtering services on Sidewinder	502
Editing sendmail files on Sidewinder	502
Configuring advanced sendmail features	506
Configuring sendmail to strip message headers	507
Configuring sendmail to use the RealTime Blackhole list	507
Allowing or denying mail on a user basis	509
Configuring sendmail to hide internal e-mail addresses	510
Enabling Sendmail TLS	511
Managing mail queues	511
Viewing the mail queue	511
Changing how long a message waits between delivery attempts	512
Manually attempting to deliver queued messages	512

Changing how long a message waits before it is returned to its sender	513
Receiving mail sent by Sidewinder	514
Setting up e-mail aliases for administrator accounts	514
Viewing administrator mail messages on Sidewinder	514

20 Virtual Private Networks 517

About the Sidewinder VPN solution	517
Protecting your information	518
What are encryption and authentication?	518
About IPsec keys	518
Planning your VPN	519
Choosing the appropriate VPN attributes	519
Choosing the appropriate authentication type	522
Ordering VPN definitions	525
Restricting VPN access with a virtual burb	525
Creating VPN policy	533
Setting up the ISAKMP service	533
Configuring client address pools	534
Creating and using a virtual burb with a VPN	534
VPN user interface reference	535
Managing VPN definitions	535
Managing client address pools	546
Managing the ISAKMP server	552
Example VPN Scenarios	555
Scenario 1: Firewall-to-firewall VPN using a shared password	555
Scenario 2: Simple deployment of remote users	557
Scenario 3: Large scale deployment of clients	560

Maintenance

21 General Maintenance Tasks 567

Setting the system date and time	568
Setting the date and time	568
Understanding Network Time Protocol	569
Configuration file backup and restore	573
About the Configuration Backup: Configuration Backup tab	574
About the Configuration Backup: Configuration Restore tab	580
About the Configuration Backup: Schedule tab	584
Activating the Sidewinder license	586
Licensing from a firewall connected to the internet	587
Licensing from a firewall on an isolated network	588
Configuring the Firewall License tabs	589
Protected host licensing and the Host Enrollment List	594
Software management	597
Understanding software management	597
About the Software Management: Manage Packages tab	599
About the Software Management: Download Packages tab	603
About the Software Management: Rollback tab	605
Editing files	606
About editing Sidewinder files	606
Using the File Editor	607
Checking file and directory permissions (ls command)	611
Changing a file's type (chtype command)	611
Creating your own scripts	611
Registering with Forcepoint Sidewinder Control Center	612
Enforcing FIPS	614
Enabling hardware acceleration	615
Configuring UPS	616

22 Certificate/Key Management 619

About Certificate/Key Management	619
Understanding Distinguished Name syntax	619

Selecting a trusted source	.622
Managing firewall certificates	.623
Creating firewall certificates	.624
Importing firewall certificates	.626
Loading manual firewall certificates	.627
Assigning new certificates for Admin Console services	.628
Managing certificate authorities	.629
Adding certificate authorities	.631
Exporting certificate authorities	.631
Managing VPN certificates	.633
Configuring the certificate server	.634
Configuring and displaying remote identities	.636
Configuring and displaying remote certificates	.638
Exporting certificates	.643
Exporting only the certificate	.644
Exporting both the certificate and private key	.644
Managing SSH keys	.645
About the Create New SSH Key window	.646
About the Import SSH Key window	.647
About the Export SSH Key window	.648

23 High Availability 649

How Sidewinder High Availability works	.649
About HA redundancy	.650
About shared cluster addresses	.650
About HA configuration options	.651
Load-sharing HA	.651
Failover HA	.653
Configuring HA	.653
Ensure HA requirements are met	.654
Configure the heartbeat interfaces	.654
Add the first Sidewinder to a new HA cluster	.655
Add a reservation for the second firewall in the HA cluster	.661
Join a Sidewinder to an existing HA cluster	.662
Post-configuration tasks	.663
Understanding the HA cluster tree structure	.665
Managing an HA cluster	.667
Modifying HA common parameters	.667
Modifying HA local parameters	.673
Scheduling a soft shutdown for a load-sharing HA cluster Sidewinder	.675
Re-establishing an HA cluster if a cluster member fails	.676
Restarting an HA cluster	.676

Troubleshooting

A Basic Troubleshooting 679

Troubleshooting rules	.679
Failed connection requests	.679
Monitoring allow and deny rule audit events	.681
Active rules and DNS	.682
Troubleshooting logging in	.682
Restoring access to the firewall	.682
Changing a forgotten password	.682
Manually clearing an authentication failure lockout	.683
Changing authentication requirements for emergency maintenance mode	.683
Troubleshooting system status	.684
Troubleshooting network status	.685
Checking network status using the Admin Console	.685
Checking network status using the command line	.692
Troubleshooting licensing problems	.695
Troubleshooting High Availability	.696
Viewing configuration-specific information	.696

Viewing status information	698
Identifying load-sharing addresses in netstat and ifconfig	700
Troubleshooting NTP	701
Troubleshooting VPNs	703
Troubleshooting transparent (bridged) mode	703
B Re-install and Recovery Options	705
About re-install and recovery	705
Recovery options	706
Configuration restore	706
Uninstall	706
Rollback	706
Disaster recovery	706
Re-install options	708
Re-installing your firewall from the virtual CD	709
Re-installing your firewall from a CD-ROM	710
Re-installing your firewall from a USB drive	711
Glossary	715
Index	719

About this Guide

Who should read this guide

This guide is intended for a Forcepoint™ Sidewinder® administrator. You should read this guide if you are responsible for configuring and managing a Sidewinder.

This guide assumes you have:

- A working knowledge of UNIX and Windows operating systems.
- A basic understanding of system administration.
- A working knowledge of the Internet and its associated terms and applications.
- An understanding of networks and network terminology, including TCP/IP protocols.

Where to find additional information

Sidewinder documentation in .pdf format is available on our web site. You can see the latest information regarding Sidewinder and other products. Open a browser and go to <https://support.forcepoint.com>.

Table 1 Summary of Sidewinder documentation

Document	Description
<i>Setup Guide</i>	Steps you through setting up your initial Sidewinder configuration.
<i>Administration Guide</i>	This is the guide you are currently reading. It provides complete administration information on all firewall functions and features. You should read this guide if you are responsible for configuring and managing a firewall.
Online help	Online help is built into the Sidewinder software. The Forcepoint™ Sidewinder® Quick Start Wizard provides help for each configuration window. The Admin Console program provides detailed context-sensitive online help.
Application notes	Provides detailed instructions for setting up specific configurations, such as setting up the firewall to work with another vendor's product or environment.
Knowledge Base	Supplemental information for all other Sidewinder documentation. Articles include helpful troubleshooting tips and commands. All manuals and application notes are also posted here.

Online help

The Sidewinder graphical user interface (known as the Admin Console) provides comprehensive online help. To access online help, click the help icon in the toolbar.

Man (or “manual”) pages provide additional help on firewall-specific commands, file formats, and system routines. To view the available information for a specific topic, enter one of the following commands:

```
man -k topic
```

or

```
apropos topic
```

where *topic* is the subject that you want to look up.

Reference materials

If you are new to system administration, you may find the following resources useful:

- *UNIX System Administration Handbook*, 3rd Edition, by Nemeth, et al. (Prentice Hall).
- *Managing Internet Information Services* by Liu, et al. (O'Reilly and Associates, Inc.)

- A standard reference on computer security is *Firewalls and Internet Security* by Cheswick and Bellovin (Addison-Wesley).
- For network management information, see *TCP/IP Network Administration* by Craig Hunt (O'Reilly & Associates, Inc.).
- For information on handling mail on UNIX networks, see *Sendmail* by Bryan Costales, with Eric Allman and Neil Rickert (O'Reilly & Associates, Inc.).
- For Domain Name System information, see *DNS and Bind* by Cricket Liu and Paul Albitz (O'Reilly & Associates, Inc.).
- For information about Internet Review for Comment (RFC) documents, refer to the following web site:

<http://www.ietf.org/rfc.html>

Note: Some of these resources are referenced throughout this guide.

Typographical conventions

This guide uses the following typographic conventions:

Table 2 Conventions

Convention	Description
Courier bold	Identifies commands and key words you type at a system prompt Note: A backslash (\) signals a command that does not fit on the same line. Enter the command as shown, ignoring the backslash.
<i>Courier italic</i> < <i>Courier italic</i> > <i>nnn.nnn.nnn.nnn</i>	Indicates a placeholder for text you type When enclosed in angle brackets (< >), identifies optional text Indicates a placeholder for an IP address you type
Courier plain	Used to show text that appears on a computer screen
<i>Plain text italics</i>	Identifies the names of files and directories Used for emphasis (for example, when introducing a new term)
Plain text bold	Identifies buttons, field names, and tabs that require user interaction
[]	Signals conditional or optional text and instructions (for example, instructions that pertain only to a specific configuration)
Caution	Be careful—in this situation, you might do something that could result in the loss of data or an unpredictable outcome.
Note	Helpful suggestion or a reference to material not covered elsewhere in the manual
Security Alert	Information that is critical for maintaining product integrity or security
Tip	Time-saving actions; may help you solve a problem

Note: The IP addresses, screen captures, and graphics used within this document are for illustration purposes only. They are not intended to represent a complete or appropriate configuration for your specific needs. Features may be enabled in screen captures to make them clear; however, not all features are appropriate or desirable for your setup.

SECTION 1

Introduction

Chapter 1, Introduction to Sidewinder

Chapter 2, Administrator Basics

1 Introduction to Sidewinder

Contents

[About Forcepoint Sidewinder](#)

[The Type Enforcement environment](#)

[Sidewinder operating characteristics](#)

[Sidewinder deployment options](#)

About Forcepoint Sidewinder

Forcepoint Sidewinder allows you to connect your organization to the Internet while protecting your network from unauthorized users and attackers, while also protecting internal users as they access the Internet. It combines an application-layer firewall, IPsec VPN capabilities, Web filtering, global-reputation-based filtering, anti-virus/anti-spyware filtering engine, and SSL decryption into one Unified Threat Management (UTM) security appliance, designed to offer centralized perimeter security.

Sidewinder provides a high level of security by using SecureOS an enhanced UNIX operating system that employs Type Enforcement security technology. SecureOS removes the inherent security risks often found in a network application running on non-security focused commercial operating systems, resulting in superior network security and no emergency security patches to apply.

The firewall prevents host identification masquerading (IP spoofing), making it very difficult for attackers to infiltrate your protected networks. It also offers advanced authentication and encryption software. Encryption allows authorized users on the Internet access to your protected network without fear of attackers eavesdropping (IP sniffing) or stealing access credentials and other valuable information.

Sidewinder allows public services such as e-mail, a public file archive (FTP), and Web (HTTP/HTTPS) access while protecting the other computers on your protected networks. It also provides powerful configuration options that allow you to control access by your employees to almost any publicly available service on the Internet.

The Type Enforcement environment

As mentioned earlier, Sidewinder runs on SecureOS, a version of BSD enhanced with a security technology called Type Enforcement. For the most part, Type Enforcement does not require any extra effort on your part. The following subsections describe areas that affect how you use the system and access files that you should be aware of.

[How Type Enforcement works](#)

[Type Enforcement's effects](#)

How Type Enforcement works

In most UNIX operating systems, logging in as super-user (root) gives you access to all system files; an intruder who knows how to acquire root privileges can access any files or applications on a system. In addition, UNIX does not have tight control over how data files are shared among the processes running on a system. This means that an intruder who managed to break into one area of a system, such as e-mail, may be able to easily gain access to other files on the system. Sidewinder Type Enforcement software is designed to plug these security holes. This is done by using the following mechanisms (each of the mechanisms is described below):

- Provides maximum network protection
- Provides Type Enforced domain processes
- Controls Type Enforced attributes applied to files and sockets
- Controls inter-domain operations, such as signals
- Controls access to system calls
- Controls the files a process can access

Maximum network protection


Type Enforcement technology provides network security protection that is unique to the industry. By using Type Enforcement within the operating system, Sidewinder provides the highest level of security.

Type Enforcement is based on the security principle of least privilege: any program executing on the system is given only the resources and privileges it needs to accomplish its tasks. On Sidewinder, there is no concept of a root super-user. Type Enforcement controls all interactions between domains and file types. Domains must have explicit permission to access specific file types, communicate with other domains, or access system functions. Any attempts to the contrary fail as though the files do not exist.

Type Enforcement domain processes

A standard UNIX system separates processes with user and group identities. Therefore, UNIX identities can be completely subverted by users who obtain root privileges. Sidewinder prevents this by providing separate, Type-Enforced domains for each process running on the system. Type-enforced domains provide more intricate control over what each process is allowed to do, as shown in [Figure 1](#).

Figure 1 Example of Sidewinder domain separation structure



Type Enforced attributes

When an administrator initially logs into a Sidewinder at a command line prompt, they are automatically placed in the *User* domain, which allows no access to sensitive files. An administrator may then switch to their defined administrative role's domain using the `srole` command (for Admn) or `srole adminro` (for AdRO). The Admn domain allows an administrator to access to all administrative functions. The AdRO domain allows read-only access to the system configuration areas, as well as the ability to generate reports. An administrator with read-only access cannot make system modifications.

This guide assumes that most commands will be issued by administrators with read/write access, and therefore only includes the `srole` command. If you are a read-only administrator and have reason to access the command line, always use `srole adminro` instead of `srole` alone.

For information on assigning administrator roles, see [Managing administrator accounts on page 44](#).

Inter-domain operations

Interactions between domains, such as signalling, are also controlled by Type Enforcement. For example, a process running in the SMTP domain cannot send a signal to the Telnet server running in the Telnet domain.

Access to system calls

A typical UNIX system has many privileged system calls that could enable malicious users to access the kernel directly and compromise the system. The firewall solves this problem with a set of flags for each domain that indicate which system calls can be made from that domain.

Files available to a process

Process-to-file access is controlled by a Domain Definition Table that maps out the various classes of data files and processes that may be running on the firewall. The table specifies which process domains can access different types of files and what type of access is allowed (such as read/write/execute). This table cannot be circumvented.

Your system is pre-configured so that domains have access only to the files they need. The Domain Definition Table cannot be changed while the Operational kernel is running. This prevents intruders from tricking the kernel into modifying the table. Also, Type Enforcement prevents intruders from installing software that may be used to circumvent Sidewinder security mechanisms.

The backup and restore functions on your system have been modified to be aware of Type Enforcement. When you restore files, they are automatically restored with the correct Type Enforcement properties.

Type Enforcement's effects

The previous section outlined how Type Enforcement works. Listed below are the major ways in which Type Enforcement affects you and other users:

- Non-administrative users will not be aware of Type Enforcement unless they try to perform unauthorized activities.
- The concept of a super-user who can have complete system control does not exist. The “root” account has no special privileges. The Admin role operating in the Admn domain has access to most system files, but is still not as powerful as root on a standard UNIX system.
- Domains make it difficult for an intruder to do damage. Breaking into the domain in which an application is executing does not provide access to the files required for administering that application.

Sidewinder operating characteristics

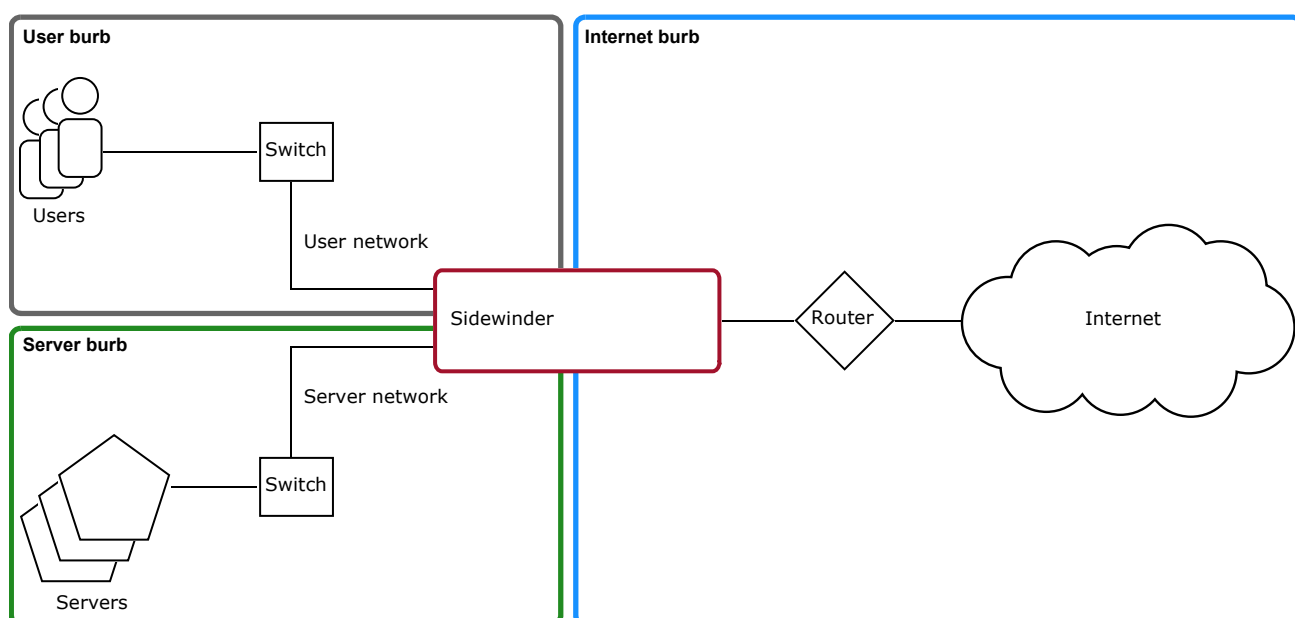
This section lists additional significant differences between Sidewinder and a standard UNIX system.

Burbs and network stack separation

While installing or managing a Sidewinder, you will notice the use of the term *burb*. *Burb* is a term that refers to an interface and all the systems it connects. Each burb must have a unique name. Unless you specify custom burb names during initial configuration, the two initial burbs are named *internal* and *external* by default.

As an example of how burbs are used, suppose your organization has two internal (protected) networks that need to be connected to the external network (Internet), but the corporate security policy requires that there be limited or no information flow between the two internal networks. In this scenario, you would configure three burbs for your Sidewinder, as shown in Figure 2. The security policy must be defined to enforce the required control over information flow between the two internal burbs and between the external burb and the individual internal burbs, while also protecting the internal burbs from unauthorized access from the Internet.

Figure 2 Multiple Type Enforced areas (burbs)




One of the unique aspects of the SecureOS is the use of multiple logical network stacks to strengthen the enforcement of the inter-burb aspects of the system security policy. A network stack consists of different layers of software responsible for different aspects of the communications. For example, one layer checks a message's routing information to ensure that it is transmitted to the correct network. Normal computing systems and firewalls that operate on an unsecured OS have only one network stack.

The SecureOS includes modifications that provide stronger separation of communication between different burbs. There are checks at all layers of the software to ensure that the network stack data from one burb is not mixed with or impacted by data associated with another burb. This logical separation of the network stacks by the security burb is augmented by Type Enforcement security policy, which is integral to SecureOS. It controls all operational aspects of the system, including enforcement of the separation data processing by the security burb. This ensures that information passes from one burb to another only if the network security policy says the specific information flow is allowed.

Figure 3 shows this logical network separation involved in the transfer of data between the network stacks associated with each burb. Before a process can interact with a network stack, the Type Enforcement security policy must indicate that the process is allowed to interact with that burb's network stack. The degree to which the firewall inspects a packet is determined by the agent processing the packet.

Figure 3 Logical network protocol stacks provide network separation



Access control

In Sidewinder, the rule set determines what traffic is permitted into and through the firewall and what is denied. Each rule requires a service. A service associates a traffic's transport layer with a specific agent that is responsible for managing the service's traffic. The transport layer information includes elements such as the protocol, the ports, and the connection or session timeouts. There are three types of agents: proxy, packet filter, and server.

Note: See the Policy section for details information on policy configuration, rules, and services.

Proxy agents and services

Sidewinder uses special programs called *proxy agents* to forward application data between two burbs, such as your internal network and the Internet. Proxy agents essentially provide a go-between that can communicate with the burbs on the firewall. For example, when a user on an internal burb tries to establish an Internet connection, the firewall intercepts the connection attempt and opens the connection on the user's behalf. All Internet connections are made by the firewall so that the internal network never communicates directly with the Internet burb. For some proxy agents, you can configure transparency on a per-service basis. For transparent connections, the client is unaware of the firewall. The firewall is implicitly included in the path based on routing. For non-transparent, the client is aware of the firewall and explicitly connects to the firewall.

Sidewinder supports HTTP, Telnet, and many other TCP-based proxies. The firewall also supports proxies for routing SNMP, NTP, DNS, and other types of services that require UDP transmissions. You can also create your own special proxies for other services. In addition, the firewall provides proxies that use multiple TCP and/or UDP sessions such as FTP, RealMedia, and Oracle SQLNet.

Most proxy agents are disabled by default. When you use a proxy service in a rule, the firewall automatically enables that proxy service's agent in the rule's source burb. That allows traffic to flow from the source to its destination. For example, you can configure rules that allow all internal users to access all Internet web sites, or you can prohibit users from accessing the web from specific internal systems or from accessing specific web sites. In addition, proxy rules can be configured to perform Network Address Translation (NAT) and redirection. Enabling NAT rewrites the source address of the packet, while enabling redirection rewrites the destination address.

Filter agents and services

You can configure the firewall to securely forward IP packets between networks using filter services in rules. Unlike proxy agents, which operate at the application layer and in most cases on TCP or UDP traffic, filters operate directly on IP packets allowing non-TCP/UDP (as well as TCP/UDP) traffic to pass between the networks. For example, with a filter service you can pass encrypted VPN sessions through the firewall.

Filtering services work by inspecting many of the fields within a packet, including the source and destination IP address, port, and protocol. Each packet that arrives at the firewall will be inspected and compared to an enabled filter rule that you have configured. Packets that match an allow rule will then be forwarded to the destination network.

You can configure filtering services to inspect TCP, UDP, and many other protocols. With TCP, UDP, and ICMP, the firewall can actively track individual sessions by performing *stateful inspection*. This ensures that only packets valid for a new session or a portion of an existing session are sent on to the final destination. In addition, packet filter rules can be configured to perform Network Address Translation (NAT) and redirection. Enabling NAT rewrites the source address of the packet, while enabling redirection rewrites the destination address.

Server agents and services

Sidewinder servers provide a variety of system functions, but generally do not pass traffic between burbs. Rules that allow access to a Sidewinder server typically have the same source and destination burbs. Therefore, proxy agents are not used to control an external (Internet) user's access to the external side of the Sidewinder. For example, when an external user accesses an SSH server that you have made publicly available on the external side of the firewall, there will be no proxy agent to intervene. For users on the Internet, proxy agents are only used when they cross burb boundaries to access systems in an internal burb.

Attack protection

The first step in protecting your network is creating a rule set based on a least-permissions philosophy and using the application-aware proxy agents to pass traffic. The next step is to use Sidewinder attack protection to defend against attacks in both allowed and denied traffic. The firewall has multiple layers that work together to protect against known and unknown attacks. Some of these defenses occur automatically, and some of them must be configured. The following sections explain the different options.

Network Defenses

Sidewinder is pre-configured to block an extensive list of suspicious traffic at the data link, network, and transport layers. Packets that do not adhere to their protocol standards are always dropped, as are packets that match known attack configurations.

Application Defenses

Application Defenses offer customizable protection at the application layer. The defenses can be used to enforce RFC (Request for Comments) standards and allowed parameters. Configurable parameters include headers, commands, versions, and file sizes. You can use these controls to deny any parameters that are not essential to your business needs and to minimize your network's attack surface; the fewer the number of parameters allowed into your network, the fewer parameters an attacker can use to attack. The controls can also provide the following key inspection services:

- Anti-virus filtering
- Reputation-based filtering

Note: The listed services are premium features.

Signature-based Intrusion Prevention Services

The Sidewinder Intrusion Prevention Service uses signature-based files to detect and prevent known network-based intrusion attacks, such as hacker-generated exploits and protocol anomalies. IPS can be added to rules to inspect allowed, incoming traffic for these attacks as the traffic enters the firewall. If an attack is detected, the rule handles the attack according to the configured response. Response options range from completely ignoring the traffic to blackholing all traffic sent from the originating host. This attack protection is particularly valuable when you cannot minimize your attack surface because your organization requires services with known vulnerabilities.

Note: This is a premium feature.

IPS Attack Responses

Even attacks that are not allowed through the firewall can cause problems if allowed to continue. For this reason, Sidewinder has IPS Attack Responses, which can be configured to notify administrators when audit events are generated by suspicious traffic. If a specified attack audit occurs a certain number of times in a given time period, the firewall can alert an administrator, blackhole all traffic from the IP address originating the attack, or both. Being aware of attempted attacks is an important part of maintaining your network's security.

Sidewinder deployment options

The internal and external network Sidewinder interfaces are defined during initial configuration. However, you can configure additional interfaces to suit the needs of your network infrastructure. The firewall can be used as:

- A gateway between your internal network and the Internet.
- A gateway between any networks with different security needs.
- A transparent firewall inside a single network.
- Any combination of the above.

For traffic to pass through your Sidewinder, it must arrive on an interface and leave on a different interface. The relationship between configured interfaces can be classified as follows:

- **Routed** – A Sidewinder interface is connected to each unique network, and the firewall allows traffic to pass between the networks like a router, enforcing your security policy.

See [Routed mode](#) for more information.

- **Transparent (bridged)** – Two Sidewinder interfaces are connected inside a single network and bridged to form a transparent interface. Traffic passes through the firewall like a switch, allowing you to enforce security policy inside the network without re-addressing the network.

See [Transparent \(bridged\) mode](#) for more information.

Note: Sidewinder supports only one configured transparent interface (bridge) at a time.

The routed and transparent modes are not exclusive—your Sidewinder can be simultaneously configured with a single bridge and additional routed interfaces. For more information, see [Hybrid mode](#).

Routed mode

In routed mode, your Sidewinder is deployed at the intersection of multiple networks.

- The firewall is connected to each network by a network interface.
- Each firewall interface must be assigned a unique IP address in the connected subnet.
- The protected networks must be unique—each network must be a different subnet.
- Hosts in a protected network communicate with other networks by using the firewall's IP address as their gateway.
- Each firewall interface is assigned to a unique burb. When traffic attempts to cross from one burb to another, the configured security policy is enforced.

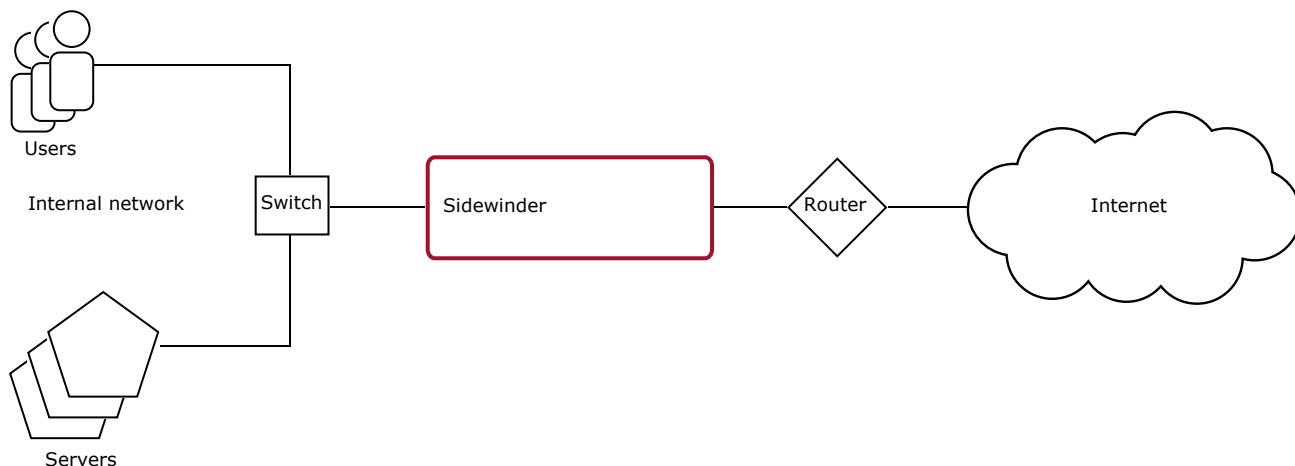
This section describes the following deployment scenarios:

- [Protecting a single network](#)
- [Protecting multiple networks](#)

Protecting a single network

Figure 4 depicts Sidewinder protecting the internal network from the Internet. This configuration uses two network interfaces. To reach the Internet, hosts on the internal network route traffic to the firewall.

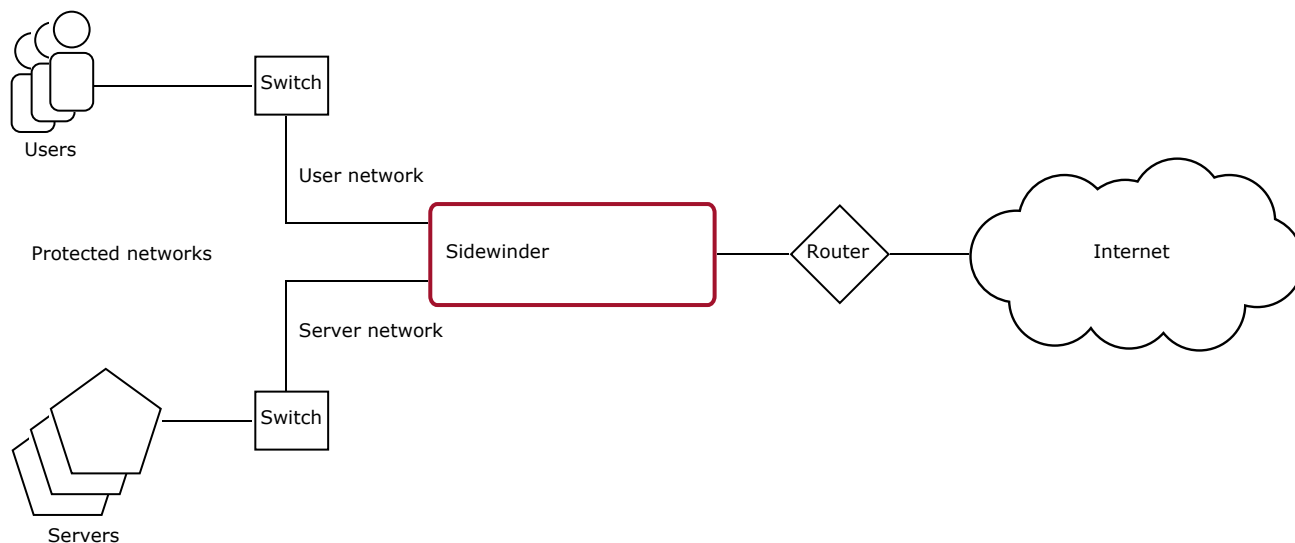
Figure 4 Protecting a single network



Protecting multiple networks

Figure 5 depicts Sidewinder protecting two otherwise separate networks, the user network and the server network, from each other and from the Internet. This configuration uses three network interfaces. To reach the Internet or one of the other protected networks, hosts route traffic to the firewall.

Figure 5 Protecting multiple networks



Transparent (bridged) mode

In transparent (bridged) mode, your Sidewinder is deployed inside a single network.

- A transparent interface is made up of two interfaces that are:
 - Connected inside the same network.
 - Assigned to unique burbs.

[Table 3](#) shows the default Sidewinder interface configuration. These interfaces, or any other two interfaces, can be used to configure a transparent interface.

Table 3 Standard interfaces

User defined interface name	NIC or NIC Group	Burb name
external_network	em0	external
internal_network	em1	internal

[Table 4](#) shows a transparent interface configured using the default interfaces. Note that bridge0 is made up of em0 and em1.

Table 4 Transparent interface

User defined transparent interface name	NIC or NIC Group
bridged_network	bridge0 (em0, em1)

- When traffic attempts to cross the transparent interface (from one burb to the other), a rule check is performed to enforce security policy.
- Since hosts inside the network are not aware that the Sidewinder is deployed, they communicate with each other as if they were directly connected by a switch.
 - If two hosts reside in the same burb (the same side of the transparent interface), they communicate directly over the network and no security policy is enforced.
 - If two hosts reside in different burbs (different sides of the transparent interface), they communicate through the firewall and security policy is enforced.

This section includes the following deployment scenarios:

- [Transparently enforcing security policy inside a single subnet](#)
- [Transparently protecting a single network](#)

Transparently enforcing security policy inside a single subnet

[Figure 6](#) depicts a single subnet (192.168.0.0/24) that contains both servers and users. For this example, assume that the network administrator has decided to introduce a firewall to protect the servers from the users. However, the network cannot be re-addressed and all of the servers and users must retain their current IP addresses. These requirements are met by Sidewinder transparent mode.

Figure 6 A single subnet containing servers and users





Figure 7 depicts a Sidewinder in transparent mode protecting the servers from the users. As traffic between the users and servers crosses the firewall's transparent interface, it also crosses from one burb to the other. This triggers a rule check which enforces security policy on the traffic. Note that while deploying the firewall in transparent mode does not require re-addressing the network, the firewall does require a management IP address. In this example, 192.168.0.10 was reserved for the firewall.


Figure 7 Transparent Sidewinder inside a single subnet



Transparently protecting a single network

Figure 8 depicts an internal network that is only protected from the Internet by a router. For this example, assume that the network administrator has decided to introduce a firewall to protect the internal network from the Internet. However, there is a requirement that the network cannot be re-addressed, and all of the servers and users must retain their current IP addresses.

Figure 8 No firewall

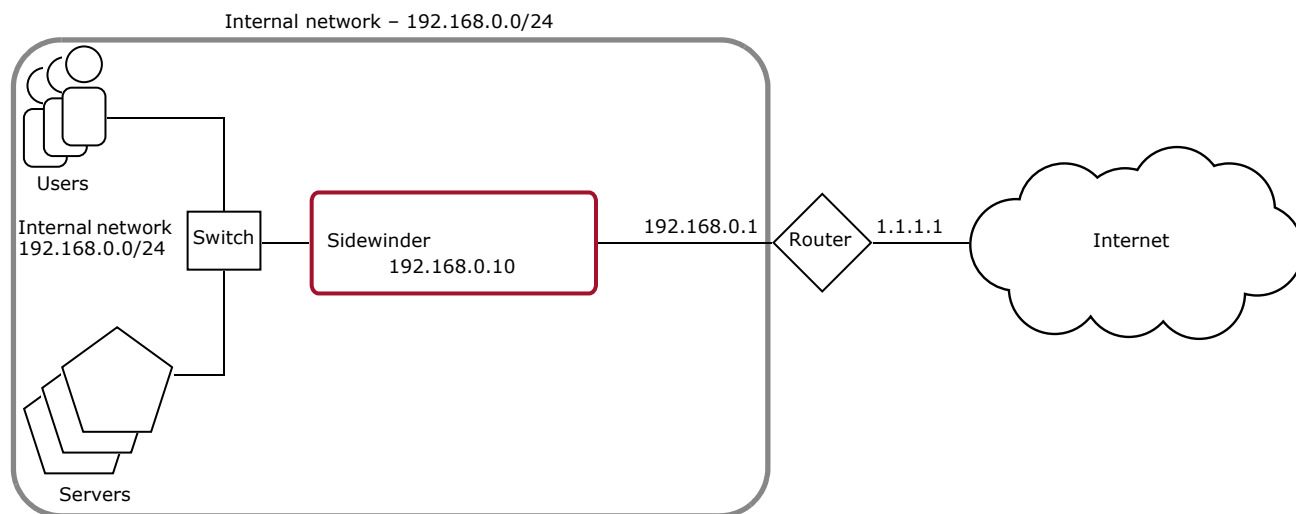


To deploy a Sidewinder to protect the internal network, two options are available:

- Deploy the firewall in router mode and re-address the networks around it.
- Deploy the firewall in transparent mode inside the internal network.

Using transparent mode has the advantage that none of the networks around the firewall need to change. In addition, the hosts on the internal network will have no knowledge that the firewall has been deployed between the switch and the router as shown in [Figure 9](#).

Figure 9 Transparent firewall protecting a single network



The default route for the hosts on the internal network is still the router (192.168.0.1), which is on the other side of the firewall. As traffic bound to the Internet from the internal hosts crosses the firewall's transparent interface, it also crosses from the internal burb to the external burb. This triggers a rule check which enforces security policy on the traffic. Note that while deploying the firewall in transparent mode does not require re-addressing the network, the firewall does require a management IP address. In this example, 192.168.0.10 was reserved for the firewall.

Hybrid mode

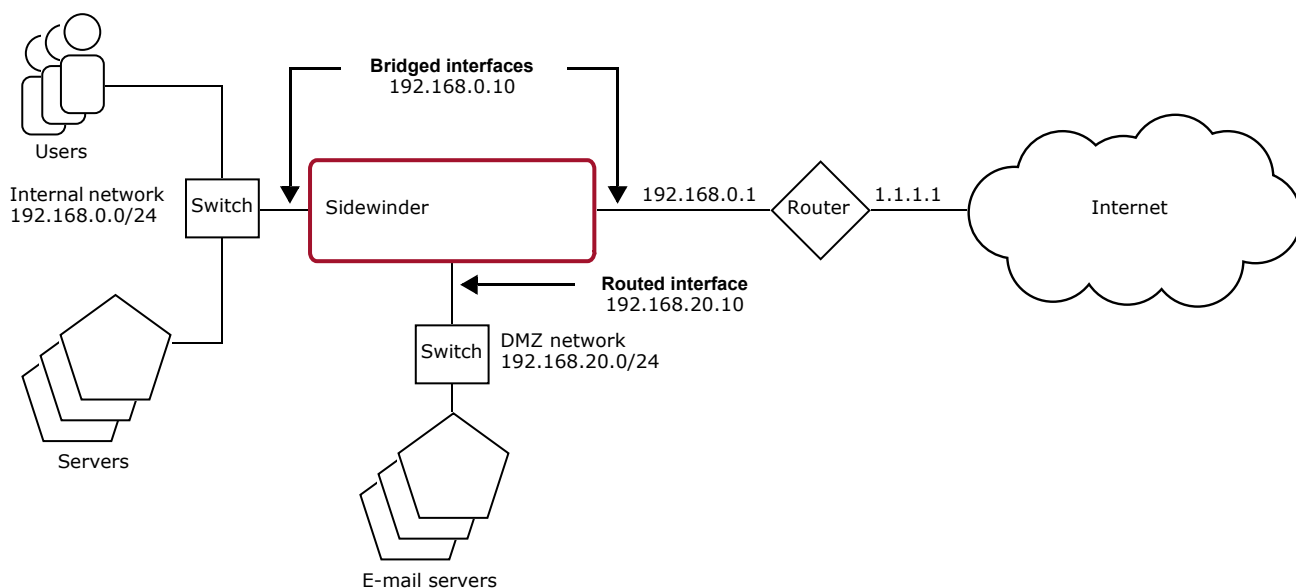
Figure 10 depicts a Sidewinder configured with a transparent interface and a routed interface. In this example, the firewall protects the internal and DMZ networks from each other and from the Internet. Note that the firewall has two IP addresses—a transparent IP address for management (192.168.0.10, assigned to both bridged interfaces) and a routed IP address on the DMZ interface (192.168.20.10).

To reach the Internet:

- Hosts in the internal network route traffic to the router's IP address (192.168.0.1) on the other side of the firewall.
- Hosts in the DMZ route traffic to the firewall's DMZ IP address (192.168.20.10).

As traffic crosses from interface to interface, it also crosses from one burb to another. This triggers a rule check which enforces security policy on the traffic.

Figure 10 Hybrid Sidewinder



2 Administrator Basics

Contents

Managing your firewall

Configuring Admin Console access

Restarting or shutting down the system

Managing administrator accounts

Changing administrator passwords

Administering Sidewinder using Secure Shell

Administering the Sidewinder using Telnet

Managing your firewall

This section explains basic Forcepoint Sidewinder management.

- *[Understanding Sidewinder management](#)*
- *[Admin Console basics](#)*
- *[Using the Admin Console](#)*
- *[Logging directly into the firewall](#)*

Understanding Sidewinder management

You can manage the Sidewinder in the following ways:

- **Admin Console** – The Administration Console (or *Admin Console*) is the graphical software that runs on a Windows computer within your network.
 - The Admin Console is installed using the “Management Tools” CD.
 - This CD also installs the Quick Start Wizard, which is used to initially configure your firewall.

See the *Setup Guide* for information on installing the Admin Console software and running the Quick Start Wizard.

Note: The Admin Console is occasionally referred to as “cobra” in some command line tools.

- **command line interface (CLI)** – If you are experienced with UNIX, you can also use the command line interface to configure and manage the firewall. Command line interface refers to any UNIX prompt. The command line interface supports many firewall-specific commands as well as standard UNIX commands you can enter at a UNIX prompt. For example, the `conf` command can perform a wide range of configuration tasks.

For help using the command line interface, refer to the following:


- *Command Line Interface Reference* at <https://support.forcepoint.com>.
- Manual (man) pages included on the firewall: log into the firewall at a command prompt, type **man** followed by the name of a command, and then press **Enter**.
- **Forcepoint Sidewinder Control Center** – An enterprise-class management appliance that allows you to centrally manage multiple Sidewinder appliances.

For more information, see the *Administration Guide* for Control Center.

Whether you use the Admin Console or the command line interface, you can manage the firewall from a number of locations. [Figure 11](#) highlights the administration interface options available to you.

Note: Normal administration is possible only when the Operational kernel is booted. The firewall in emergency maintenance mode is offline and does not pass traffic.

Figure 11 Administration options



- The firewall must allow secure sessions for the burb in which the Admin Console workstation resides.
- By default, access is enabled on the firewall's internal burb. For information on changing Admin Console access on an active firewall, see [Configuring Admin Console access](#).

Admin Console basics

To start the Admin Console on a Windows workstation, do one of the following:

- Select **Start > All Programs > Forcepoint > Sidewinder v7 Admin Console > Admin Console**.
- Double-click the Admin Console icon located on the desktop.

The main Admin Console window appears.


Use this window to connect to and manage one or more firewalls.

The main Admin Console window is divided into three areas: the toolbar, the left pane, and the right pane.

About the toolbar

The toolbar at the top of the Admin Console window contains menus and six buttons for various shortcut actions:

Figure 12 Admin Console menu and toolbar



Use the menus to perform the following actions:

Table 5 Admin Console menus

Menu	Options
File	<ul style="list-style-type: none"> • New Firewall (Ctrl+N) – Add a firewall that can be managed using the Admin Console. • Save (Ctrl+S) – Save changes. • Cancel (Ctrl+E) – Cancel changes. • Exit (Alt+X) – Exit the Admin Console.
Tools	<ul style="list-style-type: none"> • ARP – Use this feature to view the association between each MAC address on the firewall and its corresponding IP address. See About the ARP Table for more information. • Get route – Use this feature to find the first gateway in the route from the firewall to a stated destination. See About the Get Route window for more information. • DNS lookup – Use this feature to find the IP address for a host name. See About the DNS Lookup window for more information. • Ping host – Use this feature to test interface connectivity. See About the Ping Test window for more information. • TCP dump – Use this feature to capture the network traffic on selected firewall interfaces. To run tcpdump: <ul style="list-style-type: none"> a Select interfaces and set parameters that you want to capture traffic for. See About the TCP Dump window and About the TCP Dump Parameters window for more information. b Run and view the tcpdump. See About the Running TCP dump window for more information. • Traceroute – Use this feature to see all of the gateways that traffic passes through on a round trip between the firewall and a destination. See About the Traceroute window for more information.
Help	<ul style="list-style-type: none"> • [Window help] – Display information for the Admin Console window that is selected in the tree. The title for this menu option correlates to the window selected. Note: If you use a browser with a pop-up blocker turned on, you must allow blocked content to view the Sidewinder help. • About (Ctrl+H) – Display information about the current version of the Admin Console software.

Use the toolbar to perform the following actions:

Table 6 Admin Console toolbar

Icon	Action
New Firewall	Click this icon to add a firewall. For more information on adding a new firewall, see Adding a firewall to the Admin Console .
Save	Save changes to the firewall that you make in the Admin Console by clicking Save .
Rollback	Cancel (or rollback) any unsaved changes in the Admin Console by clicking Rollback .
Start ticket/ Stop ticket	Identify specific changes to the firewall by clicking Start ticket . Click Stop ticket to close the change ticket.
Refresh	Refresh or update the screen by clicking Refresh .
Help	Access online help for the current Admin Console window that is displayed by clicking Help .

About the left pane of the Admin Console window

The left pane of the window contains the Admin Console tree. You can add or delete a firewall from the tree without being connected. Once you are connected to a specific firewall, you can click any of the items in the Admin Console tree to manage that area of your firewall.

You can also right-click a firewall in the Admin Console tree to perform the following actions:

- Delete a firewall from the Admin Console.
- Connect or disconnect a firewall from the Admin Console.
- Expand or collapse all or sections of the branch items beneath a firewall icon.

About the right pane of the Admin Console window

When not connected to a firewall that is currently selected in the tree, the right pane of the Admin Console window displays configuration information for that firewall.

- **Name** – The name of the firewall.
- **IP Address** – The IP address of the firewall.
- **Port** – The port number used to connect to the firewall.
- **Version** – This is a read-only field that displays the current version of the firewall.
- **Sidewinder State** – This is a read-only field that displays the current firewall state (standalone or part of an HA cluster).
- **Connect** – Click this to connect to the selected firewall.

Admin Console conventions

When using the Admin Console connected to a firewall, the following conventions and tips will help you avoid common mistakes:

- To sort or filter a table based on the contents of a single column, right-click a column heading and select the filter criteria for which you want to filter. To customize a filter, select the **(Custom)** option. To view all items in a table, select the **(No Filter)** option.

You can also reverse the order of the table within a column by clicking the appropriate column heading. To return the table to its original order, click the column heading a second time.

Note: You cannot filter the table on the Rules window. You must open the Active Rules window.

- You can select an item to modify from a list by double-clicking it, selecting it and then clicking **Modify**, or right-clicking it and selecting **Modify**. (Read-only administrators can click **View** to view an item.)
- When a box preceding an option is filled in or contains a check mark, it is enabled or selected. When the box is empty (a check mark does not appear), the option is disabled.
- On some windows, you need to use the scroll bar to view all of the information or options.
- To delete an item from a list or table in an Admin Console window, click the item to select it, and then click **Delete**.
- When you leave a window that you have modified, you will automatically be prompted to save your changes before you exit the window. You can also save your modifications at any time by clicking the **Save** icon in the toolbar (or an **OK** button for some pop-up windows).
- When you exit a window and do not want to save your changes, click **No** when prompted to save your changes. You can also cancel your changes at any time by clicking the **Rollback** icon (or the **Cancel** button in some windows) to restore the current window's settings to the last saved version.
- For assistance on any of the Admin Console windows, click the **Help** icon located in the top portion of the window. The online help provides information about each of the Admin Console windows. To view the entire list of available help topics, click the **Contents** tab from within the help system.

Note: If you use a browser with a pop-up blocker turned on, you must allow blocked content to view the Sidewinder help.

- To exit the Admin Console, do one of the following:
 - From the **File** menu, select **Exit**.
 - Click the **X** icon in the upper right corner of the Admin Console window.
 - Press **Alt+X**.

Note: If you have any active connections when you exit the Admin Console, those connections, as well as any unsaved changes, will be lost. You will not be prompted to save before exiting.

Using the Admin Console

Be aware of these conditions when using the Admin Console to manage a firewall:

- This version of the Admin Console is not compatible with 6.x versions of the Admin Console or the Sidewinder G2 firewall.
- The firewall policy must allow Admin Console access for the burb in which the Admin Console workstation resides. By default, access is enabled on the firewall's internal burb. For information on changing Admin Console access on an active firewall, see [Configuring Admin Console access](#).

Use these procedures to add firewalls to the Admin Console tree and to connect and disconnect from a firewall:

- [Adding a firewall to the Admin Console](#)
- [Connecting to a firewall](#)
- [Disconnecting from a firewall](#)

Adding a firewall to the Admin Console


Before you can manage a firewall using the Admin Console, you must first add it to the Admin Console tree.

To add a firewall to the tree, use one of these methods to open the Add Firewall window:

- From the **File** menu, select **New Firewall**.
- In the Admin Console toolbar, click the **New Firewall** icon.
- In the Admin Console left pane, right-click the **Firewalls** icon and select **New** from the pop-up menu.

The Add Firewall window appears.

Figure 13 Add Firewall window




- 1 In the **Firewall Name** field, type a descriptive name for the firewall you are adding. For example, you might specify the host name you used during the installation process. Only alphanumeric characters, dashes (-), and underscores (_) can be used; spaces are not allowed.
- 2 In the **Firewall IP Address** field, type the IP address you want to use to access the firewall. The address must be a valid IP address for an interface on the firewall. Also, the interface must be contained within a burb selected for Admin Console access.
- 3 Click **Add** to save the information and exit this window. The firewall is displayed in the Admin Console tree in the left pane.

Connecting to a firewall

Note: You cannot connect to a 6.x version of the firewall using a 7.x version of the Admin Console.

To connect to a specific firewall, select the appropriate icon from the Admin Console tree and then click **Connect**. The Login window appears.

Figure 14 Admin Console Login window



The first time you attempt to connect to a firewall using the Admin Console, a pop-up window appears presenting you with the firewall certificate that will be used for all subsequent administrative connections. To accept the certificate, click **Yes**. If you want to verify the certificate before accepting it, you must obtain the certificate fingerprint before you log into the Admin Console. To obtain the certificate fingerprint, log into the firewall via command line and enter the `show role` command to change to the Admin domain. (If you have not configured remote access, you will need to attach a monitor and keyboard directly to your firewall.) Enter the following command:

```
cf cert view fw name=cert_name
```


The contents of the certificate are displayed. The certificate fingerprint is located at the bottom of the certificate directly beneath the END CERTIFICATE identifier. This fingerprint can be used to verify the fingerprint that is displayed when you initially connect to the firewall via the Admin Console.

To log into a firewall:

- 1 In the **Username** field, enter your user name.
- 2 In the **Authenticator** drop-down list, select the appropriate authentication method for the firewall to which you are connecting.
 - **Password** is the default authentication method. Other authentication methods must be configured on the firewall before they are available in this drop-down list.
 - If you want to have a backup authentication method, duplicate the Admin Console rule and select a different authenticator.
 - All methods other than the password method require access to a separate authentication server.
- 3 Click **OK**. The Password Authentication window appears.
- 4 Enter your password, and then click **OK**.

When you connect for the first time, the Feature Notification window displays the status of each licensed feature.

Figure 15 Feature Notification window



Tip: If you do not want this window to appear each time you connect, select the **Don't show this again** check box.

- 5 When you are finished viewing the window, click **Close**.

The main Admin Console window appears.

Disconnecting from a firewall

To end an Admin Console session for a firewall:

- In the left tree, select the firewall icon, and click **Disconnect** in the main Admin Console window.
- In the left tree, right-click the firewall icon and select **Disconnect** from the pop-up menu.

This disconnects the Admin Console from the firewall. It does not shut down the firewall.

Logging directly into the firewall

You can manage the firewall by command line interface by logging directly into the firewall. One way to do this is through SSH. A default Secure Shell Server rule allows SSH server access to the firewall. This rule must be enabled.

To log directly into the firewall:

- 1 At the login prompt, type your user name and press **Enter**. The Password prompt appears.
- 2 Type your password and press **Enter**. The User domain prompt appears:

```
firewall_name:User {1} %
```

When you initially log into the firewall using a command prompt, you are logged into the User domain by default. The User domain allows very little access, including no access to sensitive files.

- 3 To change to the Admn domain, which allows access to all firewall domains (based on your administrative role), enter the following command:

```
srole
```

- 4 To return to the previous domain role and shell, enter the following command:

```
exit
```


You are returned to the User domain.

Note: If you have read-only privileges, type `srole adminro`

Configuring Admin Console access

Sidewinder is managed from the Sidewinder Admin Console, which must be installed on a Windows workstation.

- The Quick Start Wizard enables access on the internal burb. If you want to establish an Admin Console connection to a different burb, modify the Admin Console rule. See [Modifying the Admin Console rule](#) for details.
- When the Admin Console connects to a firewall for the first time, you are prompted to accept a certificate before the connection will continue. A default SSL certificate is initially assigned to the Admin Console.

We recommend assigning a new certificate to the Admin Console before using the firewall in an operational environment. See [Configuring the Admin Console server](#) for details.

- You can configure a banner message that appears when the Admin Console connects to the firewall. This message is generally to alert users that they are accessing proprietary information. The banner window has an **Accept** button that must be clicked to proceed.

See [Configuring the Admin Console server](#) for details.

- The default port for the Admin Console is 9003.
 - See [Network Services Sentry \(NSS\)](#) for details on selecting valid ports.
 - To change the port or timeout properties for the Admin Console, see [Configuring the Admin Console server](#).

Modifying the Admin Console rule

Perform this procedure to enable Admin Console access to different burbs.

To modify the Admin Console rule:

- 1 Select **Policy > Rules**. The Rules window appears.
- 2 In the Rules list, expand the **Administration** rule group, select **Admin Console**, and click **Modify**. The Modify Server Rule window appears.

Figure 16 Modify Server Rule window: Admin Console

Rules: Modify Server Rule

Name: ☒ Enable

Description:

General

Action: ☒ Allow ☐ Deny ☐ Drop

Service: ...

Audit:

Effective Times

Time period: ...

☐ Start on: ...

☐ Expire on: ...

Source

Burb:

Endpoint:
external
internal

NAT:

☐ Preserve source port ...

Destination

Burb:

Endpoint:

Redirect:

Redirect port: ...

TrustedSource

☐ Enable TrustedSource

Malicious Suspicious Unverified Neutral Trusted

Unverified, neutral and trusted traffic will match this rule.
(Range 29 to -255)

Inspection

Application Defense: ...

Full ☒ -The security-context "filtering" aspects of the application defense are not enforced. Application layer data is examined to the minimum degree necessary to perform "proxy" activities as defined by said protocol.

None ☐

IPS Signature group: ...

Response mapping: ...

Authentication

Authenticator:

Allow users in the following groups: ...

Last modified by a on 10/06/08 11:32:10 AM EDT

OK Cancel Help


- 3 From the Source and Destination **Burb** drop-down lists, select the appropriate burb for your Admin Console connection.
- 4 Click **OK**.
- 5 Save your changes.

Configuring the Admin Console server

The Admin Console is the graphical user interface used to manage your Sidewinder. The Admin Console connects to the firewall using an SSL connection to a dedicated port (port 9003). When the Admin Console connects to a firewall for the first time, you are prompted to accept a certificate before the connection will continue. The Admin Console service also enforces the TCP idle timeout.

To configure these properties, select **Policy > Rule Elements > Services** and double-click the Admin Console service. The Admin Console service window appears.

Figure 17 Admin Console service window



You can perform the following actions:

- Change the Admin Console's port and idle timeout values. Defaults are:
 - TCP port: 9003 (See [Table 40 on page 341](#) for details on selecting valid ports.)
 - TCP idle timeout: 0 seconds (This means that there are no timeouts.)

You can click **Restore Defaults** at any time to restore the timeout value.


- Click **Properties** to change the SSL certificate and create an optional logon greeting.

Once the properties match your site's security policy, click **OK** to return to the main Services window and then save your changes.

About the Admin Console Properties window

Use this window to select the certificate that the firewall uses for Admin Console connections. You can also create a login banner message.

Figure 18 Admin Console Properties window



- 1 In the **SSL Certificate** drop-down list, select a certificate. The certificate will be one of the following:
 - The default certificate
 - A self-signed, RSA/DSA certificate that is defined on the Firewall Certificates tab of the Certificate Management window
- 2 To use a certificate that is not in the list or to view an existing certificate's properties, click **Certificates**. The Firewall Certificates window appears.

For detailed information on certificates, refer to [Configuring and displaying remote certificates on page 638](#).
- 3 [Optional] In the **Message** field, type the text you want to appear when a user connects to the firewall with the Admin Console. This message is generally to alert users that they are accessing proprietary information. The banner window has an **Accept** button that must be clicked to proceed.
 - If you select **Required**, the banner message appears each login attempt.
 - If **Required** is cleared, a **Don't show this again** option appears on the banner window.
 - Click **Preview** to see a preview of the banner message.
- 4 Click **OK** to return to the Admin Console service window.

Be sure to save your changes once you return to the main Services window.

Restarting or shutting down the system

There are four Sidewinder shutdown options:

- **Reboot to Operational Kernel**
 - The firewall boots to the Operational kernel by default. You can boot to the Operational kernel through the Admin Console or by pressing the power button.
 - You can log into the firewall via the Admin Console and perform administrative tasks.
- **Shutdown to Emergency Maintenance Mode**
 - Emergency Maintenance Mode (EMM) allows you to do repair work with other services turned off. You should use EMM only if directed by Forcepoint support.
 - The # prompt appears on the firewall, indicating that you are in a login shell and can start issuing firewall or UNIX commands.
 - The firewall in EMM is offline and does not pass traffic.
 - You must connect a console to the firewall in order to work with it. You cannot access the firewall via the Admin Console, SSH, or telnet in emergency maintenance mode.
- **Halt System**
 - The operating system shuts down, but the system remains powered on.
 - Halt System is useful if you need to connect directly to the firewall to access the BIOS.
- **Power Down System**
 - You completely shut down the firewall without restarting.
 - Power down the system before you move your firewall to a new location or make hardware changes.

You can reboot or shut down a firewall from the Admin Console or the command line.

- When the firewall is rebooted or shutdown, a record of who issued the action is logged in the `/var/log/messages` file. This applies to a reboot or shutdown issued from the Admin Console or using the **shutdown** command.
- If the boot process fails, contact Forcepoint support.

Rebooting or shutting down using the Admin Console

To reboot the firewall or to shut down the firewall completely, select **Maintenance > System Shutdown**. The System Shutdown window appears.

Figure 19 System Shutdown window

To reboot or shut down the firewall:

1 In the **Shutdown Options** area, select the action you want to perform:

- **Reboot to Operational Kernel** – Restarts the system in the Operational kernel.
- **Shutdown to Emergency Maintenance Mode** – Restarts the system in emergency maintenance mode and displays the # prompt, indicating that you are in a login shell and can start issuing firewall or UNIX commands.
 - While the firewall is in emergency maintenance mode, it is offline and does not pass traffic.
 - You must connect a console to the firewall before you can administer the system in emergency maintenance mode.
- **Halt System** – Shuts down the operating system, but the system remains powered on. Run this command if you need to connect directly to the firewall to access the BIOS.
- **Power Down System** – Completely shuts down the firewall software without restarting. Run this command before you move your firewall to a new location or make hardware changes.

2 [Optional] If you want a shutdown message to appear informing users of a pending shutdown, type the message text in the **Shutdown Message** field.

3 In the **Shutdown Time** field, select the shutdown time from the following options.

- **Shutdown Immediately** – The system will shutdown immediately when you click **Perform Shutdown**.
- **Delay Shutdown for** – The shutdown will be delayed for the amount of time specified in the **Hours** and **Minutes** fields. You can enter values in these fields that will delay the shutdown for up to 24 hours and 59 minutes.

4 Click **Perform Shutdown** to implement the shutdown.

Any connections to the Admin Console will be lost when the firewall shuts down. New connections to the firewall will not be allowed once the shutdown process has been executed.

Rebooting or shutting down using a command line interface

The **shutdown** command reboots or shuts down the system from a command line interface. Use this command to indicate how and when you want the firewall to shut down.

The table below shows some common shutdown commands from the command line.

- More information about shutdown options is available on the **shutdown** man page.
- For information on shutting down a Sidewinder that belongs to an HA cluster, see [Scheduling a soft shutdown for a load-sharing HA cluster Sidewinder on page 675](#).

Table 7 Shutdown commands on the command line

Command	Type of shutdown
shutdown -r [time]	Restarts the system in the Operational kernel. For example, shutdown -r +120 would reboot the firewall into its Operational kernel in two hours (120 minutes).
shutdown [time]	Restarts the system to emergency maintenance mode. For example, shutdown now would immediately reboot the firewall into emergency maintenance mode.
shutdown -h [time]	Shuts down the firewall without restarting. For example, shutdown -h 0601312359 would halt the firewall at one minute to midnight on January 31, 2006.
shutdown -p [time]	Completely powers off the system without restarting. For example, shutdown -p now would immediately shut down the firewall.
shutdown [-rh] -s soft_time time	<p>A load-sharing HA cluster always performs a soft shutdown. A soft shutdown provides a buffer period before the actual shutdown occurs.</p> <p>By default, the soft shutdown process will begin 30 minutes before a scheduled shutdown. If the shutdown is scheduled to occur in less than 30 minutes, the soft shutdown process will begin immediately and will remain in effect until the actual shutdown time occurs.</p> <p>You can schedule a specific shutdown time for a cluster, or a number of minutes until the shutdown, by using -s. For example:</p> <pre>shutdown -r -s +45 +60</pre> <p>(with soft shutdown in 15 minutes, with reboot in one hour)</p> <pre>shutdown -r -s 1500 1800</pre> <p>(reboot at 6:00, starting soft shutdown at 3:00)</p> <p>Note: You must include a soft shutdown time if you use the -s command.</p>

Managing administrator accounts

Each Sidewinder administrator must have an account created on the system. The initial administrator account, including user name and password for login authentication to the firewall, is created during startup configuration using the Quick Start Wizard. This section describes how to set up and maintain firewall accounts for other administrators.

Note: Only administrators have accounts directly on the firewall. People who use firewall networking services have “user” (or network login) accounts, not firewall administrator accounts. See [Authenticating groups from an internal group source on page 107](#) for information on creating non-administrative user accounts.

When you add an administrator account, you also assign the new administrator a role. The following table describes the available administrator roles. The following processes explain how to view, add, edit, or delete administrator account information or change role assignments.

Table 8 Administrator roles

Role	Authorized to:
admin	<ul style="list-style-type: none">• Access all windows, menus, and commands within the Admin Console.• Add and remove users and assign roles.• Do incremental back-ups and restore the system.• Use all other system functions and commands.
adminro	This role will allow an administrator to view all system information, as well as create and run audit reports. An administrator with readonly privileges cannot commit changes to any area of the firewall. This role is generally used as an auditor role.
no admin privileges	Maintains an existing or new administrator account with limited access to the User domain. This role is generally used to temporarily disable an administrator account.

To view and manage administrator accounts, select **Maintenance > Administrator Accounts**. The Administrator Accounts window appears.

Figure 20 Administrator Accounts window

The screenshot shows a web-based interface for managing administrator accounts. It features a table with the following data:

Username	Full Name	Role	Directory
administrator		admin	/home/administrator

Below the table, there are three buttons: **New**, **Modify**, and **Delete**. At the bottom left, there is a checkbox labeled "Delete home directory upon deletion of user". At the bottom right, there is a text input field labeled "Administrator email address:" with the letter "a" entered.

This window displays the administrator accounts currently established on the firewall.

The table identifies the administrator user name, full name, role, and home directory path for each administrator.


You can perform the following actions:

- **Create a new administrator account** – Click **New** and enter the account information in the New Administrator window.
- **Modify an existing administrator account** – Select an administrator in the table and click **Modify**, then make the desired changes in the Modify Administrator window.
- **Delete an existing administrator account** – Select an administrator in the table and click **Delete**.
 - When you delete an administrator account, the user database entry for that administrator is also removed.
 - To automatically delete an account's home directory when the account is deleted, select **Delete Home Directory Upon Deletion Of User**.

About the New/Modify Administrator window

Use this window to create or modify a firewall administrator account.

Figure 21 Administrator Information tab



- 1 In the **Username** field, type the user name for the administrator. The name can be up to 16 alpha-numeric characters.

If you are modifying an existing account, you cannot change the user name.

Note: Do **not** use uppercase characters in the **Username** field, because sendmail will automatically convert the user name to lowercase before mail is delivered. Therefore, any mail addressed to a user name that contains uppercase characters will not be forwarded.

- 2 In the **Password** field, type a password for this administrator. This is the password the administrator uses when logging into the firewall. Use the following guidelines to create a more effective password:

- Use passwords that are at least 7 or 8 characters in length.
- Use a mix of upper- and lowercase letters, and non-alphabetic characters such as symbols and numbers.
- Do not use any easily guessed words or words found in a dictionary, including foreign languages.

- 3 In the **Confirm Password** field, retype the password.

- 4 [Optional] In the **Full Name** field, type the full name of the administrator.

- 5 [Optional] In the **Office** field, type the office address of the administrator.

- 6 [Optional] In the **Office Phone** field, type the office phone number of the administrator.

- 7 [Optional] In the **Home Phone** field, type the home phone number of the administrator.

- 8 In the **Directory** field, specify the home directory for this administrator. The default value for this field is `/home/username`. This field can be modified only if you are creating a new administrator account.

9 In the **Login Shell** drop-down list, specify the UNIX shell that will be used when this administrator logs in.

10 In the **Roles** drop-down list, select the authorized role for this administrator:

- **admin** – Select this option if you want the user to have administrator privileges for all areas on the firewall.
- **adminro** – Select this option to allow read privileges only. This role will allow an administrator to *view* all system information, as well as create and run audit reports. An administrator with read-only privileges cannot commit changes to any area of the firewall.
- **no admin privileges** – Select this option to limit an administrator's access to the firewall. An administrator with no admin privileges cannot log into the firewall.

11 [Conditional] If you use CAC authentication, from the **CAC certificate** drop-down list, select the remote certificate imported for the administrator.

12 Click **Add** or **OK** and save your changes.

You are done creating or modifying this administrator account.

Changing administrator passwords

To change an administrator account password (also known as a UNIX account password), do the following:

Note: If you forget your password, you can use the emergency maintenance mode to change your password. See [Changing a forgotten password on page 682](#).

- 1** Select **Maintenance > Administrator Accounts**. The Administrator Accounts window appears.
- 2** Select the administrator account whose password you want to change, then click **Modify**. The Firewall Accounts: Modify Administrator window appears.
- 3** In the **Password** field, enter the new administrator account password, then confirm the new password.
- 4** Click **OK** and save your changes.

Administering Sidewinder using Secure Shell

Secure Shell (SSH) provides secure encrypted communication between two hosts over an insecure network, allowing you to securely manage your firewall from a remote location. This section describes how to configure and use the firewall as an SSH server and/or an SSH client.

- The procedures covered in the following sections are based on the use of OpenSSH version, which provides support for SSH version 1.5 and 2.0 sessions.
- sftp and sftp-server are included in OpenSSH and installed on the firewall.

Configuring the SSH server

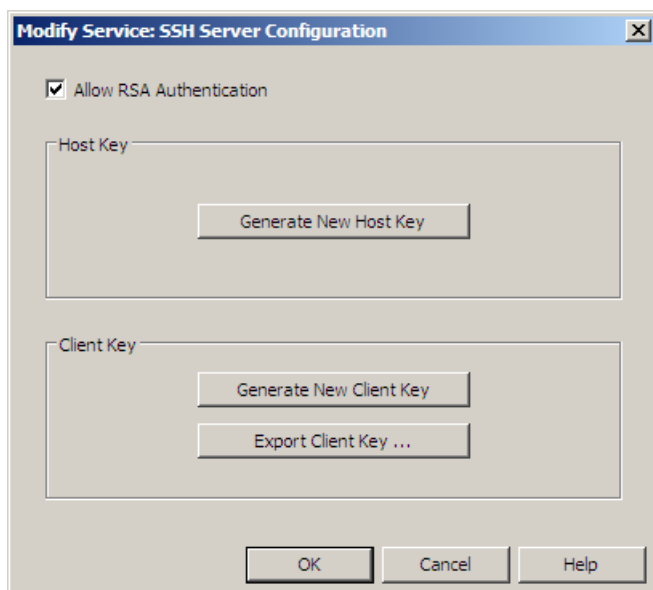
Your firewall can act as an SSH server, an SSH client, or both.

- If it will act as a server, use the `sshd` service window to generate a host key.
- If it will act as an SSH client that connects to other firewalls, use the `sshd` service window to generate a client key. See [Configuring and using the Sidewinder as an SSH client](#) for details.

To configure the SSH server:

- 1 Select **Policy > Rule Elements > Services**.
- 2 In the list of services, select **sshd**, and then click **Modify**.
- 3 If necessary, change port and timeout settings.
- 4 Click **Properties**. The SSH Server Configuration window appears.

Figure 22 SSH Server Configuration window



Use this window to generate host and client keys, and to specify whether RSA/DSA authentication is allowed. Follow the steps below.

Tip: If you plan to export client keys to other firewalls, use the Admin Console to connect to the other firewalls before starting this procedure.

- 1 If you want to allow SSH connections to be authenticated using RSA/DSA authentication, select the **Allow RSA Authentication** check box.

RSA/DSA authentication is a common encryption and authentication system that uses an exchange of public and private keys between the server and the client. It is based on the RSA/DSA algorithm. If this check box is not enabled, all SSH connections must be authenticated using the authentication method selected in the SSH rule.

- 2 To generate an SSH host authentication key that will be used when the firewall is acting as the server in an SSH connection, click **Generate New Host Key**. The firewall automatically generates the following three authentication keys: RSA1, RSA, and DSA.
- 3 To generate the SSH version 1.5 client authentication key that will be used when the firewall is acting as a client in an SSH connection, click **Generate New Client Key**.
- 4 [Conditional] To export the client key to another Sidewinder, click **Export Client Key**.

A new window appears, listing all firewalls with which your Admin Console has an active session. Select which firewalls will receive the key and click **OK**.

You can only export the client key if:

- you generated a client key as described in [Step 3](#), and
- you currently have an active Admin Console connection with one or more additional firewalls (the firewall[s] that will act as the SSH server).

- 5 Click **OK** to return to the SSH service window.

Be sure to save your changes once you return to the main Services window.

Configuring the Sidewinder as an SSH server

On the firewall, SSH is typically used by administrators to log into the firewall securely from a remote machine. In this case the firewall acts as the SSH *server*.

When configuring the SSH server you have the option to use RSA/DSA authentication. If you use RSA/DSA authentication, the authentication is accomplished via an exchange of public and private keys between the server and the client. The downside of RSA/DSA authentication is that it requires a bit more of an administrative effort. If you elect NOT to use RSA/DSA authentication, the SSH clients must enter their firewall user name and authentication information when initiating the SSH connection.

The following sub-sections provide specific information on configuring the firewall as an SSH server using RSA or DSA authentication, as well as general information on configuring the SSH server.

Configuring SSH when not using RSA/DSA authentication

If you are *not* using RSA/DSA authentication, follow the steps below to configure SSH.

- 1 Disable RSA authentication:
 - a Select **Policy > Rule Elements > Services**.
 - b In the list of services, select **sshd**, and then click **Modify**.
 - c Click **Properties**. The SSH Server Configuration window appears.
 - d Ensure that **Allow RSA Authentication** is disabled.
 - e Click **OK**.
 - f Click **OK** and save your changes.

2 Enable and modify the Secure Shell Server rule:

- a** Select **Policy > Rules**.
- b** In the Rules list, expand **Administration Server**, select **Secure Shell Server**, and then click **Modify**. The Modify Server Rule window appears.
- c** Select **Enable**.
- d** Select the desired source and destination burbs.
- e** Select an authentication method.
- f** Click **OK** and save your changes.

3 [Conditional] If a Host Key Pair does not exist, you will be prompted to confirm that the Admin Console will create an SSH host key. Click **Yes**.

Note: If the client has previously established an SSH connection to the firewall, the information associated with the previous connection must be deleted from the client.

The firewall is now ready to accept SSH connection requests. Remember that a client must have an administrator account on the firewall in order to log in.

Configuring SSH when using RSA/DSA authentication

If you are using RSA /DSA authentication to authenticate SSH, follow the steps below.

1 Configure the SSH server:

- a** Select **Policy > Rule Elements > Services**.
- b** In the list of services, select **sshd**, and then click **Modify**.
- c** Click **Properties**. The SSH Server Configuration window appears.
- d** Verify or select **Allow RSA Authentication**.
- e** If you do not currently have an SSH host key pair, click **Generate New Host Key**. Click **OK** to acknowledge that the new key pair has been created.

You must have at least one SSH host key pair for the SSH server to operate. If you have an existing key pair, you do not need to create a new one. The host key pairs are stored in the `/etc/ssh` directory and have the following file names:

<code>ssh_host_key</code>	SSH version 1.5 rsa private key
<code>ssh_host_key.pub</code>	SSH version 1.5 rsa public key
<code>ssh_host_rsa_key</code>	SSH version 2.0 rsa private key
<code>ssh_host_rsa_key.pub</code>	SSH version 2.0 rsa public key
<code>ssh_host_dsa_key</code>	SSH version 2.0 dsa private key
<code>ssh_host_dsa_key.pub</code>	SSH version 2.0 dsa public key

- f** Click **OK**.
- g** Click **OK** and save your changes.

2 Create public key directories for each user:

- a** From a command line prompt, create a subdirectory named `/.ssh` in each administrator's home directory.

Example: If an administrator named `lloyd` has a home directory named `/home/lloyd`, create the `/.ssh` subdirectory by typing the following commands:

```
srole  
cd /home/lloyd  
mkdir .ssh
```

- b** Use a text editor to create a file named `authorized_keys` in each administrator's `/.ssh` directory.

Do this using the File Editor provided in the Admin Console, or your favorite UNIX editor.

- c** Paste each user's public key into the respective `authorized_keys` file.

The method you use to get the public keys onto the firewall is up to you. You might use FTP, or you might copy/paste from one window to another.

3 Enable and modify the Secure Shell Server rule:

- a** Select **Policy > Rules**.
- b** In the Rules list, expand **Administration Server**, select **Secure Shell Server**, and then click **Modify**. The Modify Server Rule window appears.
- c** Select **Enable**.
- d** Select the desired source and destination burbs.
- e** Select an authentication method.
- f** Click **OK** and save your changes.

The firewall is now ready to accept connections from SSH clients. Remember that an administrator must have an account on the firewall in order to log in.

Configuring and using the Sidewinder as an SSH client

It is also possible for the firewall to act as an SSH *client*. For example, you might want to establish an SSH connection between two firewalls. In this case one firewall operates as the server (via the SSH server), and the other operates as an SSH client. You have the option to use RSA/DSA authentication with the SSH client.

Note: On non-Sidewinder systems, an SSH client that is run from root will bind to a reserved port. As a security feature, the firewall SSH client is not allowed to bind to a reserved port. This is prevented by Type Enforcement.

If not using RSA/DSA authentication

There is nothing to configure on the firewall if you are not using RSA/DSA authentication. To use the firewall as an SSH client, follow the steps below:

- 1 From a console attached to the firewall, log in and enter `srole` to switch to the Admin domain.
- 2 Establish the connection with the SSH server by typing one of the following commands:

```
ssh -l login_name address
```

or

```
ssh login_name@address
```

where:

login_name = the name used when logging onto the SSH server.

address = the name or address of the host with which you are establishing an SSH connection.

You have the option to use an authentication method other than the default method when connecting to another Sidewinder. Type a colon and the name of the authentication method after the *login_name* field. For example, to use SafeWord you would type:

```
ssh -l login_name:safeword address
```

If using RSA/DSA authentication

To use the firewall as an SSH client while using RSA/DSA authentication, you must perform several configuration steps before initiating the SSH connection.

Configuring the Sidewinder as an SSH client

- 1 Select **Policy > Rule Elements > Services**.
- 2 In the list of services, select **sshd**, and then click **Modify**.
- 3 Click **Properties**. The SSH Server Configuration window appears.
- 4 Click **Generate New Client Key** to generate a public and private key pair that the firewall can use when acting as an SSH client. The client public and private keys are created in the `/home/username/.ssh` directory, where *username* is the user name you used when connecting to the Admin Console. The file names vary, depending on the SSH version:
 - SSH version 1.5 – The client public key file name is *identity.pub* and the private key file name is *identity*.
 - SSH version 2.0 – The client public key file names are *id_rsa.pub* and *id_dsa.pub*. The corresponding private key file names are *id_rsa* and *id_dsa*.
- 5 [Conditional] If the SSH server that you will be connecting to is another Sidewinder, connect to that firewall using the Admin Console at this time.

If needed, click the **New Firewall** button in the top portion of the Admin Console and add the other firewall(s) to the list of firewalls you can administer.
- 6 If the SSH server that you will be connecting to is another Sidewinder, click **Export Client Key** to export the public client key to the other Sidewinders. Otherwise, use the best available method (FTP, cut and paste, etc.) to export the public client key to the SSH server.
- 7 Select the firewall to export to, and click **OK**.

Using the Sidewinder as an SSH client

- 1 At a Sidewinder command prompt, enter the following command to switch to the Admn role:

```
srole
```

- 2 Establish the connection with the SSH server by typing the following command:

```
ssh -l login_name hostname
```

where:

login_name = the user name used when logging onto the SSH server

hostname = the host name or address of the host with which you are establishing an SSH connection

See the ssh man page for more details.

On the firewall, the SSH client must be run from the Admn domain. Many SSH servers, however, do not allow root users to connect to the SSH server. To get around this, be sure to use the `-l` option when logging in. This allows you to login as a different user.

Tips on using SSH with the Sidewinder

Please note the following information about SSH on the firewall.

- There are two configuration files associated with SSH:
 - For the SSH server: */etc/ssh/sshd_config*
 - For the SSH client: */etc/ssh/ssh_config*
- See the *ssh*, *sshd*, *ssh_config*, *sshd_config*, and *ssh-keygen* man pages for additional details.
- The firewall's SSH server and client are based on the *OpenSSH* implementation. See <http://www.openssh.com> for more information.

Administering the Sidewinder using Telnet

To troubleshoot Sidewinder problems using a command line interface rather than the Admin Console, you can configure Telnet services that allow you to connect from a system within your network. You can also allow trusted users to use a Telnet client to log into Internet systems remotely.

Setting up an internal (trusted) Telnet server

Telnet provides a way to log into a system in your network from another system. All you need to know is the name of the system in which you want to log in. Once you have established a connection, you are logged in just as you would be if you were physically located at that system.

A Telnet server is defined for each burb on your firewall: one for the external (Internet) burb and one for each of the internal (or trusted) burbs. This gives you the capability to Telnet to the firewall from any system on an internal burb so you can perform administrative tasks remotely.

Note: For security reasons, the Telnet servers are not initially enabled.

Create a rule to access the trusted Telnet server. Include these selections:

- Select **telnetd (Telnet Server)** as the service.
- Select the source and destination burbs you want the Telnet server to access.
- Select an authentication method. All users accessing a Telnet server must be authenticated.

To perform firewall administration tasks, you must have an account on the firewall as described on [Managing administrator accounts](#). Aside from your account and authentication information, all you need to log into the firewall is the name or address. To log into the firewall using Telnet, see [Connecting to the Sidewinder using Telnet](#).

Setting up an external Telnet server

The Sidewinder allows you to enable an external Telnet server. An external server resides on the external network side of the firewall, and is available to Internet users once you set up the appropriate “allow” rules. (The other Telnet servers reside on the internal side of the firewall and are available only to trusted users.)

Security Alert: Setting up a Telnet server on the external side of your firewall can raise security issues. Contact Forcepoint support before attempting this.

Connecting to the Sidewinder using Telnet

Note: You must enable the Telnet server in the appropriate burb(s) before you will be allowed to Telnet. See [Setting up an internal \(trusted\) Telnet server](#).

- 1 Telnet to the firewall and log in by typing the following command, using the firewall’s host name.

```
telnet hostname
```

When prompted, enter your firewall authentication information. Depending on the authentication method configured for you on the firewall, you must provide a valid password or a special passcode or personal identification number (PIN) before you are logged on to the firewall.

- 2 Enter the following command:

```
srole
```

Enter commands from the UNIX prompt as required. For information on using individual commands, refer to the following

- *Command Line Interface Reference* at <https://support.forcepoint.com>.
- Manual (man) pages included on the firewall: log into the firewall at a command prompt, type **man** followed by the name of a command, and then press **Enter**.

SECTION 2

Policy

Chapter 3, Policy Configuration Overview

Chapter 4, Network Objects and Time Periods

Chapter 5, Authentication

Chapter 6, Content Inspection

Chapter 7, Services

Chapter 8, Application Defenses

Chapter 9, Rules

3 Policy Configuration Overview

Contents

[About policy configuration](#)

[A brief guide to planning your policy](#)

[Using groups to simplify policy management](#)

[Examining your policy using the Firewall Policy Report](#)

[About creating rules](#)

About policy configuration

Forcepoint Sidewinder policy is applied primarily by rules, which are made up of many elements. The table below shows the progression of a rule's creation using these elements and their corresponding chapters in this guide.


You are here in the Policy section	Use this chapter to...
Chapter 3, Policy Configuration Overview	understand the policy creation process.
Chapter 4, Network Objects and Time Periods	create or modify any network objects or time periods that will be used by rules.
Chapter 5, Authentication	create or modify authenticators that will be used by rules.
Chapter 6, Content Inspection	configure content inspection methods that will be used by rules.
Chapter 7, Services	create or modify services or service groups that will be used by rules.
Chapter 8, Application Defenses	create or modify Application Defenses that will be used by rules.
Chapter 9, Rules	create rules using the elements you created in the previous chapters in the policy section.

Your site's security policy is implemented and enforced by applying *rules* to all traffic that passes through the Sidewinder. Each rule is basically a mini policy that contains criteria which are used to inspect incoming or outgoing traffic. Rules determine whether that traffic will be allowed to continue to its destination. This section introduces the different ways traffic can be directed through or into the firewall.

Your security policy needs to cover what your organization wants to allow out of its perimeter (outbound traffic), what it wants to allow through its perimeter (inbound traffic), and what is allowed into the Sidewinder (management traffic, such as SSH). When planning your security policy, consider your organization's traffic requirements and how they fit into these categories. If your site has more than two burbs, you may need to create rules that start in one burb and end in another without heading out to the Internet.

The source endpoint specifies where a connection is allowed to initiate. The destination endpoint controls where a connection is allowed to go. When the firewall allows a system to initiate a connection, it automatically allows the response to that session, without needing a separate rule. For example, if you allow outbound HTTP requests, you do not need a separate rule to allow the replies to those requests; the firewall handles this for you.

Figure 23 Types of rules you can use in your security policy



Inbound policy

Inbound rules govern traffic that initiates in an untrusted network area. By default, the firewall does not allow any inbound traffic. Inbound rules represent a prominent threat to your network's security, and therefore should be controlled with authentication or directed to a sacrificial burb that only contains publicly-accessible servers that can recover quickly from attacks and do not contain confidential information. You can also increase your network's security by creating an inbound policy that is as specific as possible. The source should be the smallest possible unit, such as an IP address, or a narrow subnet or IP address range if it cannot be that precise. Other available protections include using authentication and directing traffic into a burb that does not have access to systems with confidential or business-critical information. Your inbound security policy should address any filtering, scanning, or inspection services you want the firewall to provide.

Outbound policy

Outbound rules govern traffic that is allowed to initiate on a protected, trusted burb and then heads for an external destination. Your outbound policy should focus on providing your internal employees and systems access to the resources needed for them to do their jobs. Smaller organizations can probably use the default policy, which includes commonly used services such as HTTP, HTTPS, and RealMedia™. Larger organizations probably need access to more services than those in the default policy and will need to create customized rules. Since the initiators of outbound rules are generally trusted, these rules are less likely to be candidates for filtering services, but there are exceptions. For example, if an attack or virus does manage to get into your network, inspecting outgoing traffic for malicious content can help contain damage to others.

Management policy

A small but important part of your policy will cover traffic that talks directly to a server on the firewall as opposed to passing traffic through it. Sidewinder servers provide the following: management and administration services, routing services, VPN services, communication with external clients, and inter-firewall communication in clustered and enterprise-distributed Sidewinders. Examine the default server rules and determine if your organization will need to enable any of the existing rules or create additional server rules. When creating these rules, make the endpoint as specific as possible to increase security.

A brief guide to planning your policy

Creating an effective security policy requires careful planning and implementation. The following steps are guidelines to creating the policy that is right for your organization:

- 1 Determine your site's overall security policy. This involves writing down detailed instructions about what can and cannot pass through your network perimeter. In most organizations, it is important to get the policy approved by one or more levels of management.
- 2 Once the policy is documented and approved, determine what rules are needed to put the approved policy into place.
- 3 Search your policy for patterns, such as rules that allow access to the same sources and destinations but use different services. These similarities are opportunities to create groups within your policy that will simplify the long-term task of managing your security policy.
- 4 Once all needed rules are identified, order your rule set. Put frequently-used rules at the top and infrequently-used rules at the bottom, as this optimizes processing. Also be sure to put more restrictive rules before less-restrictive rules.

Adequate preparation greatly improves the quality of your security policy and reduces future management overhead.

Using groups to simplify policy management

Using groups can be an efficient way to reduce the footprint of your security policy. A group is a way to set up a one-to-many relationship for elements that have similar security requirements. While a typical rule regulates access for a single element, a single rule that is implemented using groups can regulate access for multiple elements. Once the rules are created, the rules themselves can also be grouped to reduce management overhead. Grouping enables you to reduce the overall number of rules you define, which in turn reduces the complexity of your rule database. A less complex rule database means there is less chance of introducing errors that may affect the integrity of your security policy.

Several rule elements can be grouped to reduce the number of rules in your policy. Once you know what rules you need to implement your security policy, search for patterns of rules having similar requirements, such as traffic from different internal burbs using similar services to reach the Internet. Read the following sections to learn more about grouping different Sidewinder elements.

Service groups

A service group is a group of services of the same type; it cannot contain a mix of proxies, filters, and servers.

A rule will always apply the same properties to all services in a service group. The services in a service group can be either all allowed or all denied. It is not possible to use the same rule to allow access to a subset of services in a service group while at the same time deny access to a different subset of services. Service groups are extremely effective when implemented in a rule that regulates access for a user group or netgroup. Keep in mind, however, that all members in the user group or netgroup must conform to the same security policy (that is they will all be allowed or denied access to the same collection of services).

You can only use both service groups and authentication in a rule if all the services in the group support authentication.

Burb groups

Burb groups are a way to categorize multiple burbs that require a similar security policy. When you select a burb group as an endpoint in a rule, that rule will apply to each burb in the burb group. A source burb or a destination burb cannot contain both a burb group and an individual burb, but can contain multiple burb groups or multiple individual burbs. However, the source and destination can be different. For example, the source could contain one or more individual burbs and the destination could contain one or more burb groups.

Netgroups

Netgroups are a way to use multiple network objects in a single rule. The netgroup can be made up of any combination of available network objects: domain, Geo-Location, host, IP address, IP range, subnet, and netgroup. You may find it more convenient to create all of your network objects before defining your netgroup objects. That way, as you set up your netgroup objects, you will be able to immediately assign the desired network objects to the group.

Application Defense groups

Application Defenses can be grouped to be used in rules that use service groups. When you create an Application Defense group, you select a single Application Defense from each category (for example, HTTP, HTTPS, FTP, etc.) to populate that Application Defense group, although only the Application Defenses that apply to that rule's services will be implemented in the rule.


You can also set an Application Defense group as the default group. The purpose of this group is to be a container for each application's default settings. For example, you would make sure that the each Application Defense in group (HTTP, HTTPS, FTP, etc.) was configured using your site's most common settings for that application. The default Application Defense group is used in all new rules using an Application Defense.

Rule groups

After you plan and create all of the rules you need to enforce your security policy, you can organize them into sets, called *rule groups*. A rule group can consist of both rules and *nested* rule groups. A nested rule group is a rule group that you place within another rule group. You can nest multiple rule groups within a rule group.

[Figure 24](#) demonstrates the basic structure of a rule group that uses nested rules.

Figure 24 Basic rule group structure



Use rule groups to keep rules with similar functions together. This simplifies management overhead for when you need to enable or disable all rules for this function or change their placement in your policy.

Example of using groups in a rule

Here's an example that illustrates the power of using groups. Not all types of groups are used in the example, but the management properties are similar for those groups not included.

Assume you have a netgroup named `eng_netgroup` that consists of all subnets assigned to engineers in your organization. If you want to grant HTTP, FTP, and MS SQL access to this group, you might do so by defining three separate rules.

[Table 9](#) illustrates how these three rules might look in the rule database.

Note: In general, user groups can be used in an allow rule only if the specified service supports authentication (login, Telnet, FTP, HTTP, or secure shell [SSH]). If you want to authenticate other protocols based on user groups, use the Passport authenticator to provide single sign-on access.

Table 9 Typical rules not using groups

Name	Service	Source burb	Source endpoint	Destination burb	Destination endpoint	Application Defense
http_out	HTTP	internal	eng_netgroup	Lab	<any>	HTTP_default
ftp_out	FTP	internal	eng_netgroup	Lab	<any>	FTP_default
mssql_out	MS SQL	internal	eng_netgroup	Lab	<any>	MSSQL_default

A better option, however, is to use a service group. This enables you to accomplish the same thing with *one* rule. Create a service group that contains the HTTP, FTP, and MS SQL proxies, then use this service group when defining the rule. You can also make sure that your default Application Defense group has the proper HTTP, FTP, and MS SQL settings.

[Table 10](#) illustrates the resulting rule using the service group and the default Application Defense group.

Table 10 Sample rule using groups

Name	Service	Source burb	Source endpoint	Destination burb	Destination endpoint	Application Defense
eng_to_lab	EngServGrp (HTTP, FTP, MS SQL)	internal	eng_netgroup	Lab	<any>	default_group

Examining your policy using the Firewall Policy Report

You can open a report in a web browser showing comprehensive details of your Sidewinder policy.

Select **Monitor > Firewall Policy Report**. Click the **Firewall Policy Report** link to open the report in a web browser.

About creating rules

Rules are made up of many elements, as explained in “About rules” on page 257. When you are creating a new rule, you may need to create several new rule elements that will be used by the new rule.

The steps required to create a rule are shown below. Some of these steps may be unnecessary for your particular policy. For example, if the rule you are creating does not require authentication or the service that the rule will use is already configured, you can skip some steps.

Steps to create a rule
1 Create or modify any network objects or time periods that will be used by the rule. See Chapter 4, Network Objects and Time Periods .
2 Create or modify an authenticator that will be used by the rule. See Chapter 5, Authentication .
3 Configure any content inspection methods that will be used by the rule. See Chapter 6, Content Inspection .
4 Create or modify a service or service group that will be used by the rule. See Chapter 7, Services .
5 Create or modify an Application Defense that will be used by the rule. See Chapter 8, Application Defenses .
6 Create the rule, using the elements you created in steps 1–5. See Chapter 9, Rules .

4 Network Objects and Time Periods

Contents

[Creating network objects](#)

[Creating time periods](#)

Creating network objects

Forcepoint Sidewinder policy is applied primarily by rules, which are made up of many elements. The table below shows the progression of a rule's creation using these elements and their corresponding chapters in this guide.

You are here in the Policy section	Use this chapter to...
Chapter 3, Policy Configuration Overview	understand the policy creation process.
Chapter 4, Network Objects and Time Periods	create or modify any network objects or time periods that will be used by rules.
Chapter 5, Authentication	create or modify authenticators that will be used by rules.
Chapter 6, Content Inspection	configure content inspection methods that will be used by rules.
Chapter 7, Services	create or modify services or service groups that will be used by rules.
Chapter 8, Application Defenses	create or modify Application Defenses that will be used by rules.
Chapter 9, Rules	create rules using the elements you created in the previous chapters in the policy section.


A network object is the source or destination of a connection to or through the Sidewinder. A network object can be any of the following:

- Domain
- Geo-Location
- Host
- IP address
- IP range
- Netmap
- Subnet
- Netgroup

Each network object that you create is available for selection from the source and destination Endpoint drop-down lists on the Rules window.

To view, create, and maintain network objects, select **Policy > Rule Elements > Network Objects**. The Network Objects window appears.

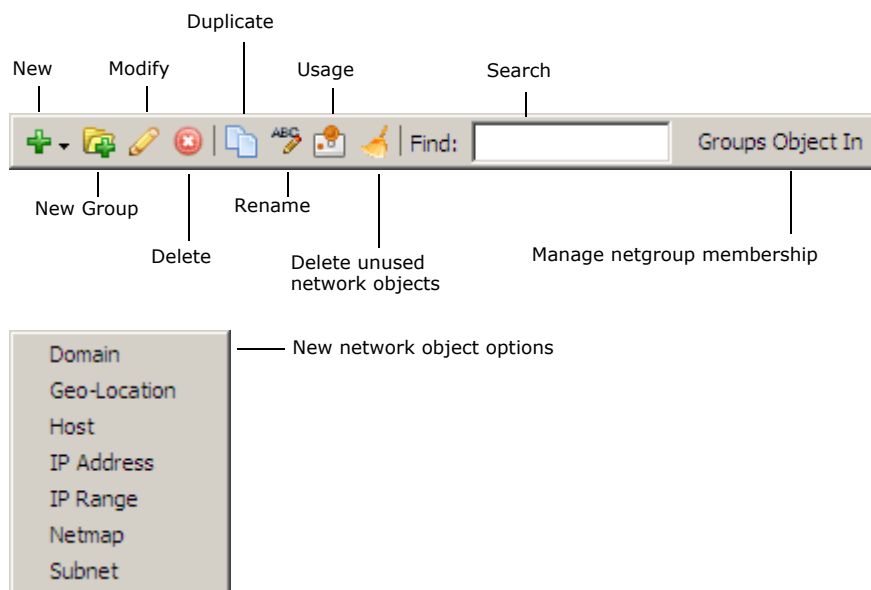
Figure 25 Network Objects window



Type	Name	Value	Description
Host	localhost	localhost	The local system.
IP Address	Firewall	127.0.0.1	Firewall network object
IP Address	external primary DNS resolver	10.65.240.246	Network object for external burb primary DNS resolver
Netgroup	DNS resolvers	external primary DNS resolver	Network group for DNS resolvers

This window lists the network objects currently configured on the Sidewinder.

Figure 26 Network Objects toolbar



Use the toolbar to perform the tasks listed in [Table 11](#).

Table 11 Network Objects toolbar tasks

Icon	Action
New	<p>Create a network object by clicking New and selecting an object from the drop-down menu. Configure the selected Network Objects window that appears:</p> <ul style="list-style-type: none"> • Domain – For information on creating a domain object, see About the Network Objects: Domain window. • Geo-Location – For information on creating a host object, see About the Network Objects: Geo-Location window. • Host – For information on creating a host object, see About the Network Objects: Host window. • IP Address – For information on creating an IP address object, see About the Network Objects: IP Address window. • IP Range – For information on creating an IP range object, see About the Network Objects: IP Range window. • Netmap – For information on creating a netmap object, see About the Network Objects: Netmap window. • Subnet – For information on creating a subnet object, see About the Network Objects: Subnet window.
New Group	<p>Create a netgroup by clicking New Group. The Netgroup window appears. See About the Network Objects: Netgroup window for more information.</p>
Modify	<p>Modify an existing network object or netgroup by selecting it from the list and clicking Modify. Make your changes in the pop-up window.</p> <p>(Read-only administrators can click View to view a network object or netgroup.)</p>
Delete	<p>Delete an existing network object or netgroup by selecting it in the list and clicking Delete.</p>
Duplicate	<p>Create a duplicate of an existing network object or netgroup by selecting it in the list and clicking Duplicate. Change the name and make any desired changes, then click Add.</p>
Rename	<p>Rename a network object or netgroup by selecting it in the list and clicking Rename. Type the new name in the pop-up window and click OK.</p>
Usage	<p>View the areas (netgroup, netmap, proxy rule) that are currently using a particular network object or netgroup by selecting it in the list and clicking Usage.</p>
Delete unused network objects	<p>Delete objects that are not in use by clicking Delete unused network objects. The Delete unused objects window appears. Select the objects that you want to delete and then click OK.</p>

Table 11 Network Objects toolbar tasks <Comment>(continued)

Icon	Action
Find	Search for specific elements in the list by typing your search criteria in the Find field. Objects with matching elements appear in the list.
Group Objects In	View or modify the group membership of a network object by selecting it and then clicking Groups Object In . See Managing netgroup membership for more information.


About the Network Objects: Domain window

Use this window to define information about a domain. Each domain you define becomes a network object that can be used in a rule.

Domain objects have features that set them apart from other network objects. Before using domain objects in rules, note the following:

- Since domains are dependent on DNS, which is out of your control, the use of domain network objects can be a security risk.
- Domain objects require a DNS lookup and therefore incur a DNS performance penalty each time they are used.
- For a proxy rule that includes a domain object to be processed correctly, that rule must be placed after the last filter rule.

Figure 27 Network Objects: Domain window



- **Name** – Type a name for this domain object (for example, “example” for *example.com*).
 - Valid values include alphanumeric characters, periods (.), dashes(-), underscores (_), and spaces ().
 - The name cannot exceed 100 characters.

Note: The name you create here is what you will see in the Endpoint drop-down list when you create a rule. You will not see any of the object’s values, so make a descriptive name to ensure that you will recognize it in the Rules window.

This field cannot be edited if you are *modifying* an existing domain.

- [Optional] **Description** – Enter any useful information for this domain object.
- **Domain** – Enter the domain to use for this object (for example, *example.com*).

Click **Add** to add the domain object, or **OK** if you modified an existing domain object.

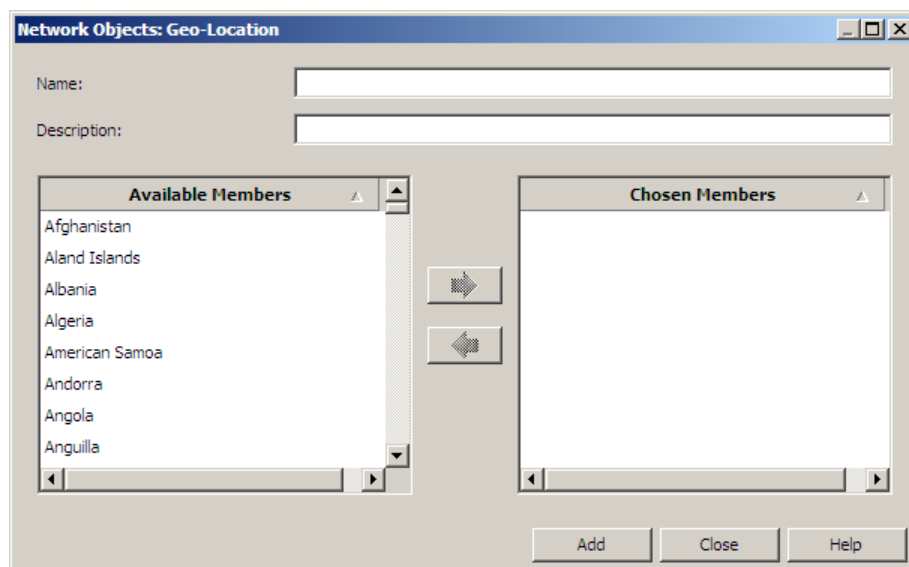
About the Network Objects: Geo-Location window

Use this window to define a Geo-Location object. Each Geo-Location object you define becomes a network object that can be used in a rule.

Geo-Location identifies the country of origin of an IP address. Use a Geo-Location object in a rule to allow or deny a network connection based on the source or destination country.

Note: Periodically update the Geo-Location database to ensure that you have the latest country database. See “Updating the Geo-Location database” on page 146 for information.

Figure 28 Network Objects: Geo-Location window



- **Name** – Enter a name for this Geo-Location object.
 - Valid values include alphanumeric characters, periods (.), dashes(-), underscores (_), and spaces ().
 - The name cannot exceed 100 characters.

Note: The name you create here is what you will see in the Endpoint drop-down list when you create a rule. You will not see any of the object’s values, so make a descriptive name to ensure that you will recognize it in the Rules window.


- **Description** – Enter any useful information for this Geo-Location object.
- The **Available Members** list displays all of the countries that you can add to this Geo-Location object. The **Chosen Members** list displays the countries that are currently members of this netgroup.
 - To add a country to this Geo-Location object, select the desired country in the **Available Members** list, then click the > arrow button to move it to the **Chosen Members** list.
 - To remove a country from this Geo-Location object, select the desired country in the **Chosen Members** list, then click the < arrow button.
 - To add or remove multiple consecutive countries at one time, select the first country, then press the **Shift** key while selecting the last country. To add or remove multiple non-consecutive countries at one time, press the **Ctrl** key while selecting each desired country.

Click **Add** to add the Geo-Location object, or **OK** if you modified an existing Geo-Location object.

About the Network Objects: Host window

Use this window to define information about a host. Each host you define becomes a network object that can be used in a rule.

Figure 29 Network Objects: Host window



Note: In IP filter rules, the localhost network object is supported, but DNS-resolvable host names should be avoided. DNS-resolvable host names become inoperative during any periods when the appropriate DNS server is unavailable or unreachable.

- **Name** – Type a name for the host.
 - Valid values include alphanumeric characters, periods (.), dashes(-), underscores (_), and spaces ().
 - The name cannot exceed 100 characters.

Note: The name you create here is what you will see in the Endpoint drop-down list when you create a rule. You will not see any of the object's values, so make a descriptive name to ensure that you will recognize it in the Rules window.

This field cannot be edited if you are *modifying* an existing host.

- [Optional] **Description** – Enter any useful information about this host.
- **Host** – Enter the hostname for this host object (for example, *mail.example.com*).
- **DNS** – Determine whether this host will use DNS:
 - **DNS** – Select this option to perform normal DNS look-ups.
 - **No DNS** – Select this option if you do not want to perform DNS lookups for this host.

Note: The dig (Domain Information Groper) command gathers information from DNS based on an IP address, and obtains the corresponding host name. A dig is useful in determining if a host is resolvable before creating a network object.

```
dig -x ipaddress any any
```

- [Conditional] **Override TTL** – If you selected DNS and you need to override the DNS time-to-live value, select this check box. Enter a time value and select a time increment for the new time-to-live value.

Note: Overriding the default DNS time-to-live value is not recommended.


- **IP Addresses For The Host** – To create and maintain IP addresses for a host, you can do the following:
 - Click **New**, then type the IP address in the pop-up window.
 - Select an IP address, then click **Modify** and type a replacement IP address in the pop-up window.
 - Select an IP address, then click **Delete** to delete an IP address.

Click **Add** to add the host object, or **OK** if you modified an existing host object.

About the Network Objects: IP Address window

Use this window to define information about an IP address. Each IP address you define becomes a network object that can be used in a rule.

Figure 30 Network Objects: IP Address window



- **Name** – Type a name for the IP address.
 - Valid values include alphanumeric characters, periods (.), dashes(-), underscores (_), and spaces ().
 - The name cannot exceed 100 characters.

Note: The name you create here is what you will see in the Endpoint drop-down list when you create a rule. You will not see any of the object's values, so make a descriptive name to ensure that you will recognize it in the Rules window.

This field cannot be edited if you are *modifying* an existing IP address.


- [Optional] **Description** – Enter any useful information about this IP address object.
- **IP Address** – Type the value of the IP address. To find the IP address for a host name, type the name and click **DNS Lookup**.

Click **Add** to add the IP address object, or **OK** if you modified an existing IP address object.

About the Network Objects: IP Range window

Use this window to define information about an IP range. The IP range you define becomes a network object that can be used in a rule.

Figure 31 Network Objects: IP Range window



- **Name** – Type a name for the IP range.
 - Valid values include alphanumeric characters, periods (.), dashes(-), underscores (_), and spaces ().
 - The name cannot exceed 100 characters.

Note: The name you create here is what you will see in the Endpoint drop-down list when you create a rule. You will not see any of the object's values, so make a descriptive name to ensure that you will recognize it in the Rules window.

This field cannot be edited if you are *modifying* an existing IP range.

- [Optional] **Description** – Enter any useful information about this IP range object.
- **Starting IP Address** – Type the value of the IP address at the beginning of the range. To find the IP address for a host name, type the name and click **DNS Lookup**.
- **Ending IP Address** – Type the value of the IP address at the end of the range. To find the IP address for a host name, type the name and click **DNS Lookup**.

Click **Add** to add the IP range object, or **OK** if you modified an existing IP range object.

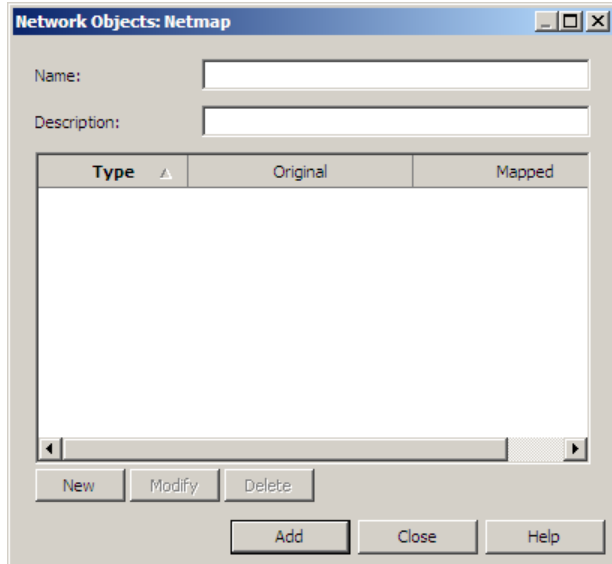
About the Network Objects: Netmap window

Use this window to define information about a netmap. Each netmap you define becomes a network object that can be used in a rule.

Netmap objects allow you to map multiple IP addresses and subnets to alternate addresses without creating numerous rules.

- A netmap consists of one or more netmap members.
- A netmap member is any IP address or subnet that you add to a netmap.
- Each member in the netmap is mapped to an alternate address or subnet that you specify.

Figure 32 Network Objects: Netmap window



- **Name** – Type a name for the netmap.
 - Valid values include alphanumeric characters, periods (.), dashes(-), underscores (_), and spaces ().
 - The name cannot exceed 100 characters.

Note: The name you create here is what you will see in the Endpoint drop-down list when you create a rule. You will not see any of the object's values, so make a descriptive name to ensure that you will recognize it in the Rules window.

This field cannot be edited if you are *modifying* an existing netmap.


- [Optional] **Description** – Enter any useful information for this netmap.
- **Netmap members list** – This list displays existing netmap members. You can perform the following actions:
 - **Create a new netmap member** – Click **New** and make a selection in the pop-up menu to create a netmap member.
 - **IP Address** – Select this option if you want to map an IP address to a different IP address.
 - **Subnet** – Select this option if you want to map a subnet address to a different subnet address.
 - **Modify an existing netmap member** – Select a netmap member in the list and click **Modify**, then make the desired selections in the pop-up window.
 - **Delete an existing netmap member** – Select a netmap member in the list and click **Delete**.
 - **Sort** – Click a column heading to sort the list by that column's content. Click again to reverse the sort order.

Click **Add** to add the netmap information, or **OK** if you modified an existing netmap.

About the Netmap Members: IP Address/Subnet Netmap Selections window

Use the IP Address Netmap/Subnet Selections window to map an IP address or a subnet to an alternate address within a netmap.

Figure 33 Netmap Members window




- 1 In the **Original** list, select the IP or subnet address that you want to map to a different address.
- 2 In the **Mapped** list, select the IP address that the original IP address will be mapped to, or select a subnet address of the same size that the original subnet address will be mapped to.
- 3 Click **Add**.

About the Network Objects: Subnet window

Use this window to define information about a subnet. Each subnet you define becomes a network object that can be used in a rule.

Figure 34 Network Objects: Subnet window



- **Name** – Type a name for the subnet.
 - Valid values include alphanumeric characters, periods (.), dashes(-), underscores (_), and spaces ().
 - The name cannot exceed 100 characters.

Note: The name you create here is what you will see in the Endpoint drop-down list when you create a rule. You will not see any of the object's values, so make a descriptive name to ensure that you will recognize it in the Rules window.

This field cannot be edited if you are *modifying* an existing subnet.

- [Optional] **Description** – Type any useful information about the object.
- **Subnet** – Enter the following information:
 - In the text field, type the subnet address.
 - In the numeric text box, enter the number of significant bits for the subnet address. You must enter an integer value in the range 0–32 (IPv4) or 0–128 (IPv6). For example, if you enter 16, only the first 16 bits of the address are important.

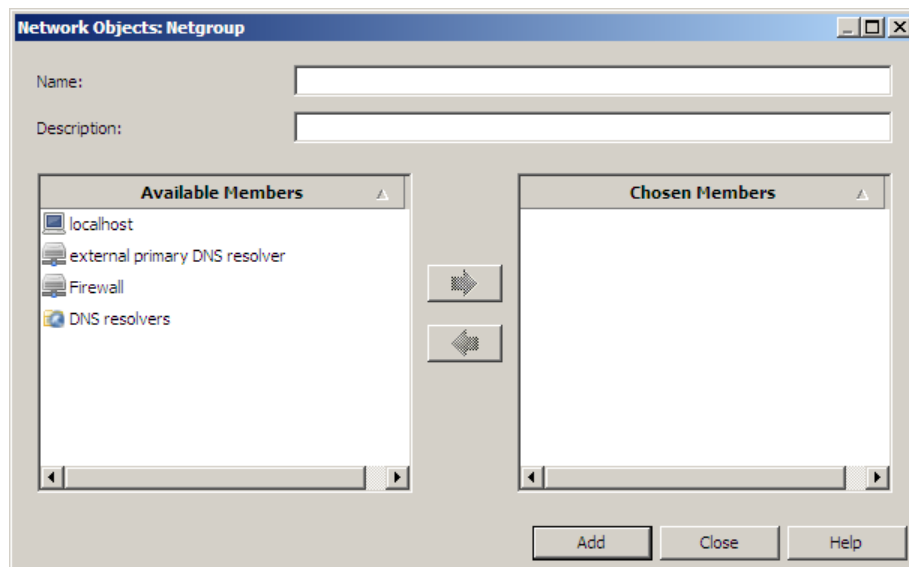
Click **Add** to add the subnet object, or **OK** if you modified an existing subnet object.

About the Network Objects: Netgroup window

Use this window to define information about a netgroup. Each group you define becomes a network object that can be used in a rule.

Tip: You may find it more convenient to create all of your network objects before defining your netgroup objects. That way, as you set up your netgroup objects, you will be able to immediately assign the desired network objects to the group.

Figure 35 Network Objects: Netgroup window



- **Name** – Type a name for the netgroup. The name will be used by rules to identify the netgroup when you set up Sidewinder connections.
 - Valid values include alphanumeric characters, periods (.), dashes(-), underscores (_), and spaces ().
 - The name cannot exceed 100 characters.

Note: The name you create here is what you will see in the Endpoint drop-down list when you create a rule. You will not see any of the object's values, so make a descriptive name to ensure that you will recognize it in the Rules window.

This field cannot be edited if you are *modifying* an existing group.


- [Optional] **Description** – Enter any useful information about this group.
- The **Available Members** list displays all of the network objects that you can add to this netgroup. The **Chosen Members** list displays the network objects that are currently members of this netgroup.
 - To add a member to this netgroup, select the desired member in the **Available Members** list, then click the > arrow button to move it to the **Chosen Members** list.
 - To remove a member from this netgroup, select the desired member in the **Chosen Members** list, then click the < arrow button.
 - To add or remove multiple consecutive members at one time, select the first member, then press the **Shift** key while selecting the last member. To add or remove multiple non-consecutive members at one time, press the **Ctrl** key while selecting each desired member.

Click **Add** to add the netgroup, or **OK** if you modified an existing netgroup.

Example of rules using netgroups

For the configuration shown in [Figure 36](#), the Sidewinder administrator has grouped all internal systems into one of three netgroups: marketing (mkt_net_group), engineering (eng_net_group), and accounting (acct_net_group).

Figure 36 Sample network configuration



Suppose you want to allow all groups access to external FTP sites but only the engineering group access to FTP host 192.55.12.3. [Table 12](#) shows the rules in the order that they should be added to the rule group. The following table shows the rules in the order that they should be added to the rule group.

Table 12 Rules for sample configuration shown in the figure above

Rule Criteria	Rule 1: allow_eng_ftp	Rule 2: deny_other_ftp	Rule 3: allow_oth_ftp
Service	FTP	FTP	FTP
Action	Allow	Deny	Allow
Source Burb	internal	internal	internal
Source Endpoint	eng_net_group	<Any>	<Any>
Destination Burb	external	external	external
Destination Endpoint	192.55.12.3	192.55.12.3	<Any>
Authenticator	SafeWord		
User Group	any (leave blank)	any (leave blank)	any (leave blank)
Time Period	Fri 7am-7pm		
Application Defense (FTP)	Allow Put/Get	Deny All	Allow Put/Get

The following list summarizes key points to consider for the proxy rules listed in [Table 12](#).


- Rule 1 allows all systems in the engineering group authenticated FTP access to IP address 192.55.12.3 on the Internet, but only on Friday between 7:00 a.m. and 7:00 p.m.
- This rule requires users to authenticate themselves via SafeWord before an FTP connection is allowed.
- Rule 2 denies all systems in the trusted burb named *internal* from FTP service to IP address 192.55.12.3 on the Internet.
- Rule 3 allows FTP service from all systems in the *internal* trusted burb to any external system in the Internet burb.

Managing netgroup membership

You can add or remove members in an existing group in two ways:

- In the Network Objects window, select the desired *netgroup* from the list and click **Modify**, then make the membership changes in the Netgroup window. See [About the Network Objects: Netgroup window](#).
- In the Network Objects window, select a *network object* from the list and click **Groups Object In**. The Group Membership window appears.

Figure 37 Group Membership window



Use the Group Membership window to see which groups the object belongs to and to add or remove the object from group membership.

The **Available** list displays all the available groups. The **Selected** list displays the groups to which the object currently belongs.

- To add this network object to another group, select the group in the **Available** list, then click the **>** arrow button to move it to the **Selected** list.
- To remove a network object from a group, select the group in the **Selected** list, then click the **<** arrow button to move the group to the **Available** list.
- To select multiple consecutive entries, press the **Shift** key while selecting the groups. To select multiple non-consecutive entries, press the **Ctrl** key while selecting the desired entries.

When you are finished, click **OK**.

Creating time periods

A time period is a rule element that can specify a segment of time a rule is in effect. The time periods you create here can be selected from the **Time period** drop-down list on the Rule window.

To create time periods for rules, select **Policy > Rule Elements > Time Periods**. The Time Periods window appears.

Figure 38 Time Periods window

The screenshot shows the 'Time Periods' window with a standard toolbar at the top. The main area is divided into two panes. The upper pane contains a table with three columns: 'Name', 'Days and Times', and 'Description'. The lower pane shows the configuration for the selected 'Business Hours' period, including input fields for 'Name' and 'Description', a 'Days and Times' list box, and 'New', 'Modify', and 'Delete' buttons.

Name	Days and Times	Description
Business Hours	Mon - Fri, 8:00 AM - 5:00 PM	
Weekdays	Mon - Fri, All day	
Weekends	Sat - Sun, All day	

Name: Business Hours

Description:

Days and Times

Mon - Fri, 8:00 AM - 5:00 PM

New Modify Delete

The upper pane lists the existing time periods. The lower pane shows the settings for the selected time period.

Use the toolbar to perform the actions listed in [Table 13](#).

Figure 39 Time Periods toolbar




Table 13 Time Periods toolbar tasks

Icon	Action
New	Create a new time period by clicking New . A pop-up window appears where you set the appropriate properties. See About the New/Modify Time Period: New Days and Times window for more information.
Modify	Modify an existing time period by selecting a time period from the list and modifying the settings in the lower pane. To modify the settings in a pop-up window, click Modify . (Read-only administrators can click View to view a time period.) See About the New/Modify Time Period: New Days and Times window for more information.
Delete	Delete an existing time period by selecting a time period from the list and clicking Delete .
Rename	Rename an existing time period by selecting a time period from the list and clicking Rename . Type a new name in the pop-up window.
Usage	View which rules are using an existing time period by selecting a time period from the list and clicking Usage . A pop-up window shows which rules use the selected time period.
Find	Search for specific elements in the list by typing your search criteria in the Find field. Time periods with matching elements appear in the list.

You can make the following modifications in the lower pane:

Description – Type a description of the time period to further identify it.

Days and Times – This list shows the parameters of the time period.

- **New** – Click this button to set day and time parameters for this time period.
- **Modify** – Click this button to modify the selected days and times.
- **Delete** – Click this button to delete the selected days and times.

About the New/Modify Time Period: New Days and Times window

Use this window to set the day and time parameters of a rule.

Figure 40 New/Modify Days and Times window

New Days and Times

☒ Continuous time period (Tues 7 AM - Wed 6 PM)

Start: Sunday 12:00 AM

End: Sunday 12:00 AM

☐ Recurring time period (Mon - Fri, 7 AM - 6 PM)

Days: Sun Mon Tue Wed Thu Fri Sat

Start: 12:00 AM ☐ All day

End: 12:00 AM

Add Close Help

- **Continuous time period** – Select this option to make a rule active for one episode per week.
 - **Start** – Select the day and time that the rule will become active each week.
 - **End** – Select the day and time that the rule will become inactive until the following week.
- **Recurring time period** – Select this option to make a rule active on specified days and times every week.
 - **Days** – Select the days that this rule will be active each week.
 - **Start** – Set the time that this rule will become active each selected day.
 - **End** – Set the time that this rule will become inactive each selected day.
 - **All day** – Select this option to make the rule active 24 hours of each selected day.

5 Authentication

Contents

[Understanding authentication](#)

[Configuring an authenticator](#)

[Telnet and FTP considerations](#)

[Setting up users to change their own passwords](#)

[Authenticating groups from an external group source](#)

[Authenticating groups from an internal group source](#)

Understanding authentication

Forcepoint Sidewinder policy is applied primarily by rules, which are made up of many elements. The table below shows the progression of a rule's creation using these elements and their corresponding chapters in this guide.

You are here in the Policy section	Use this chapter to...
Chapter 3, Policy Configuration Overview	understand the policy creation process.
Chapter 4, Network Objects and Time Periods	create or modify any network objects or time periods that will be used by rules.
Chapter 5, Authentication	create or modify authenticators that will be used by rules.
Chapter 6, Content Inspection	configure content inspection methods that will be used by rules.
Chapter 7, Services	create or modify services or service groups that will be used by rules.
Chapter 8, Application Defenses	create or modify Application Defenses that will be used by rules.
Chapter 9, Rules	create rules using the elements you created in the previous chapters in the policy section.

Authentication refers to a process that validates a person's identity before he or she is allowed to pass traffic through the firewall.

Depending on the authentication method used, a person must provide a user name and valid password and/or a special passcode or personal identification number (PIN) before being logged into a server. If a user enters an invalid password, passcode, or PIN, then the policy will not pass network traffic.

Who gets authenticated

Sidewinder authenticates two types of users:

- Administrators connecting *to* the firewall
- Proxy users connecting *through* the firewall

Administrator authentication

This is for administrators who maintain or audit the firewall. Administrators log directly into the firewall.

- The initial administrator account, including user name and password for login authentication to the firewall, is created during startup configuration using the Quick Start Wizard.
- Additional administrator accounts can be created or modified on the Administrator Accounts window.
- Administrators can use SSH to access a firewall remotely via a command line interface.

Note: We recommend using a strong authentication method for administrators logging in remotely.

Proxy authentication

This is for network users attempting to create a proxy connection from one side of the firewall to the other.

- You can authenticate internal-to-external, external-to-internal, and internal-to-internal connections.
- You can authenticate access for any service through the firewall.

- You can allow access to multiple services with a single successful authentication method by using Passport (also known as single sign-on).
- You can require authentication by selecting an authentication method on the Rules window when you create a rule.
- You can set up authentication on a user-by-user basis. Some authenticators allow you to create user groups to identify multiple users by a single name, or to add groups from an external authentication server. You can assign groups to use an authentication method for a rule in the Rules window.

See [Configuring an authenticator](#).

Weak and strong authentication

An authentication method is *weak* or *strong*, depending on the level of security it provides.

Weak authentication

An example of a weak authentication method is a fixed password, which only requires a user to enter the same password every time they log in. Even if the user carefully chooses a random password, an attacker can sniff the password as it is transmitted and masquerade as the user.

Because your internal network is thought to be trusted, fixed passwords can be adequate for internal-to-external authentication.

Strong authentication

Strong authentication uses a variety of methods to keep passwords secure. A hardware token, for example, generates a different password each time it is used.

Using multiple factors can also strengthen authentication. For example, the hardware token can require a PIN, so that the user must authenticate using something they have (the token) and something they know (the PIN).

Strong authentication is generally desired for external-to-internal proxy connections and for external administration access to the firewall.

Types of authentication methods

Sidewinder supports the following authentication methods:

- **Passport** – Passport (also known as single sign-on) associates an authenticated user with their IP address. A successful Passport authentication caches the source IP address for a specified time. Subsequent connection attempts from the same IP address are allowed without prompting for authentication.
Security level: Weak
- **Password** – Standard password authentication requires a user to enter the same password each time he or she logs in.
Security level: Weak
- **LDAP (Lightweight Directory Access Protocol)** – Four types of LDAP authentication are available: iPlanet, Active Directory, OpenLDAP, and Custom LDAP.
Security level: Weak
- **Common Access Card** – Use this authenticator to log into a Sidewinder using a U.S. Department of Defense Common Access Card (CAC).
Security level: Strong
- **Windows Domain** – You can use this authenticator if your organization operates a Windows primary domain controller (PDC) or backup domain controller (BDC).
Security level: Weak
- **RADIUS** – You can use this authenticator if your organization operates a RADIUS server.
Security level: Varies with authentication server and method
- **SafeWord** – SafeWord RemoteAccess and SafeWord PremierAccess interoperate with Sidewinder.
Security level: Varies with authentication server and method

See [Configuring an authenticator](#) for more information.

Alternate authentication methods

You can select only one authenticator in a rule. If you want alternate authentication methods for a service—for example, to ensure that you can connect to the Admin Console if an authentication server is down—you can create more rules for that connection.

To use an alternate authentication method:

- 1** Duplicate the rule allowing that connection and select a different authenticator for the duplicated rule.
- 2** Specify the alternate authentication method when logging in:
 - If it is an Admin Console connection, select the alternate method from the Authenticator drop-down list on the Login window.
 - If it is another service, at the login prompt, enter your user name followed by a colon and the name of the alternate authenticator:

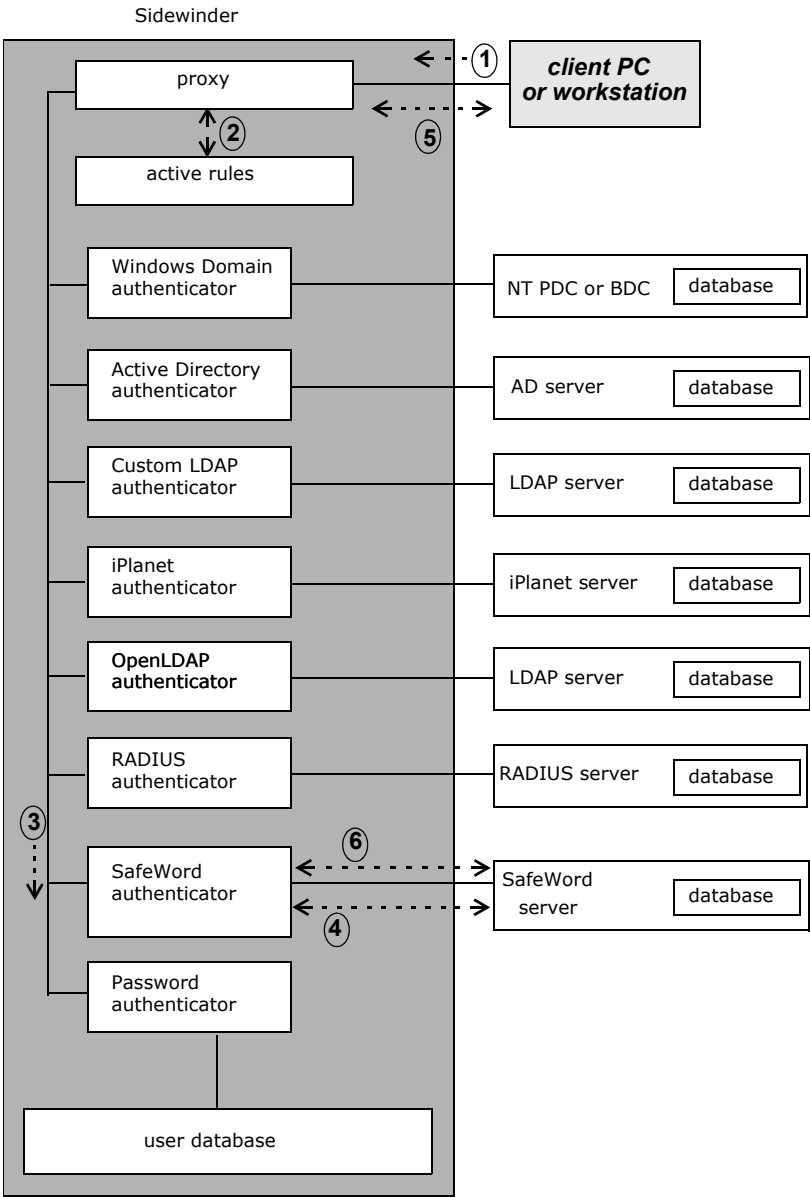
login_name:authenticator

Authentication scenario

In the following scenario, the user is authenticated using SafeWord PremierAccess, which implements a strong challenge/response authentication process. See [Figure 41](#) for an illustration. (Note that the process is different for other authentication methods.)

- 1** A user tries to make a network connection via Telnet or FTP.
- 2** The firewall checks the rules to determine whether the connection between the source and destination addresses is allowed and to determine which authenticator to use.
- 3** If the connection is allowed, the proxy contacts the appropriate authenticator in the firewall.
- 4** The authenticator passes the login request to the appropriate authentication server. The authentication server checks the database to verify the user's login name is registered.
- 5** The login challenge is sent to the user. Using client software or a hardware authenticator (token), the user types in the proper response to the prompt.
- 6** The firewall sends the response to the authentication server. The authentication server checks the response and informs the firewall to either accept or reject the login request.

Figure 41 Authentication servers supported by Sidewinder



Configuring an authenticator

Authenticators validate a person's identity before he or she is allowed to pass traffic through the firewall.

Authenticators are configured on the Authenticators window. They can then be selected on the Rules window to authenticate proxy connections.

To configure an authenticator, select **Policy > Rule Elements > Authenticators**. The Authenticators window appears.

Figure 42 Authenticators window

Name	Type	Properties	Description
Passport	Passport		Static Passport authenticator
Password	Password		Static password authenticator

Name: Type: Description:

General | Users and User Groups

Establish Passport Credentials

Authenticators to establish Passport credentials:

- ☒ Password

Default authenticator:

☒ Require Web login

☐ Active session mode

Refresh period: seconds

Grace period: seconds

Redirect delay: seconds

Web login

Port:

Login page:

Logout page:

Redirect page:

Error page:

Passport credential timeouts

Authenticate inactive users every:

Force authentication every:

Use this window to create, modify, and delete authenticators that validate login attempts by administrators and proxy users. The upper pane lists the existing authenticators. When you select an authenticator in the list, the properties of that authenticator appear in the lower pane.

Note: Passport and Password are default authenticators and cannot be deleted. They can be sorted with the rest of the list.

Figure 43 Authenticators toolbar



Use the toolbar and table in the upper pane to perform the following actions:

Table 14 Authenticators toolbar

Icon	Action
New	Create a new authenticator by clicking New and selecting an authenticator from the drop-down menu. A pop-up window appears where you set the appropriate properties.
Modify	To modify an existing authenticator, select an authenticator from the list and change the settings in the lower pane. To modify the settings in a pop-up window, click Modify . (Read-only administrators can click View to view an authenticator.)
Delete	Delete an existing authenticator by selecting it in the list and clicking Delete . You cannot delete an authenticator if it is referenced by a rule.
Rename	Rename an authenticator by selecting it in the list and clicking Rename . Type the new name in the pop-up window and click OK . You cannot rename an authenticator if it is referenced by a rule.
Usage	View the areas that are currently using a particular authenticator by selecting it in the list and clicking Usage .
Find	Search for specific elements in the list by typing your search criteria in the Find field. Objects with matching elements appear in the list.
Manage Authentication Failures	Configure the authentication failure lockout feature by clicking Manage Authentication Failures . This opens a window where you can configure the firewall to block access to a user if the number of consecutive failed authentication attempts reaches a configured number. This protects unauthorized users from multiple attempts at guessing a user's password.

For setting up specific authenticators, see the following:

- [Setting up Passport authentication](#)
- [Setting up standard password authentication](#)
- [Setting up LDAP authentication](#)
- [Setting up CAC authentication](#)
- [Setting up Windows domain authentication](#)
- [Setting up RADIUS authentication](#)
- [Setting up SafeWord authentication](#)

Setting up Passport authentication

Passport (also known as single sign-on) associates an authenticated user with their IP address. A Passport is acquired by successfully logging in using a designated authenticator.

- 1 On the Authenticators window, you configure and select authenticators that can be used to acquire a Passport.
- 2 In the New Rule window, you select Passport as the authentication method for a network connection.
- 3 After a user successfully authenticates the network connection using a designated authenticator, they acquire a Passport and their IP address is cached for a specified time. Subsequent connection attempts from the same IP address are assumed to be from the same authenticated user, and if Passport is the authentication method for the rule, the connection is allowed without prompting for authentication.

Uses for Passport

Passport can be used in the following ways:

- **Authenticator groups** – You can designate a group of authenticators that can acquire a Passport. If Passport is the authentication method in a rule, any of the selected authenticators can be used to authenticate the connection and acquire a Passport.
- **Require a web login** – You can require an HTTP connection to acquire a Passport. Users are redirected from a web request to an authentication login page, or they can go directly to the web login page. Passport authentication for other connection types are denied.

After a user has been authenticated, a “Successful Login” browser window appears and the user is redirected to the requested web page. Any type of connection with a Passport authentication method is then allowed for the life of the Passport.

- **Active session mode** – You can use active session mode with web login to require the Passport holder to maintain an open network connection to the firewall. This increases security when multiple users share the same IP address, for example, if a computer is shared or if users connect through a VPN concentrator.

When active session mode is enabled, the “Successful Login” browser window must remain open during the life of the Passport. Other browser windows must be used to access web sites. If the user was redirected to the web login page, the “Successful Login” browser window contains a link to the requested web page.

A heartbeat message periodically tests the HTTPS connection and refreshes the “Successful Login” web page. If the connection is broken, the Passport is revoked. The Passport can also be revoked by clicking **Stop** on the browser window, closing the browser window, or rebooting the computer. When a Passport is revoked, all of the sessions that were authorized by that Passport are closed.

- **Other authentications** – Because a Passport holder does not need to be authenticated for subsequent connections, Passport can be used for encrypted services or for services that do not have an authentication mechanism, such as ping.

Revoking a Passport

Passports can be revoked in these ways:

- A Passport can expire after a configured time has passed.
- A user can be prompted to re-authenticate after a configured idle period.
- An administrator can revoke a Passport directly.

Configuring Passport

To set up Passport authentication: In the list in the upper pane of the Authenticators window, select **Passport**. The Passport: General tab appears in the lower pane.

- Passport is a default authenticator. It cannot be deleted.
- The Passport rule is part of the initial active policy of the firewall and is enabled by default. The rule allows authentication to the Passport server. Do not modify this rule.

Figure 44 Passport: General tab

To configure the Passport authenticator:

- 1 [Optional] Enter identifying information: In the **Description** field, you can modify the description to help you more easily identify this authenticator. You cannot change the name or type for Passport.
- 2 Select Passport authenticators:
 - a In the **Authenticators to establish Passport credentials** list, select the authenticators that can be used to acquire a Passport. Configured authenticators populate this list.
 - b From the **Default authenticator** drop-down list, select the authenticator used by default for connections that have Passport as the authenticator.
 - The default authenticator should be the authenticator most commonly used by users.
 - Other authenticators selected in the **Authenticators to establish Passport credentials** list can be used to authenticate a connection and acquire a Passport. See [Using an alternate Passport authentication method](#) for instructions.
 - If the default authenticator is the authentication method in a rule, a successful authentication does not acquire a Passport. Passport must be the selected authentication method in the rule.
- 3 [Optional] To require an HTTP connection to acquire a Passport, select **Require Web login**.

Users are redirected from a web request to the authentication login page. Passport authentication for other connection types are denied. After a user has been authenticated, a “Successful Login” browser window appears and the user is redirected to the requested web page.

4 [Optional] To require the Passport holder to maintain an open network connection to the firewall, select **Active session mode**.

- Use the **Refresh period** field to configure how frequently a heartbeat message is sent to the “Successful Login” web page. A heartbeat message periodically tests the HTTPS connection and refreshes the page. If the connection is broken, the Passport is revoked.

Note: Time-outs vary among web browsers. A high refresh period could result in revoked Passports for some browsers due to the HTTPS connection timing out.

- Use the **Grace period** field to configure how many seconds the HTTPS connection can be broken before the Passport is revoked.

5 Select how long a web redirect page remains open after a successful Passport login: In the **Redirect delay** field, enter or select the appropriate number of seconds.

If a user makes a web request and has not yet been authenticated for Passport, they are redirected to the authentication login page. After successful authentication, the “Successful Login” browser window states that the user will be redirected to the requested page in the configured number of seconds.

This option is not available for active session mode. If **Active session mode** is enabled, the Successful Login window contains a link to open the requested page in a new browser window.

6 [Optional] Set the port and banner messages for users who log in through the web login page:

Note: If **Active session mode** is enabled, a different set of web pages is available.

- **Port** – Type the port number that will be used to log into the web. The default port is 8111.
- **Login page** – Click **Edit** to modify the message displayed for successfully logging in. Click **View** to see the web page.
- **Logout page** – Click **Edit** to modify the message displayed for successfully logging out. Click **View** to see the web page.
- **Redirect page** – Click **Edit** to modify the message displayed for successfully logging in after being redirected from a web request. Click **View** to see the web page.
- **Error page** – Click **Edit** to modify the message displayed if a page cannot be found. Click **View** to see the web page.

See [Accessing the web login and logout pages](#) for more information.

7 Set the timeout parameters for Passport users:

- **Authenticate inactive users every** – Set how long a user can be inactive before they must log into Passport again. (Not available for active session mode.)
- **Force authentication every** – Set the length of time between mandatory authentications. This setting applies even if a user is currently active.

8 [Optional] Click **Manage Passports** to view the current Passport-authenticated (cached) users, and to expire user Passport authentication for one or more users.

9 [Optional] To restrict proxy connections to a specific group of users that are created and managed on the firewall, click the **Users and User Groups** tab to create a user group.


See [Authenticating groups from an internal group source](#) for detailed information on users and user groups.

10 Save your changes.

Managing Passports

Use this window to view the current Passport-authenticated (cached) users. In this window, you have the option to revoke user Passport authentication for one or more users.

Figure 45 Manage Passports window



The following fields are displayed in the table:

- **Name** – This column displays the name(s) of all users who currently have cached authentication.
- **External Group** – [Conditional] This column displays the external group to which a user belongs. This applies only when a user authenticates with an authentication method that supports external groups.
- **Authenticator** – This column displays the type of authentication used by a user.
- **IP Address** – This column displays the source IP Address from which the authentication originated.
- **Issued** – This column displays the time at which a user was initially authenticated and obtained a Passport.
- **Last Used** – This column displays the time at which a user last accessed a service that required authentication.

You can immediately revoke Passport authentication for selected users by doing the following:

- To revoke the Passport authentication cache for all users listed in the table, click **Revoke All Passports**.
- To revoke the Passport for a single user or group of users, select the users you want to revoke by clicking the appropriate table row(s).
- To revoke multiple users, press and hold the **Ctrl** key as you select users. Then click **Revoke Passport(s)** to expire the selected users from the Passport.

When you revoke the Passport for users, those users will be required to re-authenticate before they can again access any Passport-authenticated service.

Note: Subsequent authentication requests by an expired user will be cached when they re-authenticate, allowing them to again use Passport authentication.

Accessing the web login and logout pages

When an HTTP connection is required to acquire a Passport, users are redirected from a web request to the authentication login page. Users can also access the authentication login page by directing their browser to:

https://firewall_address:8111/login.html

If a user wants to log out of the Passport cache manually (before their Passport authentication cache expires), they can point their browser to:

https://firewall_address:8111/logout.html

Note: If active session mode is enabled, this page is not available. Click **Stop** on the “Successful Login” page or close the browser window to log out of the Passport cache.

If a browser is configured for the proxy, you will need to configure that browser to NOT proxy requests going to the firewall on port 8111. The following steps provide an example of configuring an exception using Netscape.

- 1 Open Netscape and select **Edit > Preferences > Advanced > Proxies**.
- 2 Select **Manual Proxy Configuration**.
- 3 In the **No Proxy For** field, type the URL for the firewall (for example, *firewall_name.example.com*).
- 4 Click **OK** to save the information and exit.

Using an alternate Passport authentication method

If you need to use an authentication method other than the default for Passport authentication (for backup or test purposes, for example), you can enter a configured authenticator at the login prompt.

- The name of the authenticator can be abbreviated as long as it is unique. For example, *pass* is ambiguous because it matches *Password* and *Passport*, but *passw* would make it unique.
- The name of the authenticator is case-insensitive.

- 1 Configure the alternate authentication method in the Authenticators window.
- 2 On the Passport: General tab, select the alternate authentication method in the **Authenticators to establish Passport credentials** list.
- 3 When you attempt a connection and the login prompt appears, enter your user name followed by a colon and the name of the alternate authenticator:

login_name:authenticator

Setting up standard password authentication

Standard password authentication requires a user to enter the same password each time he or she logs in. The password is maintained in the user database on the firewall.

Standard password authentication is typically used for internal-to-external SOCKS5, Telnet, FTP, and HTTP connections, and for administrators logging into the firewall from the internal (trusted) network.

Since internal users are generally trusted, a weak authentication method like password may be all that is required. You may want to authenticate internal-to-external connections not so much for security reasons but to track usage of the system.

To set up standard password authentication: In the list in the upper pane of the Authenticators window, select **Password**.

Note: Password is a default authenticator. It cannot be deleted.

The Password: General tab appears in the lower pane.

Figure 46 Password: General tab

The screenshot shows the 'Password: General' configuration window. At the top, there are three fields: 'Name' (Password), 'Type' (Password), and 'Description' (Static password authenticator). Below these are two tabs: 'General' (selected) and 'Users and User Groups'. The 'General' tab is divided into two main sections: 'Login settings' and 'Password requirements'. The 'Login settings' section includes 'Login prompt' (Username:), 'Password prompt' (Password:), 'Expiration message' (Your password has expired.), and 'Password expiration' (90 days). The 'Password requirements' section includes 'Minimum Password Length' (8), a radio button for 'Allow simple passwords' (selected), and a radio button for 'Require complex passwords'. Below these are two spinners: 'Require 2 of the four character groups in every password' and 'Require at least 2 characters(s) per required group in every password'. A note states: 'The four character groups are: lowercase letters, uppercase letters, numbers, and special characters such as #.' An example valid password is shown as 'abcdefgh'.

To configure the Password: General tab:

1 Enter identifying information:

- **Name** – You cannot change the name of the Password authenticator.
- **Type** – This shows the type of authenticator. You cannot modify this field.
- **Description** – If desired, you can modify the description to help you more easily identify this authenticator.

2 Configure the login settings:

- **Login prompt** – Enter the text to appear asking for user identification.
- **Password prompt** – Enter the text to appear asking for a password.
- **Expiration message** – Enter the text to appear when a password has expired.
- **Password expiration** – Enter the number of days a password remains valid.

3 Configure the password requirements:

- **Minimum password length** – Enter the minimum number of characters a password must contain.
- **Allow simple passwords** – Select this option if you do not want to specify any other password requirements.
- **Require complex passwords** – Select this option to configure and enforce complex password requirements.
- **Require # of the four character groups in every password** – Specify the number of character groups required for passwords. The character groups are:

- lowercase
- uppercase
- numbers
- special characters (all printable characters that can be typed from the keyboard, such as ^ % \$ @ !, etc.)

If you specify **2**, passwords must use characters from any two of the four character groups.

- **Require at least # character(s) per required group in every password** – Specify the number of characters required from each character group.

If you specify **3** characters from each group, and two character groups are required, passwords must contain three characters from two different groups, such as **a13c7b**.

- 4 [Optional] To restrict proxy connections to a specific group of users that are created and managed on the firewall, click the **Users and User Groups** tab to create a user group.

See [Authenticating groups from an internal group source](#) for detailed information on users and user groups.

- 5 Save your changes.

Setting up LDAP authentication

Use LDAP (Lightweight Directory Access Protocol) to provide fixed password authentication for SOCKS5, Telnet, FTP, and HTTP sessions through the firewall. It can also be used to authenticate logins and SSH logins to the firewall.

There are four LDAP types you can use:

- **iPlanet** – Select this option if using an iPlanet LDAP server.
- **Active Directory** – Select this option if using an Active Directory LDAP server. You can set up an LDAP directory server containing users and passwords. Use any valid combination of LDAP attributes and values as an optional filter string to distinguish authorized firewall users.
- **Open LDAP** – Select this option if using an Open LDAP server. OpenLDAP Software is a free, open source implementation of LDAP developed by the OpenLDAP Project.
- **Custom LDAP** – Select this option to customize the Directory User Identifier and Directory Member Identifier, the attributes used in the LDAP server searches.

To set up LDAP authentication: In the upper pane of the Authenticators window, click **New** and select the appropriate LDAP type from the drop-down list.

The LDAP: General tab appears. For more information, see:

- [About the LDAP: General tab](#)
- [About the LDAP: Search tab](#)

About the LDAP: General tab

Use this tab to configure your firewall to work with an LDAP server.

Figure 47 LDAP: General tab

The screenshot shows the 'Authenticators: New Authenticator' window. At the top, there are input fields for 'Name', 'Type' (set to 'OpenLDAP'), and 'Description'. Below these are four tabs: 'General', 'Search', 'Users and User Groups', and 'Groups'. The 'General' tab is selected. It contains a table titled 'OpenLDAP servers' with columns 'Rank', 'Host', and 'Port'. To the right of the table are up and down arrows. Below the table are 'New', 'Modify', and 'Delete' buttons. To the right of the table is a 'Login options' section with two radio buttons: 'Connect to server(s) anonymously' and 'Connect to server(s) with username/password' (which is selected). Below these are three input fields for 'Username', 'Password', and 'Confirm password'. There are also two buttons: 'Server Timeouts/Retries' and 'Console and Telnet LDAP Logins'. At the bottom right of the dialog are 'Add', 'Close', and 'Help' buttons.

The left pane displays a list of any LDAP servers currently configured for the firewall, with the following columns:

- **Rank** – Which server the firewall will try first.
 - If the server returns any response, no further servers are queried.
 - If the server does not respond, the next server in the list is tried.
- **Host** – The host IP address for the LDAP server.
- **Port** – The port number the LDAP server should use. The default port is 389.

Click a column heading to sort the list by that column's content. Click again to reverse the sort order.

To configure the LDAP: General tab:

1 Enter identifying information:

- **Name** – Type a name to identify this authenticator. If you are modifying this authenticator, you cannot change the name.
- **Type** – This shows the type of authenticator. You cannot modify this field.
- **Description** – Type a description to help you more easily identify this authenticator.

2 Define and rank the LDAP servers.

Note: The maximum number of LDAP servers allowed at one time is four.

You can do the following:

- **Create a new server** – Click **New** and enter the IP address and port of the new LDAP server in the pop-up window. The default port is 389.
- **Modify an existing server** – Select the server and click **Modify**. Make the desired changes in the pop-up window.
- **Delete an existing server** – Select the server and click **Delete**.
- **Rank the servers** – Select a server and use the up and down arrows to change the rank.

3 Select how the firewall will connect to LDAP servers by selecting one of the following options:

- **Connect to server(s) anonymously** – Select this option if the LDAP server allows the firewall to connect and search subcontainers without providing login information.
- **Connect to server(s) with username/password** – Select this option if the LDAP server requires the firewall to submit the specified user name and password in order to connect and search subcontainers.
 - **Username** – Type the login name required by the LDAP server.

If you are configuring an Active Directory authenticator, specify a full distinguished name (DN) in this field. For example: *user@example.com*
 - **Password** – Type a password required by the LDAP server.
 - **Confirm password** – Type the password again.
- **Server Timeouts/Retries** – Click this to configure the login limit. Enter the login timeout in seconds.
- **Console and Telnet LDAP Logins** – Click this to specify what you want to appear as prompts during the login process. The defaults are *Username:* and *Password:.*

4 [Optional] Select a group source.

- To create internally managed groups that you can specifically allow in proxy connections, select **internal**, then click the **Users and User Groups** tab.

See [Authenticating groups from an internal group source](#) for detailed information on users and user groups.
- To add externally created groups that you can specifically allow in proxy connections, select **external**, then click the **Groups** tab.

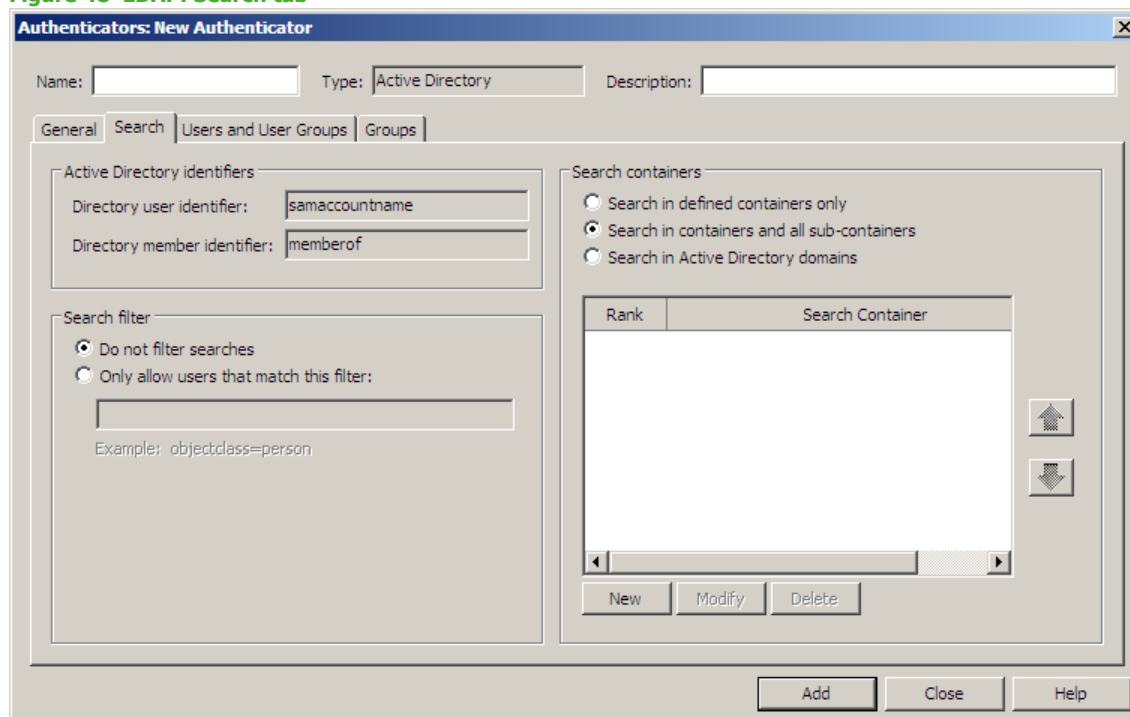
See [Authenticating groups from an external group source](#) for information on external authentication groups.

5 Click **Add** or **OK** and save your changes.

About the LDAP: Search tab

Use this tab to define the search parameters for LDAP authentication.

Figure 48 LDAP: Search tab



1 [Custom LDAP only] Enter the LDAP identifiers:

- **Directory user identifier** – Enter the attribute used in the LDAP database for user names. The LDAP server searches for a match on the user name assigned to this attribute.
- **Directory member identifier** – Enter the attribute used in the LDAP database for group names. The LDAP server searches for a match on the group name assigned to this attribute.

Note: In iPlanet, Active Directory, and OpenLDAP, these are default attributes that cannot be modified.

2 Define the search filter option:

- **Do not filter searches** – Select this option to disable filtering of the LDAP or Active Directory tree.
- **Only allow users that match this filter** – Select this option to filter users based on the profile filter displayed here.

3 Select which containers will be searched:

- **Search in all containers and sub-containers** – Select this option to search all listed containers and their subcontainers.
- **Search in defined containers only** – Select this option to limit searches to containers listed here.
- [Active Directory only] **Search in Active Directory domains** – Select this option to search only in Active Directory domains listed here. Each domain must be listed separately.

You can perform the following actions:

- **Create a new search container** – Click **New** and make entries in the pop-up window. Enter either a single container name or a concatenated container name.
Note: The search string format depends on the type of server selected. Microsoft Active Directory searches use a format similar to sales.example.com. Standard LDAP searches use a format similar to dc=sales,dc=example,dc=com.
- **Modify an existing search container** – Select the search container and click **Modify**. Make the desired changes in the pop-up window.
- **Delete an existing search container** – Select the search container and click **Delete**.
- **Change the search container's rank** – Select the search container and use the up and down arrows.

Setting up CAC authentication

Use the CAC authenticator to log into Sidewinder using a U.S. Department of Defense Common Access Card.

Figure 49 CAC: General tab

Authenticators: New Authenticator

Name: Type: Description:

General | Users and User Groups | Webservice Configuration

Expire one-time password after seconds

One-time password size:

Trusted Certificate Authorities (CA):

- ☒ DOD_CA-11
- ☒ DOD_CA-12
- ☒ DOD_CA-13
- ☒ DOD_CA-14
- ☒ DOD_CA-15
- ☒ DOD_CA-16
- ☒ DOD_CA-17
- ☒ DOD_CA-18
- ☒ DOD_CA-19
- ☒ DOD_CA-20
- ☒ DOD_CA-21
- ☒ DOD_CA-22

☐ Allow expired certificates

☐ Allow unknown certificate status

☐ Perform OCSRP checking

- ☐ Check full path through OCSRP
- ☐ Use nonce checking for OCSRP
- ☐ Do not check timestamps on OCSRP response
- ☒ Use Certificate OCSRP Responders
- ☒ Use CA OCSRP Responders

Default OCSRP Urls (optional):

Setting up Windows domain authentication

Use a Windows primary domain controller (PDC) or backup domain controller (BDC) to provide password authentication for login, SOCKS5, Telnet, FTP, HTTP, and SSH sessions to the firewall. You can also configure transparent browser authentication (NTLM) for browsers that support it.

Note: Be sure the domain controller does not allow blank or default logins that can be easily guessed by outsiders.

You can also use transparent browser authentication. For more information about configuring your organization's PDC or BDC to use transparent browser authentication on the firewall, see the related application note located at <https://support.forcepoint.com>.

Note: Transparent browser authentication is also known as NTLM or integrated Windows authentication.

To set up Windows authentication: In the upper pane of the Authenticators window, click **New** and select **Windows** from the drop-down list.

The Windows: General tab appears.

Figure 50 Windows: General tab

The screenshot shows the 'Authenticators: New Authenticator' window with the 'Windows' type selected. The 'General' tab is active, displaying a table for 'Windows domain controllers' with columns for Rank, IP address, Port, and Name. To the right of the table are 'Login options' including prompts for Username and Password, a failed authentication message field, and radio buttons for authentication methods: Domain (MSNT), Transparent (NTLM), and Both. At the bottom are buttons for 'Add', 'Close', and 'Help'.

Rank	IP address	Port	Name
------	------------	------	------

Buttons: New, Modify, Delete

Buttons: Add, Close, Help

The Windows domain controllers table lists the Windows domain controllers currently configured for the firewall.

- **Rank** – Which Windows domain controllers the firewall will try first.
 - If the server returns any response, no further servers are queried.
 - If the server does not respond, the next server in the list is tried.
- **IP Address** – The Windows domain controller's IP address.
- **Port** – The port used by the Windows domain controller. The default port is 139 and cannot be changed.
- **Name** – The name of the Windows domain controller.

Click a column heading to sort the list by that column's content. Click again to reverse the sort order.

To configure the Windows: General tab:

1 Enter identifying information:

- **Name** – Type a name to identify this authenticator. If you are modifying this authenticator, you cannot change the name.
- **Type** – This shows the type of authenticator. You cannot modify this field.
- **Description** – Type a description to help you more easily identify this authenticator.

2 Define and rank the Windows domain controllers.

Note: The maximum number of Windows domain controllers allowed at one time is four.

You can do the following:

- **Create a new controller** – Click **New** and make entries in the pop-up window:
 - **IP address** – Type the IP address used by the Windows domain controller.
 - **Windows domain controller name** – Type the name of the Windows domain controller. Type only the host or computer name, not the fully qualified name. (You can determine the name by going into the Network window on the Windows controller.)
 - **Port** – The port cannot be changed.
- **Modify an existing controller** – Select the controller and click **Modify**. Make the desired changes in the pop-up window.
- **Delete an existing controller** – Select the controller and click **Delete**.
- **Rank the controllers** – Select a Windows domain controller and use the up and down arrows to change the rank.

3 Modify the Login options:

- **Login prompt** – This is the login prompt that displays to users.
- **Password prompt** – This is the password prompt that displays to users.
- **Failed authentication message** – This is the message that displays if a user's authentication attempt fails.

4 Select prompted or transparent browser authentication:

- **Domain (MSNT)** – Select this option to prompt users for a user name and password. This is typically used for older browsers that do not support transparent authentication.

Security Alert: The user password is not encrypted in this method.

- **Transparent (NTLM)** – Select this option if you want transparent browser authentication. If a user has already been authenticated by the Windows domain, they are not prompted for a user name and password when using a rule that requires this authenticator.

If this option is selected and the user's browser does not support transparent authentication, the authentication will fail. No further rule matching is attempted.

- **Both** – Select this option to attempt both authentication methods. Transparent authentication is attempted first. If it is not supported, domain authentication is used.

5 [Optional] To restrict proxy connections to a specific group of users that are created and managed on the firewall, click the **Users and User Groups** tab to create a user group.

See [Authenticating groups from an internal group source](#) for detailed information on users and user groups.

Setting up RADIUS authentication

RADIUS is a standard protocol used to authenticate users before they are allowed access to your system.

- You can use RADIUS to provide authentication for SOCKS5, Telnet, FTP, and HTTP sessions through the Sidewinder.

Authentication

Configuring an authenticator

- You can use RADIUS to authenticate logins and SSH logins to the firewall.

Note: SafeWord RemoteAccess and SafeWord PremierAccess are RADIUS servers that have been certified for full interoperability with the firewall.

To set up RADIUS authentication: In the upper pane of the Authenticators window, click **New** and select **RADIUS** from the drop-down list.

The RADIUS: General tab appears.

Figure 51 RADIUS: General tab

The screenshot shows the 'Authenticators: New Authenticator' window. At the top, there are fields for 'Name:', 'Type:' (set to 'RADIUS'), and 'Description:'. Below these are three tabs: 'General', 'Users and User Groups', and 'Groups'. The 'General' tab is active. It contains a 'RADIUS servers' table with columns 'Rank', 'Host', 'Port', and 'Shared Secret'. To the right of the table are 'Login options' and 'Group source' sections. The 'Login options' section has fields for 'Login Prompt:' (Username:), 'Password Prompt:' (Password:), and 'Failed Authentication Message:' (Login incorrect). The 'Group source' section has radio buttons for 'internal' (selected) and 'external'. At the bottom of the window are 'Add', 'Close', and 'Help' buttons.

Rank	Host	Port	Shared Secret
------	------	------	---------------

The Radius Servers table lists the RADIUS servers currently configured for the Sidewinder. The columns indicate the following:

- Rank** – Which server the firewall will try first.
 - If the server returns any response, no further servers are queried.
 - If the server does not respond, the next server in the list is tried.
- IP address** – The host IP address for each server entry.
- Port Number** – The port number for each server entry. The default port is 1812.
- Shared Secret** – The text string or phrase that matches the shared secret of the listed RADIUS server.

Click a column heading to sort the list by that column's content. Click again to reverse the sort order.

To configure the RADIUS: General tab:

1 Enter identifying information:

- **Name** – Type a name to identify this authenticator. If you are modifying this authenticator, you cannot change the name.
- **Type** – This shows the type of authenticator. You cannot modify this field.
- **Description** – Type a description to help you more easily identify this authenticator.

2 Define and rank the RADIUS servers.

Note: The maximum number of RADIUS servers allowed at one time is four.

You can do the following:

- **Create a new server** – Click **New** and make entries in the pop-up window:
 - **IP address** – Type the host IP address for each server entry.
 - **Port Number** – Type the port number for each server entry. The default port is 1812.
 - **Shared Secret** – Type the text string or phrase that matches the shared secret of the listed RADIUS server.
- **Modify an existing server** – Select the server and click **Modify**. Make the desired changes in the pop-up window.
- **Delete an existing server** – Select the server and click **Delete**.
- **Rank the servers** – Select a server and use the up and down arrows to change the rank.

3 Modify the Login options:

- **Login prompt** – This is the login prompt that displays to users when they log in using RADIUS.
- **Password prompt** – This is the password prompt that displays to users when they log in using RADIUS.
- **Failed authentication message** – This is the message that displays if a user's authentication attempt fails.

4 [Optional] Select a group source.

- To create internally managed groups that you can specifically allow in proxy connections, select **internal**, then click the **Users and User Groups** tab.

See [Authenticating groups from an internal group source](#) for detailed information on users and user groups.

- To add externally created groups that you can specifically allow in proxy connections, select **external**, then click the **Groups** tab.

See [Authenticating groups from an external group source](#) for information on external authentication groups.

Setting up SafeWord authentication

The SafeWord RemoteAccess and SafeWord PremierAccess authentication servers interoperate with Sidewinder.

- To configure SafeWord PremierAccess authentication on Sidewinder, you must first install and configure the SafeWord PremierAccess Authentication Server.

With SafeWord PremierAccess, you can use fixed passwords or passcode authentication for Telnet and FTP sessions through the firewall, and for administrator login attempts directly to the firewall or through an SSH session. You can authenticate HTTP sessions using either fixed passwords or passcodes without the challenge/response option (not all tokens support this option).


Refer to the appropriate product documentation.

- To configure SafeWord RemoteAccess authentication, use the RADIUS authenticator. See [Setting up RADIUS authentication](#) for more information.

To set up SafeWord authentication: In the upper pane of the Authenticators window, click **New** and select **SafeWord** from the drop-down list.

The SafeWord: General tab appears.

Figure 52 SafeWord: General tab



The left pane displays a list of SafeWord servers currently configured for Sidewinder, with the following columns:

- Rank** – Which server Sidewinder will try first.
 - If the server returns any response, no further servers are queried.
 - If the server does not respond, the next server in the list is tried.
- IP Address** – The host IP address for each server entry.
- Port** – The port number for each server entry. The default port number for SafeWord PremierAccess is 5030.

Click a column heading to sort the list by that column's content. Click again to reverse the sort order.

To configure the SafeWord: General tab:

1 Enter identifying information:

- Name** – Type a name to identify this authenticator. If you are modifying this authenticator, you cannot change the name.

- **Type** – This shows the type of authenticator. You cannot modify this field.
- **Description** – Type a description to help you more easily identify this authenticator.

2 Define and rank the SafeWord servers.

Note: The maximum number of SafeWord servers allowed at one time is four.

You can do the following:

- **Create a new SafeWord server entry** – Click **New** and enter the IP address and port in the pop-up window. The default port number for SafeWord PremierAccess is 5030.
- **Modify an existing server entry** – Select the server and click **Modify**. Make the desired changes, then click **OK**.
- **Delete an existing entry** – Select the entry and click **Delete**.
- **Rank the servers** – Select a server and use the up and down arrows to change the rank.

3 [Optional] Select a group source.

- To create internally managed groups that you can specifically allow in proxy connections, select **internal**, then click the **Users and User Groups** tab.
See [Authenticating groups from an internal group source](#) for detailed information on users and user groups.
- To add externally created groups that you can specifically allow in proxy connections, select **external**, then click the **Groups** tab.
See [Authenticating groups from an external group source](#) for information on external authentication groups.

4 Click **Add** or **OK** and save your changes.

Telnet and FTP considerations

There are some special considerations that users should be made aware of regarding Telnet and FTP authenticated connections through Sidewinder.

- **Changing user passwords and PINs for authentication methods**

Sidewinder supports changing user passwords and PINs only under the Telnet proxy. For example, users can change their SafeWord PremierAccess PIN via the Telnet proxy. (Refer to the documentation for your authentication method for information on the commands used to change passwords and PINs.) Passwords and PINs cannot be changed using the FTP, HTTP, or SOCKS5 proxy. The user must either initiate a Telnet proxy session or they can contact their system administrator.

- **Switching authentication methods during a login session**

The firewall allows you to use multiple authentication methods for a given service (for example, users might use SafeWord PremierAccess or Password for Telnet authentication). When logging in, if a user specifies the incorrect authentication method and authenticator, they can change to another authentication method by typing `:authenticator` after the user name.

- **Non-authenticated nontransparent FTP proxy prompts for authentication**

Administrators should instruct end users that they will be prompted to supply a user name, authentication method, and destination, even if the associated allow rule does not require authentication. This is because the non-transparent FTP proxy needs the login and destination information in order to determine which rule will allow the connection.

When end users attempt to connect to the FTP server, the firewall sends them the following prompt:

```
220-Firewall ftp proxy. You must login to the proxy first.  
220 Use proxy-user:auth-method@destination.  
Name (si_ipaddr:proxy-user):
```

Instruct users to respond to the `Name (si_ipaddr:username) :` prompt by entering the @ sign followed by the FTP server's IP address, as shown in this example:

```
Name (si_ipaddr:proxy-user) :@172.1.1.25
```

Users who incorrectly put a user name before the prompt are still allowed access to the FTP server through the non-transparent FTP rule that does not require authentication. The firewall handles entries containing user names that do not match any existing FTP rule and entries without a user name in the same manner.

Setting up users to change their own passwords

The firewall changepw server allows users to use a web browser to change their Sidewinder, SafeWord PremierAccess, or LDAP login password.

To allow this process, you must do the following:

- Create a change password rule that allows users to change their passwords.
- Inform users how they can change their own passwords using a web browser.

Create a change password rule

To create a change password rule: Select **Policy > Rules** and select the appropriate settings from the table below.

Table 15 Proxy rule settings to allow users to change their login passwords

Criteria	Setting
Action:	Allow
Service:	changepw
Source Burb:	Desired burb (for example, internal)
Destination Burb:	Desired burb (for example, internal)
Source Endpoint:	Site dependent
Destination Endpoint:	localhost (a default host object)
Redirect:	Firewall (IP)

How users can change their own password

Using standard password authentication, you can authenticate trusted and Internet users who request SOCKS5, FTP, HTTP, and Telnet access via proxies. As an administrator, you should inform those users how they can change their own password from their terminal or workstation by using a web browser. However, there are some restrictions:

- Users can change their own password only if using standard password, SafeWord PremierAccess, or LDAP authentication.
- To allow users to change their login passwords, you must first create a rule for the firewall to allow this.

1 Start a web browser.

2 Configure your browser *not* to proxy requests going to the firewall on port 1999. For example, if you are using a Netscape browser do the following:

- Open Netscape and select **Edit > Preferences > Advanced > Proxies**.
- Select **Manual Proxy Configuration**.
- In the **No Proxy For** field, type the URL for the firewall (for example, *myfirewall.example.com*).
- Click **OK** to save the information and exit.

3 Open an HTTP connection to Sidewinder. For example:

`http://myfirewall.example.com:1999/`

A pre-defined HTML change password form appears.

4 Enter your user name.

5 Enter your current password. This is your current password for establishing network connections.

6 Enter your new password. This will be your new password for establishing network connections.

7 Re-enter the new password. This confirms the spelling of the new password.

8 Select one of the following password types:

- If you are changing a Sidewinder login password, select **Password**.

Authentication

Setting up users to change their own passwords

- If you are changing a SafeWord PremierAccess login password, select **SafeWord**.
- If you are changing an LDAP password, select **LDAP**.

9 Click **Send Request**.

This sends the change password request to the firewall. You will be notified if the request failed or if it is accepted. If the request is accepted, the password database is updated and the new password must be used for all future connections.

Authenticating groups from an external group source

A group is a logical grouping of one or more users, identified by a single name. You can restrict proxy connections to specific groups created and managed on an external authentication server.

You can authenticate groups from external servers using LDAP, RADIUS, and Safeword authenticators:

- A group is created on an external authentication server. In the Admin Console, you add the matching group name on the Groups tab of an LDAP, RADIUS, or Safeword authenticator.
- When you select this authenticator in the Rules window, you can also select one or more groups that were added to the Groups tab. Proxy connections are restricted to users in the matching group(s) on the external authentication server.
- An external group added to an authenticator is not available globally. An external group is unique to the authenticator it is added to. If you want to use the same group for another authenticator, it must also be added to the Groups tab of that authenticator.

To add or modify external group names to an authenticator:

1 Select **Policy > Rule Elements > Authenticators**.

2 Open a new or existing LDAP, RADIUS, or Safeword authenticator window:

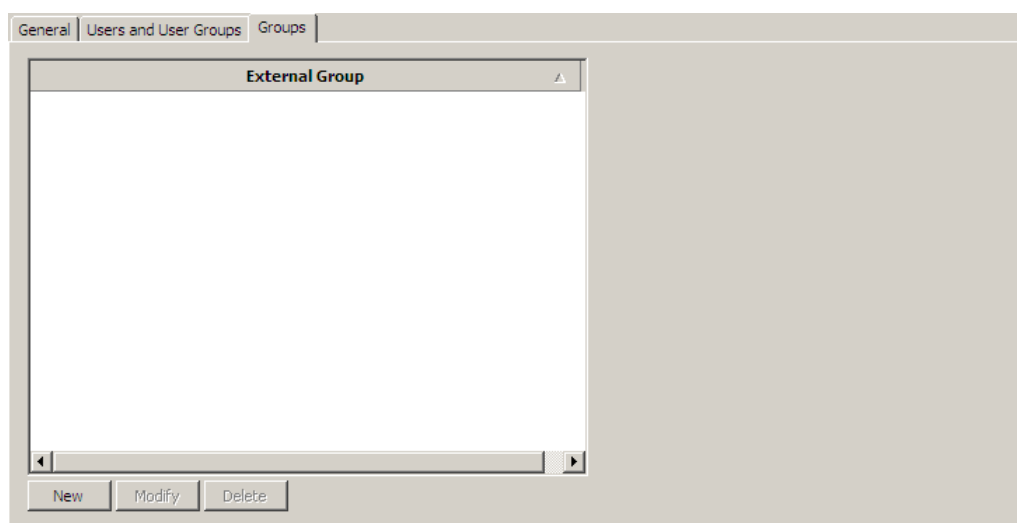
- Click **New** and select the appropriate authenticator.
- In the upper pane, select an existing authenticator.

3 On the General tab, select the **external** group source.

Note: If the authenticator is being used in a rule, you cannot change the group source.

4 Click the **Groups** tab.

Figure 53 Authenticator: Groups tab



You can perform the following actions:

- **Add a new external group** – Click **New**, then type the name of a group that matches the group name residing on an external LDAP, RADIUS, or SafeWord authentication server. If you enter multiple groups, put each group on a separate line.
- **Modify an existing external group** – Select an external group name and click **Modify**, then make the appropriate changes in the Modify Group window.
- **Delete an existing external group** – Select an external group name and click **Delete**.

Authentication

Authenticating groups from an external group source

RADIUS group options [RADIUS authenticators only]:

Enter the attributes defined in the dictionary files on the RADIUS server. The firewall looks for these attributes in the RADIUS server's response.

- **Group type** – Enter an attribute type. The default is 26, which is a vendor-specific attribute.
 - **Vendor ID** – If the group type is 26, enter a vendor ID from the RADIUS server's dictionary files.
 - **Vendor type** – If the group type is 26, enter a vendor type from the RADIUS server's dictionary files.
- **Group delimiters** – If the RADIUS server sends attributes in a single string, enter the character(s) that separate the groups in the string. Multiple characters can be entered in this field consecutively, with no space or separators.

Save your changes.

Authenticating groups from an internal group source

A user is a person who uses the networking services provided by the firewall. A user group is a logical grouping of one or more users, identified by a single name. You can restrict proxy connections to specific groups created and managed on the firewall.

You can authenticate user groups using any firewall authenticator.

- You create users and add them to user groups on the Authenticators windows.
- When you select an authenticator on the Rules window, you can also select one or more of these groups. Proxy connections are restricted to users in the selected group(s).
- Users and groups created on the Users and User Groups tab of an authenticator are available to all authenticators.

Note: When using an internal group source, users created and maintained on the firewall for LDAP, RADIUS, Windows, or Safeword authenticators must also be maintained on their external servers.

- You create administrators in the Administrator Accounts window. All administrator accounts that are created appear in the Users and User Groups tab.
 - On a newly installed firewall, the only user to appear in the Users and User Groups tab is the administrator created during installation.
 - If you delete an administrator in the Administrator Accounts window, that administrator is also deleted from the Users and User Groups tab.
 - If you delete an administrator in the Users and User Groups tab, that administrator is also deleted from the Administrator Accounts window.

To add or modify users and user groups:

1 Select **Policy > Rule Elements > Authenticators**.

2 Open a new or existing authenticator window:

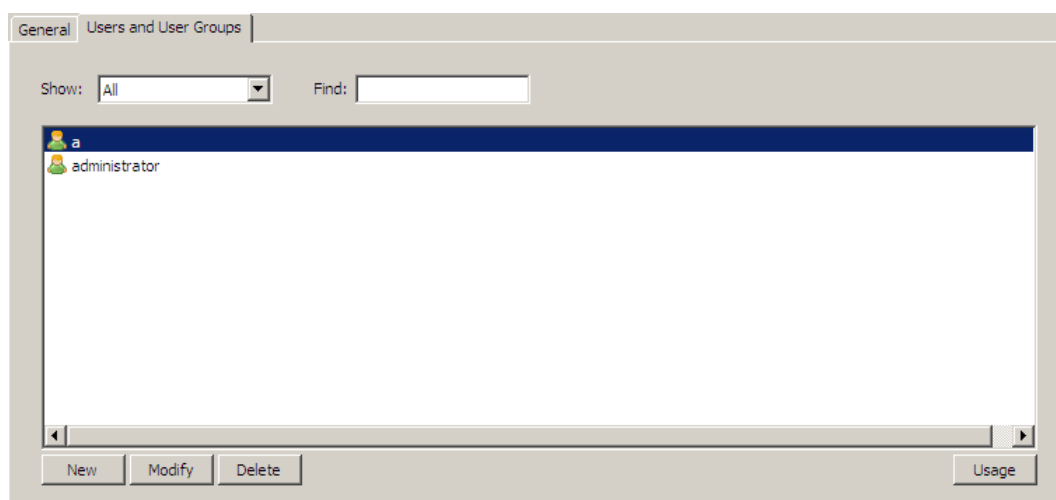
- Click **New** and select the appropriate authenticator.
- In the upper pane, select an existing authenticator.

3 On the General tab, select the **internal** group source (not necessary for Passport, Password, and Windows authenticators).

Note: If the authenticator is being used in a rule, you cannot change the group source.

4 Click the **Users and User Groups** tab.

Figure 54 Users and User Groups tab



About the Users and User Groups tab

Use this tab to create and manage users and user groups.

To manage the list of users and user groups:

- Display only users (**Users**), only groups (**Groups**) or both users and groups (**All**) by using the **Show** drop-down list.
- Filter the list by typing letters in the **Find** field. Only users or user groups that contain the corresponding string of letters appear in the list. For example, if you type **br** in the **Find** field, only users and groups whose name contains “br” will appear in the list. The **Find** field is case sensitive.
- To see which areas of the firewall are using a selected user or group, select the entry in the list and click **Usage**.

You can perform the following tasks in this tab:

- [Create a new user](#)
- [Create a new group](#)
- [Modify an existing user or user group](#)
- [Block consecutive failed authentication attempts](#)
- [Delete an existing user or user group](#)

Create a new user

- 1 In the lower pane, click **New**. The Create New User/Group window appears.
- 2 Select **New User**.
- 3 Select a template: Select **Use empty template** or select a user from the list and select **Copy from an existing user**.
- 4 Click **OK**. The User Objects window appears.
- 5 Enter the appropriate information for the new user.
- 6 Click **OK**, then save your changes.

Create a new group

From the Users and User Groups tab:

- 1 In the lower pane, click **New**. The Create New User/Group window appears.
- 2 Select **New Group**.
- 3 Click **OK**. The Group Objects window appears.
- 4 Enter the appropriate information for the new user group.
- 5 Click **OK**, then save your changes.

For more information, see:

- [About the Create New User/Group window](#)
- [About the Group Objects: Group Information tab](#)
- [About the User Objects: User Information tab](#)

Authentication

Authenticating groups from an internal group source

Modify an existing user or user group

From the Users and User Groups tab:

- 1 In the lower pane, select a user or user group from the list. If necessary, use the **Show** drop-down list or the **Find** field to narrow the list choices.
- 2 Click **Modify**. The User Objects or Group Objects window appears.
- 3 Make the necessary changes.
- 4 Click **OK**, then save your changes.

For more information, see:

- [About the Group Objects: Group Information tab](#)
- [About the User Objects: User Information tab](#)

Block consecutive failed authentication attempts

From the Authenticators window:

- 1 From the toolbar, click **Manage Authentication Failures**. The Authentication Failure Lockout Properties window appears.
- 2 Select **Enable** to enable the lockout feature.
- 3 In the Lockout Threshold field, type the number of failed login attempts allowed for a single user before that user is locked out of the firewall.
- 4 [Conditional] To clear the lock for a user, select the user and click **Clear**. Click **Clear All** to clear all users in the list.

For more information, see [About the Authentication Failure Lockout Properties window](#).

Delete an existing user or user group

From the Users and User Groups tab:

- 1 Select a user or user group from the list. If necessary, use the **Show** drop-down list or the **Find** field to narrow the list choices.
- 2 Click **Delete**. Save your changes.

Note: If you select an administrator to delete, that administrator will also be deleted from the Administrator Accounts window.

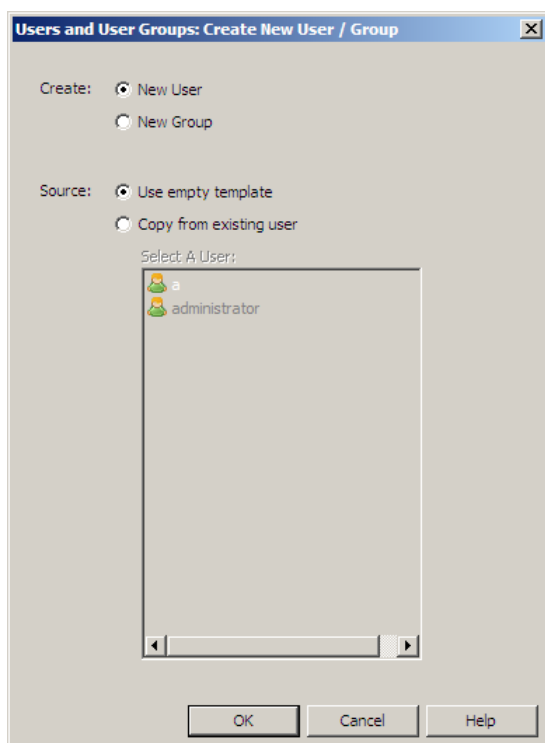
Authentication

Authenticating groups from an internal group source

About the Create New User/Group window

Use this window to select whether you want to create a user or a user group. The selections you make in this window will open the appropriate window to enter information.

Figure 55 Create New User/Group window



1 Select one of the following options in the **Create** field:

- **New User** – Select this option to create a new user.
- **New Group** – Select this option to create a new user group.

2 [New User only] Select a source:

- If you want to enter all new information, select **Use empty template**.
- If you want to create a new user account using the information contained in an existing user account, select the **Copy from existing user** option and then select the user account that you want to copy.

This option will copy the following information fields from the existing user's account:

- Organization
- User Fields 1–4
- Description
- Employee ID
- Group Membership

You will still need to enter information for the **Username** and **Password**, as these fields contain information specific to each individual user.

3 Click **OK**.

- If you are creating a new user group, the Group Objects window appears. See [About the Group Objects: Group Information tab](#).
- If you are creating a new user, the User Objects window appears. See [About the User Objects: User Information tab](#).


Authentication

Authenticating groups from an internal group source

About the Group Objects: Group Information tab

Use this tab to create or modify user groups.

Figure 56 Group Objects: Group Information tab

The screenshot shows a dialog box titled "Users and User Groups: Group Objects". It has two tabs: "Group Information" (selected) and "User Group Membership". The "Group Information" tab contains two text input fields: "Group Name:" and "Comments:". At the bottom of the dialog are three buttons: "OK", "Cancel", and "Help".

- **Group Name** – Type a name for this group.
 - Valid values include alphanumeric characters, periods (.), dashes(-), underscores (_), and spaces ().
 - The first and last character of the name must be alphanumeric.
 - The name cannot exceed 100 characters.

Note: You cannot edit the name of an existing group from this window. To change a group name, delete the group and add it back using the new name. Be sure to add the group back to any rules that used the deleted group name.

- [Optional] **Comments** – Type additional information about the user group.

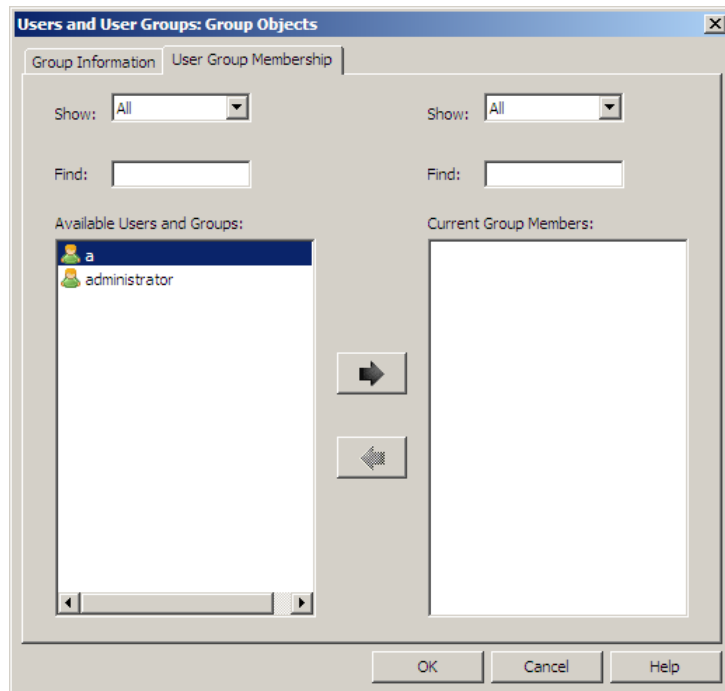
Use the **User Group Membership** tab to add or remove users or groups as members of this group.

When you are done creating or modifying this user group, click **OK** and save your changes.

About the Group Objects: User Group Membership tab

Use this tab to add or remove users or groups as members of this group. A group within a group is called a *nested group*.

Figure 57 Group Objects: User Group Membership tab



To filter the list:

- Use the **Show** drop-down list to display only users (Users), only groups (Groups) or both users and groups (All).
- Filter the list by typing letters in the **Find** field. Only users or user groups that contain the corresponding string of letters appear in the list. For example, if you type br in the Find field, only users and groups whose name contains "br" will appear in the list. The Find field is case sensitive.

To add or remove members of the selected group:

- To add a user or group as a member of this group, select an entry in the **Available Users and Groups** list, and then click the > arrow button.

Select multiple consecutive entries by pressing the **Shift** key as you select the entries. To select multiple non-consecutive entries, press the **Ctrl** key as you select the desired entries.

- To remove a user or group from this group, select the entry in the **Current Group Members** list, and then click the < arrow button.

Use the **Group Information** tab to enter information about the current group.

When you are done creating or modifying this user group, click **OK** and save your changes.

Authentication

Authenticating groups from an internal group source

About the User Objects: User Information tab

Use this tab to enter descriptive information about a user.

Figure 58 User Objects: User Information tab

The screenshot shows a dialog box titled "Users and User Groups: User Objects" with a close button (X) in the top right corner. It has two tabs: "User Information" (selected) and "User Password". The "User Information" tab contains several text input fields: "Username:", "Description:", "Employee ID:", "Organization:", "User Field 1:", "User Field 2:", "User Field 3:", and "User Field 4:". At the bottom of the dialog are three buttons: "OK", "Cancel", and "Help".

You can perform the following actions in this window:

- **Username** – Type the name the user will enter when he or she requests a connection that requires authentication. This entry can consist of up to 16 alphanumeric characters (upper or lower case). Apostrophes are *not* allowed (for example, *O'Hare*).
- [Optional] **Description** – Type any information about the user that may be helpful.
- [Optional] **Employee ID** – Type an employee ID number, if applicable.
- [Optional] **Organization** – Type the organization that the user is associated with, if applicable.
- [Optional] **User Field 1–4** – Enter any additional information that your organization requires. For example, if you will be generating chargeback reports for authenticated FTP, Telnet, or HTTP connections, you might enter account numbers in these fields.

You cannot modify the field names.

Use the **User Password** tab to enter password information for a user.


When you are done creating or modifying this user, click **OK** and save your changes.

About the User Objects: User Password tab

Use this tab to enter password information for a user.

Note: This password is used only for the Password authenticator. For other authenticators, the password is determined on the external server.

Figure 59 User Objects: User Password tab



- **Password** – Create the user's password using one of these methods:
 - **Manually create a password** – If you want to manually create a password, type a password in the **Password** field, then retype the password in the **Confirm Password** field. The password must not exceed 64 characters.
 - **Automatically generate a password** – If you want the firewall to automatically create a password, click **Generate Password**. Be sure to note the password that appears in the Generated Password window before clicking **OK**. Once you click **OK**, the password will no longer be visible, but can be changed at any time.
- **Expire Password** – Click this if you want the user's password to expire so they are required to change it. The **Expire Password** button changes to a **Reinstate Password** button.
- **Reinstate Password** – Click this if you need to re-instate a user's expired password. The **Reinstate Password** button changes to an **Expire Password** button.
- **Discard Password Info** – Click this to delete a user's password account from the database. For example, this can be used if you are changing a user's authentication method from password to SafeWord and need to remove the previous password information.

Use the **User Information** tab to enter descriptive information about a user.

When you are done creating or modifying this user, click **OK** and save your changes.

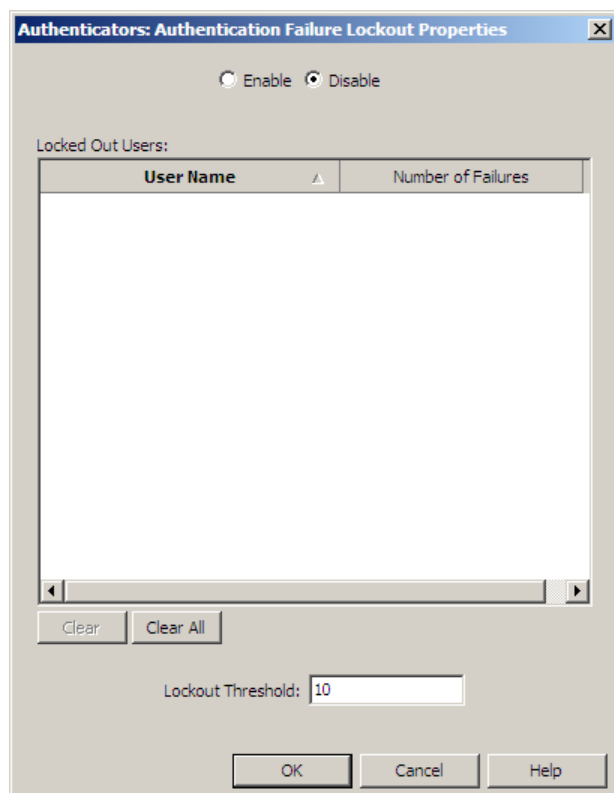
About the Authentication Failure Lockout Properties window

Use this window to configure the authentication failure lockout feature on your firewall. This feature allows you to configure the firewall to block access to a user if the number of consecutive failed authentication attempts reaches a configured number. This protects unauthorized users from multiple attempts at guessing a user's password.

Authentication

Authenticating groups from an internal group source

Figure 60 Authentication Failure Lockout Properties window



You can perform the following actions:

Note: If all administrators become locked out of the firewall, see [Manually clearing an authentication failure lockout](#).

- **Enable or disable the lockout feature** – To enable this feature, select the **Enable** radio button. To disable this feature, select the **Disable** radio button.

When this feature is enabled, any time a user account reaches the specified authentication attempt threshold without a successful authentication, that user will be locked out until the lock is cleared by an administrator. The lock can also be cleared if the locked out administrator logs in at the firewall using the correct login information.

- **View locked out users** – The **Locked Out Users** area lists any users who are currently locked out of the firewall due to exceeded authentication failures. It will also display the number of failed login attempts for each user.
- **Configure the lockout threshold** – Use the **Lockout Threshold** field to specify the number of failed login attempts that can occur for a single user account before that user is locked out of the firewall.


Note: When a user is locked out, their authentication method will become invalid. They will **not** be notified that they are locked out.

- **Clear user locks** – To clear the lock for a user, select the user and click **Clear**. Click **Clear All** to clear all users in the list.

Authentication

Authenticating groups from an internal group source

Figure 61 Locked-out user displayed in Authentication Failure Lockout Properties window



6 Content Inspection

Contents

[About content inspection](#)
[Configuring IPS inspection](#)
[Configuring virus scanning services](#)
[About TrustedSource](#)
[Updating the Geo-Location database](#)
[Configuring SmartFilter for Sidewinder](#)

About content inspection

Forcepoint Sidewinder policy is applied primarily by rules, which are made up of many elements. The table below shows the progression of a rule's creation using these elements and their corresponding chapters in this guide.

You are here in the Policy section	Use this chapter to...
Chapter 3, Policy Configuration Overview	understand the policy creation process.
Chapter 4, Network Objects and Time Periods	create or modify any network objects or time periods that will be used by rules.
Chapter 5, Authentication	create or modify authenticators that will be used by rules.
Chapter 6, Content Inspection	configure content inspection methods that will be used by rules.
Chapter 7, Services	create or modify services or service groups that will be used by rules.
Chapter 8, Application Defenses	create or modify Application Defenses that will be used by rules.
Chapter 9, Rules	create rules using the elements you created in the previous chapters in the policy section.

The Sidewinder content inspection methods provide additional security features by examining the content of a connection after it has matched a rule.

The following content inspection methods are available:

- **Intrusion Prevention System (IPS)** – IPS is a signature-based inspection tool that identifies attacks before they pass through the Sidewinder. See [Configuring IPS inspection on page 117](#).
- **Virus scanning** – The anti-virus service is a licensed add-on module that uses a firewall-hosted virus scanner that allows you to configure rule-based MIME, virus, and spyware scanning. See [Configuring virus scanning services on page 134](#).
- **TrustedSource** – ® TrustedSource™ is a reputation service that assigns a reputation score to an IP address based on the behavior attributes of the traffic it generates. A reputation score is like a credit score that indicates the trustworthiness of an IP address. See [About TrustedSource on page 138](#).
- **Geo-Location** – Geo-Location identifies the country of origin of an IP address. You can create a Geo-Location network object and apply it to a rule to allow or deny a connection based on the source or destination country. See [Updating the Geo-Location database on page 146](#).

Once a content inspection method is configured, it becomes available for selection on rules or in some cases Application Defenses.

Configuring IPS inspection

The Sidewinder Intrusion Prevention System (IPS) feature is a signature-based inspection tool that identifies attacks before they pass through the firewall. IPS plays an important role in protecting hosts and services that have known vulnerabilities and exploits, yet are required components of your organization.

Before the firewall will compare traffic to its IPS signatures, you must ensure the following conditions have been met:

- The IPS and IPS Signature features must be licensed. To verify that these features are licensed, select **Maintenance > License**, and click the **Firewall** tab. If you are not licensed for IPS and IPS Signature, contact your sales representative.
- The signature files are current. Select **Policy > IPS** and then click the **Signature Updates** tab. We strongly recommend that you enable automated signature download and install.
- You must create the appropriate signature groups and response mappings. Select **Policy > IPS** and click the **Signature Groups** and **Response Mapping** tabs.
- The rules governing the traffic you want inspected must have the appropriate signature categories and response mappings selected.

The following sections explain how Sidewinder IPS inspection is designed, how it interacts with other Sidewinder attack protection tools, how it is used in rules, and how to configure its basic components and signature file download schedule.

- [Understanding signature-based IPS](#)
- [Adding IPS inspection to rules](#)
- [About signature file updates](#)
- [Using IPS with other Sidewinder attack protection tools](#)
- [Configuring a response mapping](#)
- [Configuring a signature group](#)
- [Managing signatures](#)
- [Configuring IPS signature file updates](#)

Understanding signature-based IPS

The Sidewinder IPS inspection uses signatures to detect and prevent known network-based intrusion attacks, such as hacker-generated exploits. How the firewall responds to an attack is configurable; options range from allowing but auditing the attack to blackholing all traffic coming from the attacker.

IPS inspection is controlled on a per-rule basis. Each proxy, filter, or server rule that uses IPS inspection is assigned a *signature group* and a *response mapping*. The signature group is used to limit scanning to relevant signatures. The response mapping specifies the action to take when a packet or session is identified as an attack.

The foundation of IPS inspection is its *signatures*. The signatures are the data for recognizing attacks. Each signature has a *category* attribute, a *threat level* attribute, and a *class type* attribute.

The signature category is classified by the network service targeted for attack, and consists of a main category and a subcategory. One or more categories can be added to a signature group. For example, to create a signature group to add to an inbound rule for an Oracle server, create a group named *Oracle* that includes the categories *DB:Oracle*, *Component:Encoder*, and *Component:Shellcode*. The firewall also provides default signature groups based on common attack targets, such as the Database Servers group and the Internal Desktops group.

Within each category and response mapping, the signatures have a threat level attribute: *IPS* or *IDS*. This threat level indicates a relationship between confidence level and severity. Signatures classified as IPS detect attacks that are considered dangerous. Signatures classified as IDS detect attacks that are either considered minor, such as probe or discovery activity, or they are suspected attacks, meaning the signature may be likely to incorrectly identify legitimate traffic as an attack. The default signature groups and response mappings include both the IPS and IDS threat levels.


Signatures classified as *Policy* identify network traffic that you want to control based on your organization's security policy, such as instant messaging or P2P communication. Policy signatures are added individually to a signature group—they are not included in the default signature categories since they are specific to an organization.

The class type identifies the attack's intended purpose, such as *Root Level Exploit* or *Discovery*. Based on class type and threat level, you configure the response the firewall will take when an attack matches a signature. Options are to allow the packet or session, deny it, drop it, or blackhole it. These options generally include an IPS audit that records the action. In general, the response should correspond to the severity of the attack. Categories labeled IDS may generate some false positives or may be probing or discovery attacks. Therefore, attacks of this threat level should generally never be blackholed.

For example, to create a response mapping that protects against root level exploits against an Oracle server, create a mapping named *Oracle* and set Root Level Exploit type IDS to *Allow* and Root Level Exploit type IPS to *Blackhole* for 10,000 seconds. The process, illustrated in Figure 62, is as follows: An Oracle attack matches an Oracle proxy rule. That rule is configured for signature-based IPS inspection. The packet is compared to the signatures in the signature group and a match is found. The firewall then checks the rule's response mapping for instructions on responding to the attack. If the identified attack matches a signature with a threat level of IDS, the connection is allowed through but generates an IPS audit event. If the identified attack matches a signature with a threat level of IPS, the connection is blackholed for 10,000 seconds, so all traffic from the source's IP address is blackholed for that length of time.

Figure 62 shows the order in which IPS processing is performed.

Figure 62 IPS process flow



Adding IPS inspection to rules

As explained in the previous section, IPS inspection is controlled on a per-rule basis. Inspecting all traffic using IPS signatures can greatly reduce your firewall's performance. Enabling IPS inspection only when needed allows you to focus your firewall's resources on traffic that is most likely to contain attacks, such as HTTP traffic. Use signature groups, which limit scanning to relevant areas of the signature file database, to improve inspection efficiency.

Note: If your policy does call for extensive IPS inspection, you may be able to install a hardware accelerator to improve performance. This option is not available on low-end models. Contact your sales representative for more information.

When planning your security policy, determine what traffic and systems are likely to be targets for network-based attacks. IPS is most commonly used to inspect inbound connections, since attacks typically come from external, untrusted sources. If an internal server, such as a web server on your DMZ, were to be compromised, scanning its outbound connections is useful for containing damage and preventing attacks from spreading to other systems. Enable IPS on the rules governing likely targets. Traffic that does not have IPS inspection enabled will not be inspected for network-based attacks.

Tip: If you want to blackhole an attack that is identified by the signature-based IPS when it first occurs, set that action in the response mapping. If you want to blackhole an attack only if it occurs multiple times, set that action in the IPS Attack Responses (**Monitor > IPS Attack Responses**).

The following figure is an example of a rule with IPS inspection enabled. When HTTP traffic destined for the vulnerable_web_server reaches the firewall, the firewall checks that traffic against signatures in the "Web Server Attacks" signature group. When the traffic's pattern matches an attack, the firewall checks the "Exploit Protection" response mapping to see how it should respond to that attack's associated class type.

For more information on enabling IPS inspection in rules, see [Chapter 9, Rules](#).

Figure 63 A rule with IPS enabled

The screenshot displays the configuration for a firewall rule. The **General** tab is active, showing the Action set to **Allow**, Service set to **http (HTTP Proxy)**, and Audit set to **Standard (recommended)**. The **Effective Times** section shows the Time period set to **<Any>**. The **Source** section shows the Burb set to **external**, Endpoint set to **<Any>**, and NAT set to **localhost (Host)**. The **Destination** section shows the Burb set to **external**, Endpoint set to **External Vulnerable Web Server (IP)**, and Redirect set to **Vulnerable Web Server (IP)**. The **TrustedSource** section shows the **Enable TrustedSource** checkbox unchecked. The **Inspection** section shows the **Application Defense** set to **<Default Group>**, the **IPS Signature group** set to **Web Servers**, and the **Response mapping** set to **Exploit Protection**. Two arrows point to the **IPS Signature group** and **Response mapping** fields with the following text:

Searches signatures related to web server attacks.

Checks this response mapping to see what it should do with the connection.

About signature file updates

Since new attacks are being identified all the time, it is important to update the signatures frequently. When new signatures are added to the firewall, they will go into effect based on how the existing signatures categories and response mappings are configured. Therefore, if a new signature comes in and, based on its category and class type, is associated with a signature group that is assigned to a rule, that signature will go into effect immediately. Any attack matching that signature will be handled based on the response mapping for the signature's class type.

- Signatures with any risk of false positive are always given a threat level of IDS. Therefore, do not deny, drop, or blackhole traffic for class types with a threat level of IDS.
- Policy signatures are not included by default in any signature category group and general class types are not applied to them. Therefore, new policy signatures must be specifically added to category groups in order to use them.

Using IPS with other Sidewinder attack protection tools

There are several different approaches to protecting your internal network. One approach is to prohibit any traffic from entering your network. While this solution is secure, it is also impractical. Another approach is to attempt to scan all incoming traffic for known attacks, viruses, etc., but this can slow down the firewall, and therefore your network connectivity.

The best solution is first use tools to minimize your network's attack surface, and then use scanning to protect services that must be allowed. You can reduce your network's attack surface by creating the minimum number of rules necessary to allow essential inbound traffic and limiting the source and destination endpoints to hosts or address ranges. In addition, Application Defenses can be used to further refine what traffic is allowed into your network by prohibiting unnecessary commands, header, protocol versions, and other parameters. Once your policy is sufficiently restrictive, use IPS and other signature-based services such as anti-virus to inspect traffic destined for vulnerable yet essential services.

For example, an administrator is running a web server that requires allowing inbound HTTP traffic. The administrator knows that the Content Length header and the Content Location header are often used in attacks. The Content Location header is not required by the web server, and therefore does not need to be allowed into the network. The administrator uses the HTTP Application Defense to deny that header. The Content Length header is required, so the administrator allows it but adds IPS inspection to the rule allowing that traffic to make sure known attacks using that header are blocked.

While a small attack surface and inspection tools are a strong defense, you should still use IPS Attack Responses to monitor attack activity. Even attacks that are not allowed through the firewall are noteworthy as they may be an attempt from a hacker who will later try a more sophisticated attack. IPS Attack Responses can send out alerts when your network is under attack. These alerts will notify you of situations that may require a configuration change to increase the security of your network or investigation into the reason for the attack. For information on monitoring attack audits, see [Chapter 13, IPS Attack and System Event Responses](#).

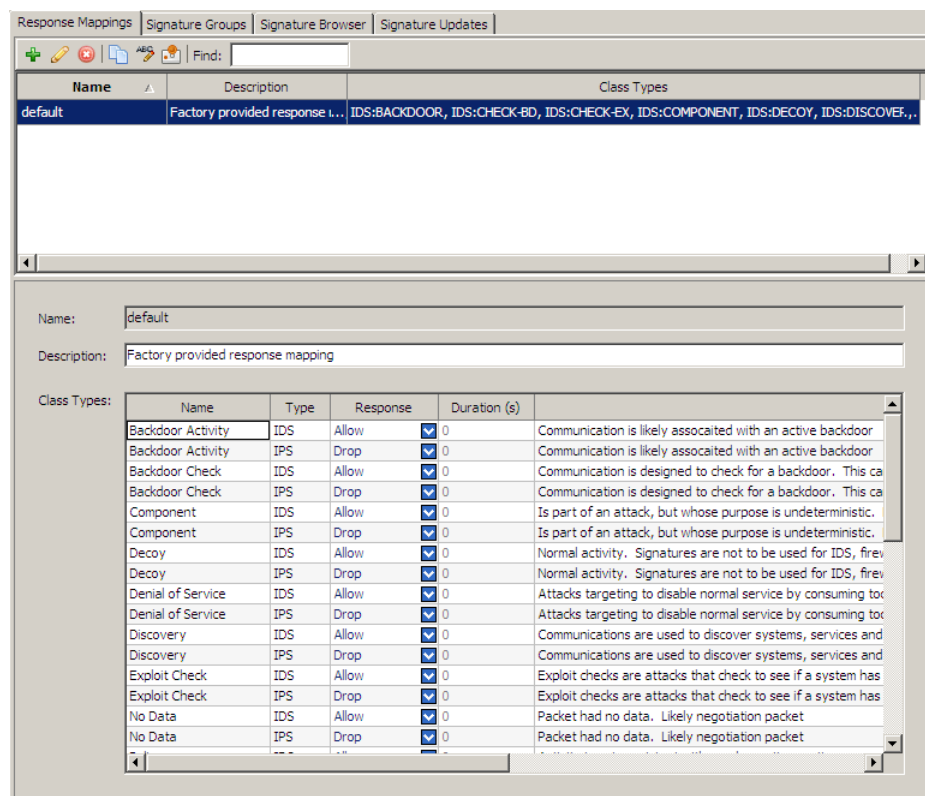
Configuring a response mapping

A response mapping contains a list of class types, their threat level, and their response settings. Each class type refers to a set of known network-based attacks. Class types classified as IPS detect confirmed attacks that are also considered dangerous. Class types classified as IDS detect either suspected attacks or traffic that is considered less dangerous, such as probe or discovery activity. Class types classified as Policy identify traffic based on organizational security practices.

Response mappings are configured on the Response Mapping tab. They can then be selected on the Rules window to indicate how the firewall will respond when an attack is detected.

To configure a response mapping, select **Policy > IPS**. The Response Mappings tab appears.

Figure 64 IPS: Response Mappings tab



The upper pane contains the toolbar and the existing response mappings. When you select a mapping, its properties appear in the lower pane.

Figure 65 Response Mappings toolbar



Use the toolbar and table in the upper pane to perform the actions listed here:


Table 16 Response Mappings toolbar

Icon/ Menu item	Action
New	Create a new response mapping by clicking New . The New Response Mapping window appears.
Modify	Modify a response mapping: <ul style="list-style-type: none"> Select it and modify its properties in the lower pane. Double-click it and modify it in the new window. Select it, click Modify, and edit it in the new window. Note: Read-only administrators can view a response mapping.
Delete	Delete a response mapping by selecting it and clicking Delete .
Duplicate	Create a copy of an existing response mapping by selecting the mapping, clicking Duplicate , and customizing the copy as needed.
Rename	Rename a response mapping by selecting it and clicking Rename .
Usage	View what rules currently use a response mapping by selecting a mapping and clicking Usage .
Find	Search for a specific element(s) in the list using the Find field. Type your search criteria, and response mappings with matching elements will appear in the list. Clear this field to see the full list again.

Create or modify response mappings

When you click **New**, **Modify**, or **Duplicate**, the New/Modify Response Mapping window appears:

Figure 66 New/Modify Response Mapping window



To configure a response mapping:

- 1 In the **Name** field, enter a name that identifies the purpose of the response mapping. For example, if you create two mappings to address different threat levels to your web servers, you would name them “web server high” and “web server low.”

Valid values include alphanumeric characters, dashes (-), underscores (_), and spaces (). However, the first and last character of the name must be alphanumeric. The name cannot exceed 256 characters. You can rename the mapping later.

- 2 [Optional] In the **Description** field, enter any useful information about this mapping. For example, a mapping that allows but audits probe and discovery attacks would have the description *Probe-Discovery audit only*.
- 3 In the **Class Types** area, identify the class types to which you want the firewall to respond by setting the responses to one of the following:

- **Allow no audit** – Allows the traffic to pass and does not generate an IPS audit event. This is the default for all class types when creating a new response mapping.
- **Allow** – Allows the traffic to pass and generates an IPS audit event. Use this setting for traffic that is an anomaly and appears suspicious but is not an identifiable attack.
- **Drop** – Denies only those packets that are suspect while allowing trusted packets. The firewall will not alert the attacker that the connection was closed. This generates an IPS audit event.
- **Deny** – Similar to Drop except that this response sends a TCP reset informing the originating host the connection was deliberately closed. This generates an IPS audit event.

Caution: Use this setting only when troubleshooting or when instructed by Technical Support. Sending a TCP reset or other connection-denied response could notify the attacker that the firewall has recognized the attack, prompting the attacker to switch to a new attack.

- **Deny no audit** – Similar to Deny except that this response does not generate an IPS audit event.
- **Blackhole** – Denies all traffic from the host originating the hostile traffic for a set period of time. This generates an IPS audit event. The firewall will not alert the attacker that the connection was closed. Use this setting when you are sure all traffic coming from an address is malicious.

In the **Duration** field, enter the time in seconds that the traffic will be denied.

- Valid values are 0 and 1–100000 seconds.
- To blackhole the host indefinitely, enter **0**. The host remains blackholed until it is deleted from the blackhole list in the dashboard or the firewall is restarted.

Tip: See the Dashboard to manage the blackholed IP addresses.

- 4 Click **Add**.
- 5 Save your changes.

This response mapping is available for use in a rule.

Configuring a signature group

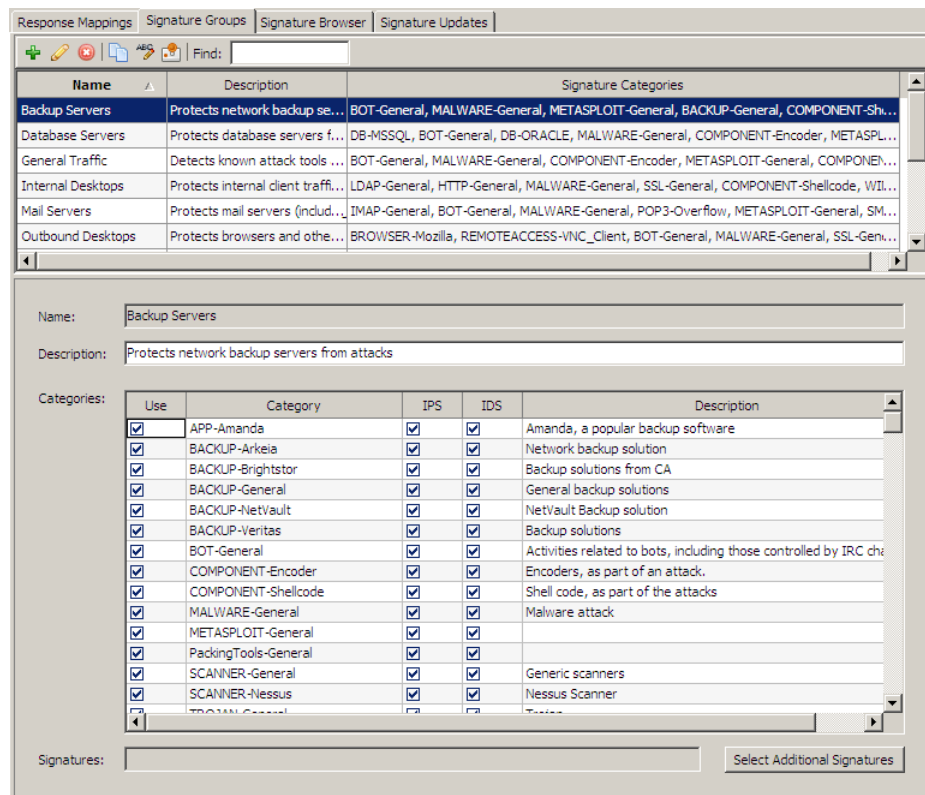
A signature group can contain one or more signature categories. A signature category is a category of signatures that all involve the same type of attack. The IPS engine provides the categories and may update them occasionally.

You can also add individual signatures to a signature group. This gives you finer control in creating a signature group, and it allows you to add Policy signatures, which are not included in the default signature categories since they are specific to an organization.

Signature groups are configured on the Signature Groups tab. They can then be selected on the Rules window to focus IPS inspection on relevant attacks.

To configure a signature group, select **Policy > IPS** and click **Signature Groups**. The Signature Groups tab appears.


Figure 67 IPS: Signature Groups tab



The upper pane contains the toolbar and the existing signature groups. When you select a signature group in the list, the properties of that group appear in the lower pane.

Note: Policy signatures must be added to a signature group by using the **Select Additional Signatures** button.

Figure 68 Signature Groups toolbar



Use the toolbar and table in the upper pane to perform the actions listed here:

Table 17 Signature Groups toolbar

Icon/ Menu item	Action
New	Create a new signature group by clicking New . The New Signature Group window appears.
Modify	Modify a signature group: <ul style="list-style-type: none"> Select it and modify its properties in the lower pane. Double-click it and modify it in the pop-up window. Select it, click Modify, and edit it in the pop-up window. (Read-only administrators can click View to view a signature group.)
Delete	Delete a signature group by selecting it and clicking Delete .
Duplicate	Create a copy of an existing signature group by selecting the group, clicking Duplicate , and customizing the copy as needed.
Rename	Rename a signature group by selecting it and clicking Rename .
Usage	View what rules use a given signature group by selecting a group and clicking Usage .
Find	Search for a specific element(s) in the list using the Find field. Type your search criteria, and signature groups with matching elements will appear in the list. Clear this field to see the full list again.

Create or modify signature groups

When you click **New**, **Modify**, or **Duplicate**, the New/Modify Signature Group window appears.

Figure 69 New/Modify Signature Group window

The window is titled "Signature Category Groups: New Signature Group". It contains the following elements:

- Name:** A text input field.
- Description:** A text input field.
- Categories:** A table with columns: Use, Category, IPS, IDS, and Description.

Use	Category	IPS	IDS	Description
<input type="checkbox"/>	ANTIVIRUS-ClamAV	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	
<input type="checkbox"/>	APP-Amanda	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Amanda, a popular backup software
<input type="checkbox"/>	APP-Ethereal	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Ethereal
<input type="checkbox"/>	APP-iTunes	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Apple iTunes player
<input type="checkbox"/>	APP-Realplayer	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Media player from RealNetworks
<input type="checkbox"/>	APP-Realserver	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	RealNetworks RealServer player
<input type="checkbox"/>	APP-WinAMP	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	WinAMP
<input type="checkbox"/>	APP-WMP	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	MS Windows Media Player
<input type="checkbox"/>	AUTHENTICATION-General	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Authentication
<input type="checkbox"/>	AUTHENTICATION-KERBEROS	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Kerberos
<input type="checkbox"/>	AUTHENTICATION-XTACACS	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	XTACACS
- Signatures:** A text input field and a "Select Additional Signatures" button.
- Buttons:** "Add", "Close", and "Help" at the bottom right.

Use this window to create, modify, or duplicate a signature group.

To configure a signature group:

- 1** In the **Name** field, enter a name that describes the purpose of the signature group. For example, if you wanted a signature category that searches both HTTP and FTP attack signature files, you would name it *HTTP_FTP*.

Valid values include alphanumeric characters, dashes (-), underscores (_), and spaces (). However, the first and last character of the name must be alphanumeric. The name cannot exceed 256 characters. You can rename the mapping later.
- 2** [Optional] In the **Description** field, enter any useful information about this group. For example, a signature category designed to inspect Oracle-related connections would be named *Oracle* and include the categories *DB:Oracle*, *Component:Encoder*, and *Component:Shellcode*
- 3** Configure the **Categories** area:
 - a** In the **Use** column, select each category to include in the signature group.
 - b** For each selected category, select IPS, IDS, or both:
 - Select **IPS** to identify attacks that are an exact match to a signature file.
 - Select **IDS** to identify attacks that are considered minor, such as probe or discovery activity, or suspected attacks, meaning the signature may have incorrectly identified legitimate traffic as an attack.


Both options are selected by default.
- 4** [Optional] Click **Select Additional Signatures** to open a pop-up window and enable individual signatures to add to the signature group. The added signatures appear in the read-only Signatures field.
- 5** Click **Add** and save your changes.

This signature group is now available for use in a rule.

Add individual signatures to a signature group

When you click **Select Additional Signatures**, the Select Additional Signatures window appears.

Figure 70 Select Additional Signatures window



Use this window to add individual signatures to a signature group.

Adding signatures individually gives you finer control in creating a signature group, and it also allows you to add Policy signatures, which are not included in the default signature categories since they are specific to an organization.

You can perform the following tasks:

Filter the table

The table lists available signatures on the firewall, along with information such as category, class type, and type. You can control what appears in the table for easier viewing and faster table loading.

- To view signatures of specific categories, click **Filter Categories**. In the pop-up window, select the signature categories that you want to view. See [Filtering categories](#) for more information.
- To search for a specific element(s) in the table, type your search criteria in the **Find** field and then click **Find Now**. Signatures with matching elements will appear in the table. Clear this field and click **Find Now** to see the full table again.

Note: User-added signatures (enabled with a green check mark) appear no matter what is in the Find field.

View signature vulnerabilities

The Vulnerability column of the table lists a number assigned by Common Vulnerabilities and Exposures (CVE). Two types of identifiers can appear for a signature:

- If **CVE** precedes the number, the vulnerability has been reviewed and accepted by CVE and is an official entry in the CVE list.
- If **CAN** or nothing precedes the number, the vulnerability is under review by CVE and is not yet an official entry in the CVE list.

Select a signature and click **View Vulnerabilities** to open a CVE web page with detailed information about the vulnerability for that signature.

Note: The **View Vulnerability** button is disabled if no identifier exists for the selected signature or if multiple signatures are selected.

Enable and disable signatures

- A blue check mark in the Enabled column indicates that the signature is implicitly included in a category used by the signature group. These signatures cannot be disabled.
Note: To disable an implicit signature, use the Signature Browser tab. This disables the signature globally, meaning it will not be used by any rule to scan traffic.
- A green check mark in the Enabled column indicates that the signature has been added to the signature group. These signatures can be disabled.

To change the status of a signature:

- 1 Select a signature in the table and click **Enable** or **Disable**.
 - You can select multiple signatures by pressing and holding the **Ctrl** key while selecting the appropriate signatures.
 - You can select a range of signatures by selecting the first signature in the range, pressing and holding the **Shift** key, and then selecting the last signature in the range.
- 2 Click **OK**. You return to the Signature Groups tab and the enabled signatures appear in the read-only **Signatures** field.

Managing signatures

Use the Signature Browser tab to view and manage available signatures. You can perform the following actions:

- Filter signatures for easier viewing.
- Enable or disable signatures globally.
- View signature vulnerabilities on the Common Vulnerabilities and Exposures (CVE®) web site.

To manage signatures, select **Policy > IPS** and click the **Signature Browser** tab. The Signature Browser tab appears.

Figure 71 IPS: Signature Browser tab

Enabled	Name	Category	Class Type	Type	Date Added	Description
<input checked="" type="checkbox"/>	MSLAUGH.GET	WORM-General	WORM	IPS	08/17/2005	MSLAUGH.EXE is the pe
<input checked="" type="checkbox"/>	MSBLASTER.GET	WORM-General	WORM	IPS	08/17/2005	The Blaster/Lovsan/Po
<input checked="" type="checkbox"/>	TEEKIDS.GET	WORM-General	WORM	IPS	08/17/2005	TEEKIDS.EXE is the pay
<input checked="" type="checkbox"/>	TEEKIDS.START	WORM-General	WORM	IPS	08/17/2005	TEEKIDS.EXE is the pay
<input checked="" type="checkbox"/>	TEEKIDS.EXE	WORM-General	WORM	IPS	08/17/2005	TEEKIDS.EXE is the pay
<input checked="" type="checkbox"/>	ENBIEI.GET.vB	WORM-General	WORM	IPS	08/17/2005	ENBIEI.EXE is the payl
<input checked="" type="checkbox"/>	ENBIEI.START	WORM-General	WORM	IPS	08/17/2005	ENBIEI.EXE is the payl
<input checked="" type="checkbox"/>	SMBSERVER.Wildcard	SMB-General	DISCOVER	POLICY	08/17/2005	
<input checked="" type="checkbox"/>	NETBIOS.DISCOVERY.POLICY	SMB-Policy	DISCOVER	POLICY	08/17/2005	
<input checked="" type="checkbox"/>	WAREZ.FTP.DELE-TEST.SCAN	FTP-General	SUSPECT	IPS	08/17/2005	An attempt has been m
<input checked="" type="checkbox"/>	WAREZ.FTP.STOR-TEST.SCAN	FTP-General	SUSPECT	IPS	08/17/2005	An attempt has been m
<input checked="" type="checkbox"/>	WINSERVER.OFFSET.B.EXPLOIT	WINS-General	REXPLOIT	IPS	08/17/2005	
<input checked="" type="checkbox"/>	RPC.BINDINIT.B.CHECK	RPC-General	CHECK-EX	IDS	08/17/2005	
<input checked="" type="checkbox"/>	SDBOT.P2.BACKDOOR	WORM-General	WORM	IPS	08/17/2005	The SDBot exploit is a v
<input checked="" type="checkbox"/>	RADMIN.IDSTRING.CHECK	REMOTEACCES...	DISCOVER	IDS	08/17/2005	
<input checked="" type="checkbox"/>	SMTP-TEST.HELO.POLICY	SMTP-Policy	DECOY	POLICY	08/17/2005	SMTP client identifying
<input checked="" type="checkbox"/>	SSH_2dot0.CONNECT_REQ.A.POLICY	SSH-Policy	DECOY	POLICY	08/17/2005	SSH v2.0/libssh connec
<input checked="" type="checkbox"/>	SSH.BYE.A.POLICY	SSH-Policy	DECOY	POLICY	08/17/2005	Client has ended SSH s
<input checked="" type="checkbox"/>	MSLAUGH.P1.START	WORM-General	WORM	IPS	08/17/2005	MSLAUGH.EXE is the pe
<input checked="" type="checkbox"/>	MSLAUGH.P2.START	WORM-General	WORM	IPS	08/17/2005	MSLAUGH.EXE is the pe
<input checked="" type="checkbox"/>	SMB.VERIFYDIALECT.ALLNOTMS.A.PC...	SMB-Policy	DECOY	POLICY	08/17/2005	SMB control message d
<input checked="" type="checkbox"/>	H0no.DAMEWARE.A.P1.CHECK	WORM-General	CHECK-EX	IPS	08/17/2005	Buffer overflow in Dam
<input checked="" type="checkbox"/>	NETMANTIC.DAMEWARE.A.EXPLOIT	WORM-General	REXPLOIT	IPS	08/17/2005	Buffer overflow in Dam
<input checked="" type="checkbox"/>	NETMANTIC.DAMEWARE.B.EXPLOIT	WORM-General	REXPLOIT	IPS	08/17/2005	Buffer overflow in Dam

You can perform the following tasks:

Filter the table

The table lists available signatures on the firewall, along with information such as category, class type, and type. You can control what appears in the table for easier viewing and faster table loading.

- To view signatures of specific categories, click **Filter Categories**. In the pop-up window, select the signature categories that you want to view. See [Filtering categories](#) for more information.
- To search for a specific element(s) in the table, type your search criteria in the **Find** field and then click **Find Now**. Signatures with matching elements will appear in the table. Clear this field and click **Find Now** to see the full table again.

Note: Disabled signatures appear no matter what is in the Find field.

View signature vulnerabilities

The Vulnerability column of the table lists a number assigned by Common Vulnerabilities and Exposures (CVE). Two types of identifiers can appear for a signature:

- If **CVE** precedes the number, the vulnerability has been reviewed and accepted by CVE and is an official entry in the CVE list.
- If **CAN** or nothing precedes the number, the vulnerability is under review by CVE and is not yet an official entry in the CVE list.

Select a signature and click **View Vulnerabilities** to open a CVE web page with detailed information about the vulnerability for that signature.

Note: The **View Vulnerability** button is disabled if no identifier exists for the selected signature or if multiple signatures are selected.

Enable and disable signatures globally

By default, all signatures are enabled and can be used by a rule to scan traffic. An enabled signature is indicated by a green check mark in the Enabled column.

If a signature is disabled, the Enabled check box is cleared and the signature will not be used when scanning traffic, even if it is part of a signature group referenced in a rule. Disabling may help avoid false positives based on signature, for example, if a certain signature is identifying legitimate traffic as an attack.


To change the status of a signature, select a signature in the table and click **Enable** or **Disable**.

- You can select multiple signatures by pressing and holding the **Ctrl** key while selecting the appropriate signatures.
- You can select a range of signatures by selecting the first signature in the range, pressing and holding the **Shift** key, and then selecting the last signature in the range.

Filtering categories

When you click **Filter Categories**, the Category Filter window appears.

Figure 72 Category Filter window



Use this window to populate the Signature Browser tab with signatures of selected categories.

- Select and clear categories individually by clicking the check box in the View column, or use the **Select All** and **Deselect All** buttons.
- Search for a specific element(s) in the list using the **Find** field. Type your search criteria, and signature categories with matching elements will appear in the list. The buttons become **Selected Filtered** and **Deselect Filtered**.

Clear this field to see the full list again.

- Select the **Remember this filter selection** check box to retain the selected categories. The next time you open an IPS Signature Browser, the same category filter will be used.

When you are done selecting the categories of signatures you want to view, click **OK**. Only signatures of the selected categories appear in the Signature Browser.

Configuring IPS signature file updates

Use this tab to configure the IPS signature file update properties. The firewall can automatically download and install updates at intervals that you determine. You can also manually download and install updated signature files at any time.

Note: We recommend downloading the latest signature files prior to enabling IPS inspection on any active rules.

Figure 73 IPS: Signature Updates tab

The screenshot shows the 'Signature Updates' tab in a web interface. At the top, there are four tabs: 'Response Mappings', 'Signature Groups', 'Signature Browser', and 'Signature Updates'. The 'Signature Updates' tab is active. Below the tabs, there are two main sections. The first section, labeled 'Source', contains two text input fields: 'Download Site:' with the value 'downloads.securecomputing.com' and 'Directory:' with the value 'cgi-bin/sigupdate.py'. The second section, labeled 'Download and Install', contains a button 'Download and Install Signatures Now'. Below this button, there is a checkbox 'Enable Automated Signature Download and Install'. If checked, there are three radio buttons for 'Frequency': 'Hourly' (selected), 'Daily', and 'Weekly'. Below the frequency options, there is a 'Day:' dropdown menu set to 'Sunday' and a 'Time:' field set to '11:23 AM'. Below these, there is another checkbox 'Enable Email Notification'. If checked, there is a 'Recipients:' text input field. At the bottom of the 'Download and Install' section, there is a button 'Show Installed Signatures File Version Number Now'.

Use this window to configure the IPS signature file update properties.

Note: While most sites will use provided IPS signature files, the firewall also supports using user-defined files. User-defined IPS signature files can only be created or updated using the command line interface. See Knowledge Base article [8959](#) for details.

To configure updates:

1 In the **Source** area, verify/modify the following fields:

Caution: Changing these defaults may prevent the firewall from obtaining updated signatures file, resulting in inadequate IPS protection.

- **Download Site** – This is the site from which the package will be downloaded. The default site is <https://support.forcepoint.com/Downloads>.

Note: If the download fails, troubleshoot the problem by verifying that the site name resolves to an IP address and is reachable from the Sidewinder.

- **Directory** – The path name on the download site that contains the update. The default directory is: `cgi-bin/sigupdate.py`

2 [Conditional] To configure automatic signature file updates, follow the sub-steps below. To manually download and install the signature files, skip to [Step 3](#).

a Select **Enable Automated Signature Download and Install**.

b In the **Frequency** field, specify how frequently you want to download and install updated signature files:

- (Recommended) To download and install every hour, select **Hourly**.
- To download and install every day, select **Daily**.
- To download and install once a week, select **Weekly**. Also specify the day of the week on which you want downloads to occur.

c For all frequency options, in the **Time** field, specify the time of day you want the firewall to download and install the updates.

Note: Downloading and installing updates has a minimal impact on your system. Traffic that is received while the download and installation are in process will be inspected using the current version. Once installation is complete, all traffic will be scanned using the updated information.

d If you want to receive e-mail notification when the updates are downloaded and installed, select the **Enable Email Notification** check box. If you select this option, you will also need to specify an e-mail address in the **Recipient** field.

3 [Conditional] To update the signature files manually, click **Download and Install Signatures Now**. A progress bar appears while the files are downloaded, then a message appears stating that the update is complete.

4 To view the current version of the signature file you are using, click **Show Installed Signatures File Version Number Now**. An Info window appears displaying the current installed version. When you are finished viewing the version, click **OK**.

5 Save your changes.

The IPS engine is now using the current signature files.

Configuring virus scanning services

The anti-virus service is a licensed add-on module that uses a firewall-hosted virus scanner that allows you to configure rule-based MIME, virus, and spyware scanning. Use scanning services on HTTP and HTTPS traffic, FTP files, and mail messages. When using scanning services, you can specify the number of server processes to be dedicated to various data sizes, allowing the firewall to process data more efficiently. You can also configure how often to update the signature files. Before the firewall will scan traffic for viruses, you must ensure the following conditions have been met:

- The Anti-Virus feature must be licensed. To verify that the feature has been licensed, select **Maintenance > License**, and click the **Firewall** tab. If you are not licensed for Anti-Virus, contact your sales representative.
- The rules governing the traffic you want filtered must have the appropriate Application Defenses options selected:
 - To scan web traffic, create rules using an HTTP or HTTPS application defense with the MIME/Virus/Spyware option configured. See [Creating HTTP or HTTPS Application Defenses on page 198](#) for more information.
 - To scan mail messages, create rules using the sendmail server and a Mail (Sendmail) application defense with the MIME/Virus/Spyware option configured. See [Creating Mail \(Sendmail\) Application Defenses on page 212](#) for more information.
 - To scan FTP session, create rules using an FTP application defense with the Virus/Spyware Scanning option configured. See [Creating FTP Application Defenses on page 228](#).

To configure scanning services, select **Policy > Application Defenses > Virus Scanning**. The Virus Scanning window appears with the Signature tab displayed.

Configuring virus scanning signature updates

Use this tab to configure the anti-virus update properties. The firewall can automatically download and install updates at intervals that you determine. (This window mainly deals with updating signature files, but has an option to make sure the virus engine is also up-to-date.) You can also manually download and install updated signature files at any time.

Note: We recommend downloading the most recent engine patch (for example, 701MCV01) and the latest signature files prior to enabling anti-virus services.

Figure 74 Virus Scanning: Signature tab

The screenshot shows the 'Signatures' tab with the 'Advanced' sub-tab selected. The 'Source' section contains two text fields: 'Download Site' with the value 'av.sidewinder.downloads.forcepoint.com' and 'Directory' with the value 'cgi-bin/avupdate'. The 'Download and Install' section features a 'Download and Install Signatures Now' button. Below this are three checkboxes: 'Enable Automated Signature Download and Install' (checked), 'Enable Automated Scanner Engine Updates' (unchecked), and 'Enable Email Notification' (unchecked). The 'Enable Automated Signature Download and Install' section includes a 'Frequency' dropdown set to 'Hourly', a 'Day' dropdown set to 'Wednesday', and a 'Time' field set to '02:30 PM'. At the bottom of the 'Enable Email Notification' section is a 'Recipient' text field. A 'Show Installed Signatures File Version Number Now' button is located at the bottom of the window.

To configure the anti-virus update properties:

1 In the **Source** area, verify/modify the following fields:

Caution: Changing these defaults may prevent the firewall from obtaining updated signatures file, resulting in inadequate virus and spyware protection.

- **Download Site** – This is the name of the site from which the package will be downloaded. The default site is support.forcepoint.com/Downloads.

If the download fails, troubleshoot the problem by verifying that the site name resolves to an IP address and is reachable from the Sidewinder.

- **Directory** – The path name on the download site that contains the update. The default directory is *cgi-bin/avupdate*.

- 2 [Conditional] To configure automatic virus updates, follow the sub-steps below. To manually download and install the signature files, skip to [Step 3](#).

Automatically updating both the scanner engine and the signature files is strongly recommended. If your engine is out-of-date, the firewall will not install the most recent files.

Note: Failure to regularly update your anti-virus engine and signature files will result in inadequate virus and spyware protection. For best results, also select **Automatically check for and load packages** (**Maintenance > Software Management > Download Packages** tab).

a Select **Enable Automated Signature Download and Install**.

b In the **Frequency** field, specify how frequently you want to download and install updated signature files:

- (Recommended) To download and install every hour, select **Hourly**.
- To download and install every day, select **Daily**.
- To download and install once a week, select **Weekly**. Also specify the day of the week on which you want downloads to occur.

c For all frequency options, in the **Time** field, specify the time of day you want to download and install the updates.

Note: Downloading and installing updates has a minimal impact on your system. Traffic that is received while the download and installation are in process will be scanned using the current version. Once installation is complete, all traffic will be scanned using the updated scanner information.

d Select **Enable Automated Scanner Engine Updates** to automatically check for new loaded (but not installed) anti-virus engine updates (for example, patch 701MCV01) when installing new virus signature files. If an uninstalled engine update exists, the firewall will install it the next time it installs the new signature files. This installation does not interrupt system processes.

e If you want to receive e-mail notification when the updates are downloaded and installed, select the **Enable Email Notification** check box. If you select this option, you will also need to specify an e-mail address in the **Recipient** field.

f Proceed to [Step 5](#).

- 3 [Conditional] To update the virus definition manually, do the following:

a Click **Download and Install Signatures Now**. A pop-up window appears.

b Click **Background** to perform the update in the background, or click **Wait** to receive a notification and status pop-up when the update is complete. Proceed to [Step 5](#).

- 4 To view the current version of the signature file you are using, click **Show Installed Signatures File Version Number Now**. An Info window appears displaying the current installed version. When you are finished viewing the version, click **OK**.

- 5 Save your changes.

The virus scanner is now using a supported engine and the current signature files.

Configuring the advanced virus scanning features

Use this tab to configure how your firewall distributes scanner processes for incoming and outgoing traffic. This is done by configuring the number of scanners to be run to service each of the defined file size ranges.

Figure 75 Virus Scanning: Advanced tab

File Size Range	Scanners
Up to 40K	2
Up to 100K	2
Up to 1M	1
Unlimited	1

Modify

Scan Buffer Size (KB): (Recommended: 50 KB, Maximum: 64 KB)

Archive Scan Buffer Size (MB): (Recommended: 128 MB, Maximum: 512 MB)

Maximum Number of Files to Scan in an Archive:

☐ Scan Encrypted Files

Show Installed Engine Version Now

- The **File Size Range** column displays the size limits for each range. Ranges are **Up to 40K**, **Up to 100K**, **Up to 1MB**, and **Unlimited**.
- The **Scanners** column displays the number of scanner processes dedicated to each range.

You cannot modify the existing size ranges or add new size ranges in the Admin Console.

Files are handled by the first file size range that is greater than the file's size. For example, a 39K file will be processed by a scanner process assigned to the *Up to 40K* file size range, but a 40K file will be processed by a scanner process assigned to the *Up to 100K* file size range.

Tip: While using additional scanners may speed up virus scanning, it can slow down your firewall's overall performance. Try using more restrictive MIME/Virus/Spyware rules, configured on the Application Defenses, to make virus scanning more efficient.

This tab also allows you to view the current virus scanner engine version.

To configure virus scanning's advanced properties:

- 1 To configure the number of scanner processes for a particular size range, select the file size range in the table and click **Modify**. The Edit Scanners window appears.

In that window's **Scanners** field, specify the number of scanner processes you want to dedicate for the selected group. Valid values are between 1 and 10, and the total number of scanner processes should not exceed a combined total of 20 processes. (Configuring more than 20 total processes may have a negative impact on performance, particularly on smaller firewalls.) Click **OK** to return to the Advanced tab.

Note: If you decrease the number of scanners, you must restart the virus scanner on the **Monitor > Service Status** window.

- 2 In the **Scan Buffer Size** field, specify the size of information (in KB) that can be held in the memory buffer before a backup file is created to temporarily hold the traffic for processing. This value must be between 8KB and 64KB. The default value is 50KB.

- 3 In the **Archive Scan Buffer Size** field, specify the amount of memory to be used to contain the contents of archive files before the anti-virus engine temporarily writes the contents to disk to perform the virus scan. The default is 128 MB.
- 4 In the **Maximum Number of Files to Scan in an Archive** field, specify the maximum number of files to be scanned within an archive (such as a .zip file, etc.). If the number of files in an archive exceeds the number specified in this field, scanning does not take place.
- 5 [Optional] The **Scan Encrypted Files** option controls how the Virus Scanner behaves when it scans password-protected files (primarily .xls and .zip files), which the scanner classifies as encrypted. This is relevant for mail attachments, HTTP traffic, and FTP transmissions. Determine how the scanner will handle encrypted files by doing one of the following:
 - If you leave this option clear, the scanner generates an error and rejects the password-protected files.
 - If you select this option, the scanner ignores those errors and scans any unencrypted parts of the file. If no virus is found, the file is allowed.
- 6 To view the virus scanner engine version number that is currently installed, click **Show Installed Engine Version Number Now**. A pop-up window appears displaying the current version. To close the pop-up window, click **OK**.
- 7 Save your changes.

The changes to virus scanning's advanced properties are now applied.

About TrustedSource

TrustedSource is a reputation service from that assigns a reputation score to an IP address based on the behavior attributes of the traffic it generates. A reputation score is like a credit score that indicates the trustworthiness of an IP address.

TrustedSource uses servers around the world to gather and analyze billions of packets dynamically to determine reputation scores. For each IP address on the internet, TrustedSource calculates a reputation value based on such attributes as sending behavior, blacklist and whitelist information, and spam trap information.

Note: See the TrustedSource web site at www.trustedsource.org for more information about the service.

Using TrustedSource on your Sidewinder can:

- block spam e-mail from botnets.
- help prevent hosts on your network from being infected with botnet agents.
- identify hosts on your network that have been compromised in botnet or pharming attacks.
- protect critical servers from access by authorized users inadvertently using external machines that are compromised.

For more information, see the TrustedSource application note at <https://support.forcepoint.com>.

Using TrustedSource on a Sidewinder

Use TrustedSource on a Sidewinder to more accurately filter network traffic passing through the firewall.

- You can use TrustedSource in a rule to inspect traffic for a reputation score ranging from Trusted to Malicious.
- You can filter incoming mail connections by allowing messages only from senders with a reputation score below a defined threshold.

Using TrustedSource in rules

TrustedSource can be enabled on an inbound or outbound rule that uses a proxy or server service. When a packet is examined by the rule, the firewall queries a TrustedSource server to get the reputation score of all IP addresses involved in the connection.

- You can whitelist objects to exempt them from TrustedSource rule requirements. This is useful for routable internal addresses or trusted external sources.

You create a whitelist on the TrustedSource window: **Policy > Application Defenses > TrustedSource**.

- Private IP addresses are not evaluated by TrustedSource or examined in rules (for example, 10.x.x.x, 172.16.x.x, 192.168.x.x).

- TrustedSource queries are cached, so another query to a TrustedSource server is not made for an IP address that has recently been examined by the firewall.

Traffic is not explicitly allowed or denied based on a TrustedSource reputation score. The score is one of the elements in the rule that is examined for a match.

- In an *allow* rule, the **Unverified** to **Trusted** side of the TrustedSource slider is active by default. IP addresses with a good reputation will match this rule.
 - If the reputation score is within the **Unverified** to **Trusted** range marked by the slider, and all other elements in the rule match, the connection is allowed. No other rules are queried.
 - If the reputation score is left of the **Unverified** to **Trusted** range marked by the slider, it is not a match. The connection is passed to the next rule.

Figure 76 TrustedSource on an allow rule

The screenshot displays the configuration interface for a firewall rule, divided into two main sections: 'General' and 'TrustedSource'.

General Section:


- Action:** Three radio buttons are present: 'Allow' (selected), 'Deny', and 'Drop'.
- Service:** A dropdown menu is set to 'http (HTTP Proxy)', with a small '...' button to its right.
- Audit:** A dropdown menu is set to 'Standard (recommended)'.

TrustedSource Section:

- Enable TrustedSource:** A checkbox is checked.
- Slider:** A horizontal slider is positioned between the 'Unverified' and 'Trusted' markers. Below the slider, five reputation levels are labeled: 'Malicious', 'Suspicious', 'Unverified', 'Neutral', and 'Trusted'.
- Text:** To the right of the slider, a text box states: 'Unverified, neutral and trusted traffic will match this rule. (Range 29 to -255)'.

- In a *deny* or *drop* rule, the **Suspicious** to **Malicious** side of the TrustedSource slider is active by default. IP addresses with a bad reputation will match this rule.
- If the reputation score is within the **Suspicious** to **Malicious** range marked by the slider, and all other elements in the rule match, the connection is denied or dropped. No other rules are queried.
- If the reputation score is right of the **Suspicious** to **Malicious** range marked by the slider, it is not a match. The connection is passed to the next rule.

Figure 77 TrustedSource on a deny or drop rule



A reputation is expressed in five classes:

- **Trusted** – The IP address is a source of substantial amounts of legitimate traffic.
- **Neutral** – The IP address is a source of legitimate traffic, but may send small amounts of unusual traffic or traffic requiring further inspection.
- **Unverified** – The IP address may be a legitimate sender, but data gathered to date has been either inconclusive or insufficient to make a firm reputation decision.
- **Suspicious** – The IP address has exhibited substantial suspicious behavior in the past, and connections should be treated with caution appropriate to the application protocol in question.
- **Malicious** – The IP address has a history of malicious behavior.

Using TrustedSource to filter e-mail

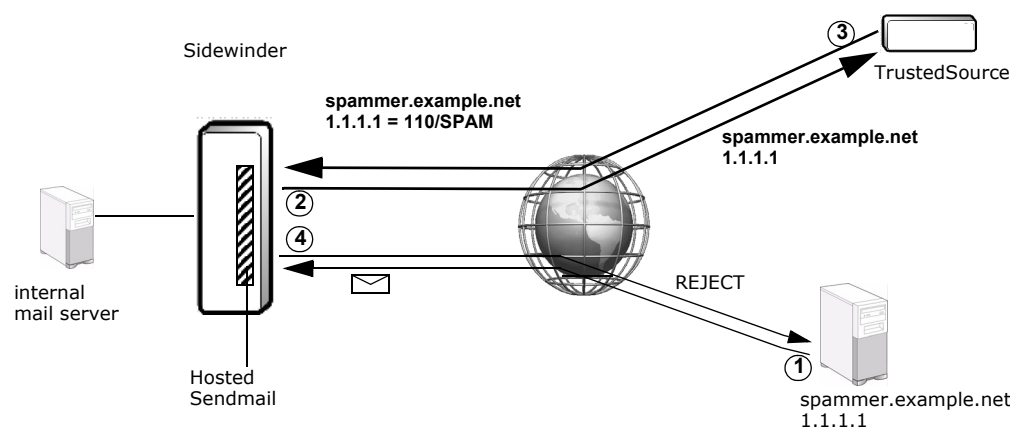
If you use sendmail, you can use TrustedSource to filter incoming mail connections.

- 1 You enable TrustedSource Filtering on the TrustedSource window and set a threshold value for incoming mail.
- 2 A sending server contacts a Sidewinder running hosted sendmail.
- 3 The firewall sends a modified DNS query using the server's IP address to a TrustedSource server to get its reputation score.
- 4 The firewall compares the score to the threshold value you set.
 - If the score is lower than that threshold, e-mail messages from the server are accepted.
 - If the score is higher than the threshold, the firewall rejects the message, audits the violation, and closes the connection.

To determine reputation scores, TrustedSource uses servers around the world to gather and analyze billions of messages dynamically. TrustedSource assigns a score between 0 and 120 to an IP address based on the type of mail, legitimate and spam, this particular host generates. The TrustedSource servers are constantly communicating so as one server identifies a spam flood in progress, it can alert all TrustedSource servers moments after the attack starts and update that sender's reputation score.

The process works like a real-time blackhole list. The Sidewinder administrator can configure what score is a tolerable threshold for your network. A sending server contacts a Sidewinder running hosted sendmail. The firewall then sends a modified DNS query using the server's IP address to a TrustedSource server to get its reputation score. The firewall then compares the score to the threshold value. If the score is lower than that threshold, e-mail messages from the server are accepted. If the score is higher than the threshold, the firewall rejects the message, audits the violation, and closes the connection. This process is illustrated in the following figure:

Figure 78 Example of a TrustedSource e-mail query



To filter mail using TrustedSource, the firewall must be located on your network's perimeter, be configured for hosted sendmail, and have functioning DNS with access to the Internet. Licensing is handled by the TrustedSource server as opposed to the Sidewinder license. Once enabled, TrustedSource automatically starts filtering all inbound mail; you do not need to alter the existing mail rules or create new rules.

With spammers, rejecting one mail message and closing the connection is rarely enough to protect your network from them. Even though the malicious content is prevented from entering your network, the server typically attempts to resend its message. The processing effort and bandwidth to continuously query TrustedSource and reject each message can constitute a Denial of Service attack.

We recommend that in addition to enabling TrustedSource filtering, you configure an IPS attack response that is triggered by the audit violation and that blackholes all traffic coming from the untrusted server. In addition to silently dropping that host's incoming connections, blackholing immediately closes all existing connections with that host. This is particularly useful if the sender's reputation score was updated after the spam flood began.

Configuring TrustedSource

Use the TrustedSource window to configure global TrustedSource settings for rules and mail filtering.

- Create a whitelist for TrustedSource queries, directly query a host's reputation, and adjust reputation boundaries. Settings you configure apply to all rules that have TrustedSource enabled.
- Enable TrustedSource Filtering and set the reputation threshold for inbound mail. TrustedSource filtering will be performed on all inbound mail.

Ensure that your firewall:

- is using hosted sendmail.
- has DNS set up with access to the Internet.
- is on your network's perimeter.

To configure TrustedSource, select **Policy > Application Defenses > TrustedSource**. The TrustedSource window appears.

Figure 79 TrustedSource window

The screenshot shows the TrustedSource configuration window with the following sections:

- TrustedSource Whitelist**
 - Do not perform TrustedSource on:
 - ☒ IP Address objects
 - ☒ IP Range objects
 - ☒ Subnet objects
 - ☒ Host objects
 - ☐ Burbs except the following: [text field]
 - Do not perform TrustedSource on these objects: [empty list box]
 - [Edit button]
- Audit**
 - ☐ Audit traffic allowed by TrustedSource
- Tools**
 - Query reputation at trustedsource.org:
 - Host: [text field]
 - [Query button]
- TrustedSource Filtering**
 - ☐ Perform TrustedSource filtering on inbound mail
 - Reputation threshold: [80]
- Advanced Settings**
 - Adjust reputation boundaries:
 - Malicious: 49 (255 to 50)
 - Suspicious: 29 (49 to 30)
 - Unverified: 14 (29 to 15)
 - Neutral: -1 (14 to 0)
 - Trusted: -1 (-1 to -255)
 - [Color gradient bar]
 - Default reputation if TrustedSource servers are unavailable: [30]
 - [Restore Defaults button]

Configuring TrustedSource settings for rules

TrustedSource is used in rules to examine reputation scores of IP addresses for inbound and outbound traffic.

- You create a TrustedSource whitelist and configure reputation boundaries on the TrustedSource window. Settings on the TrustedSource window are global and apply to all rules that have TrustedSource enabled.
- You enable TrustedSource for individual rules on the New/Modify Rule window. See [Creating, modifying, and duplicating rules on page 273](#) for more information.

You can perform these actions on the TrustedSource window:

- [Create a TrustedSource whitelist.](#)
- [Query a host's reputation.](#)
- [Adjust the reputation boundaries.](#)

Create a TrustedSource whitelist.

In the TrustedSource Whitelist area, select objects to include in the TrustedSource whitelist. Selected objects will not be examined for TrustedSource reputation scores and will be exempt from a rule's TrustedSource matching requirement.

- **Add object types to the whitelist.**

To include all objects of a type in the TrustedSource whitelist, select the object in the **Do not perform TrustedSource on** list. You can include the following:

- **IP Address objects**
- **IP Range objects**
- **Subnet objects**
- **Host objects**

These objects are selected by default because your security policy most likely defines allow and deny rules for these objects.

To exclude an object type from the TrustedSource whitelist, clear the object's check box. All objects of that type will be included in TrustedSource queries and will be subject to a rule's TrustedSource matching requirements.

- **Add individual objects to the whitelist.**

a Clear the object type in the **Do not perform TrustedSource on** list.

b Click **Edit** and select objects that you want to whitelist in the pop-up window.

- **Select burbs to be examined by TrustedSource.**

Select burbs to exclude from the whitelist. These burbs will be examined by TrustedSource and will be subject to a rule's TrustedSource matching requirements.

- You may want to have external burbs and internal burbs with routable IP addresses evaluated by TrustedSource.
- Private IP addresses are not evaluated by TrustedSource or examined in rules (for example, 10.x.x.x, 172.16.x.x, 192.168.x.x).

To exclude burbs from the whitelist:

a Select **Burbs except the following**.

b Select burbs to exclude from the whitelist:

- To exclude a single burb, select it from the drop-down list.
- To exclude multiple burbs, click the ... button and select burbs from the pop-up window.

- **Include audit for allowed traffic.**

Select the **Audit traffic allowed by TrustedSource** check box to include the reputation scores of an allowed connection's IP addresses in the audit log. If the check box is selected and TrustedSource is used to look up the reputation of the source and/or destination IP address of a connection that is allowed, it appears in the audit log like this example:

```
dest_reputation: 20
```

An allow audit message appears in the audit log only if TrustedSource was used in the rule matching process. It will not appear in the audit log for allowed connections under these conditions:

- Both the source and destination IP addresses are on the TrustedSource whitelist.
- The connection is allowed by a rule before a rule that uses TrustedSource.
- The connection does not match another element in the rule using TrustedSource (for example, the destination burb did not match), but is allowed by a subsequent rule that does not use TrustedSource.

Query a host's reputation.

Use the Tools area to directly query a host's reputation on the TrustedSource web site: Enter the host name in the **Host** field and click **Query**. A TrustedSource Feedback web page for the specified host opens.

Adjust the reputation boundaries.

Most users will not need to change the reputation boundaries or the default reputation if TrustedSource servers are unavailable. If you do make adjustments and want to revert to the original settings, click **Restore Defaults**.

Configuring TrustedSource mail filtering

Before you enable TrustedSource filtering for inbound mail, make sure that your firewall:

- is using hosted sendmail.
- has DNS set up with access to the Internet.
- is on your network's perimeter.

Also ensure that you have obtained a TrustedSource subscription. If you do not have a TrustedSource subscription, contact your channel partner or sales representative.

To enable TrustedSource mail filtering:

- 1 Select **Perform TrustedSource filtering on inbound mail**.
- 2 In the **Reputation threshold** field, set the threshold to a value from 0 to 120. Messages from senders with reputation scores above that value are rejected. The default threshold is 80.

Trustworthy senders receive low scores and untrustworthy senders receive high scores. The values map to five reputation classes:

Table 18 TrustedSource reputation classes

Value	Class	Class description
0	Inoffensive	The IP address is a legitimate sender or a source of substantial amounts of legitimate e-mail.
1–25	Neutral	The IP address is likely a legitimate sender but may send small amounts of e-mail requiring further inspection.
26–50	Unverified	The IP address may be a legitimate sender but displays a few properties suggesting further content inspection of e-mails received from that address.
51–80	Suspicious	The IP address shows many spam sender characteristics, and e-mail received from this address may be subject to higher scrutiny.
80+	Spam	The IP address has either been used to send spam or should not send any e-mail messages in general.

- 3 Click **Save**.

The firewall now uses the TrustedSource reputation service to filter inbound e-mail.

To blackhole senders with ratings above the set threshold

- 1 Select **Monitor > IPS Attack Responses**.
- 2 Select the **TrustedSource** attack response.

Its preconfigured settings are:

- **Attack Frequency** – Always Respond
- **Alerts** – Send e-mail, and wait 120 seconds between alerts

- **Strikeback** – Blackhole each host responsible for 100% of the attacks for 21600 seconds (6 hours)

3 Right-click the TrustedSource attack response and select **Enable**.

4 Save your changes.

The Sidewinder now blackholes hosts that have TrustedSource scores that do not meet the set threshold and are trying to send mail to your network. Use the Blackholed IPs feature on the Dashboard to manage blackholed IP addresses.

For more information on TrustedSource, visit www.trustedsource.org.

Updating the Geo-Location database

Geo-Location identifies the country of origin of an IP address. You can create a Geo-Location network object and apply it to a rule to allow or deny a connection based on the source or destination country. See [About the Network Objects: Geo-Location window on page 66](#) for information about creating a Geo-Location network object.

A Geo-Location database on the Sidewinder stores the country IP information that is examined by your policy. Use the Geo-Location Settings window to update the Geo-Location database with the latest country IP information. You can also schedule automatic updates and configure e-mail to notify you when updates are downloaded and installed.

To configure Geo-Location database updates, select **Policy > Application Defenses > Geo-Location Settings**. The Geo-Location Settings window appears.

Figure 80 Geo-Location Setting window

You can configure the following settings:

Update the source of the Geo-Location database downloads

Caution: Changing these defaults may prevent the firewall from obtaining updated databases.

In the Update Source area, you can configure the following fields:

- **Site** – Enter the name of the site the database will be downloaded from. The default site is support.forcepoint.com/Downloads.

If the download fails, troubleshoot the problem by verifying that the site name resolves to an IP address and is reachable from the Sidewinder.

- **Directory** – Enter the path on the download site that contains the update. The default directory is *cgi-bin/geoupdate.py*.

Click **Restore Defaults** to restore these fields to the default locations.

Manually update the Geo-Location database

Click **Update Database Now**. A confirmation message appears with the new database version.

Schedule automatic database updates

- 1 Select **Enable automated database updates**.
- 2 In the **Frequency** field, select how frequently you want to download and install updated database files. hourly, daily, or weekly. Specify the day of the week for weekly downloads.
- 3 In the **Time** field, specify the time of day you want to download and install the updates.

Configure e-mail to notify you when updates are downloaded and installed

- 1 Select **Enable email notification**.
- 2 In the **Recipient** field, enter the e-mail address that will receive the notifications.

Check the version of the current database

Click **Show Database Version**. An information window appears with the database version.

Configuring SmartFilter for Sidewinder

SmartFilter is not actively supported in Sidewinder 7.x.

7 Services

Contents

[About services](#)

[Using the main Services window](#)

[Configuring proxy agents and services](#)

[Configuring packet filter agents and services](#)

[Configuring server agents](#)

[Configuring additional proxy agent properties](#)

About services

On the Forcepoint Sidewinder, policy is applied primarily by rules, which are made up of many elements. The table below shows the progression of a rule's creation using these elements and their corresponding chapters in this guide.

You are here in the Policy section	Use this chapter to...
Chapter 3, Policy Configuration Overview	understand the policy creation process.
Chapter 4, Network Objects and Time Periods	create or modify any network objects or time periods that will be used by rules.
Chapter 5, Authentication	create or modify authenticators that will be used by rules.
Chapter 6, Content Inspection	configure content inspection methods that will be used by rules.
Chapter 7, Services	create or modify services or service groups that will be used by rules.
Chapter 8, Application Defenses	create or modify Application Defenses that will be used by rules.
Chapter 9, Rules	create rules using the elements you created in the previous chapters in the policy section.

A firewall service associates a traffic's transport layer with a specific agent that is responsible for managing the service's traffic. The transport layer information includes elements such as the protocol, the ports, and the idle timeout. Rules use services, along with source and destination information, to determine what traffic that rule will allow or deny. You create a service by selecting an agent, assigning it specific transport-layer properties, and then giving it a name and saving it.

An *agent* is responsible for handling traffic and can be one of these types:

- Proxy (see [Configuring proxy agents and services](#))
- Filter (see [Configuring packet filter agents and services](#))
- Server (see [Configuring server agents](#))

The *proxy* and *filter* agents can be used to create new services. Their configurable service properties vary widely. An agent's properties can be very basic, such as the Ping Proxy agent, which only allows configuration of response timeout and fast path information. Other agents have more options, such as the Telnet Proxy agent, which includes ports, timeouts, fast path information, and connection transparency (transparent, non-transparent, or both).

The *server* agents (also called *daemons*) cannot be used to create new services. All server services are created during the initial configuration and cannot be deleted. You can modify the server services' basic properties, such as port and timeout values. Some servers also have advanced properties that may need to be configured, such as the sendmail server's configuration files.

Some of the agents have *global properties*, meaning the values apply to every service using that agent. When this value involves a connection, the connection total is a sum off all connections using that agent, even if they are distributed through multiple services and rules. Proxy agents' global property setting controls the number of proxy instances running based on the expected connection volume. Filter agents' global property settings controls the maximum number of TCP sessions, UDP sessions, and the port range reserved for filter sessions.

All services are disabled until they are used in an enabled rule. The first time a server, filter, or proxy service is used in an enabled rule, the service is enabled (posts a listen) in the source burb or burbs. When all rules using a given service are disabled or deleted, the service is automatically disabled.

Services

Using the main Services window

Note: To view which services are currently running and if they're running as expected, go to **Monitor > Service Status**. This window displays which services are enabled and where they are being used (which rules, ports, etc.). It also gives you the ability to stop or restart a service, if necessary. See [Chapter 12, Service Status](#) for more information on monitoring service status.

When planning your security policy, study the agents and the default services to determine which ones you will need and what values to assign them. Consider the following:


- Decide what type of inspection is needed for each allowed service. Proxy agents inspect traffic at the application layer. Filter agents tend to inspect traffic at the transport layer.
- When possible, use an application-aware or protocol-aware proxy agent instead of a generic proxy agent. When choosing between the Generic Proxy and the TCP/UDP Packet Filter agent, always try to use the proxy agent because it does not allow the client to connect directly to the server. Instead, the firewall maintains a separate connection to server on the client's behalf, thereby providing more security. See [Configuring proxy agents and services](#) for more information.
- Consider how traffic will get from one burb to another. Ensure the appropriate routing is in place and that you know what connection types are needed (transparent, non-transparent, or both).
- Review the server services to see which ones your policy requires, and which of those servers need modification. Some servers have advanced properties, such as the ability to add extended authentication to the ISAKMP server or to modify the single sign-on (SSO) server's banners.

Security Alert: There is a security risk involved with using non-application aware services. The firewall has greater control over traffic managed by proxies because it can manipulate independent proxy connections on each side of the firewall.

Using the main Services window

To view the available services, select **Policy > Rule Elements > Services**. The main Services window appears.

Figure 81 The main Services window



Name	Agent	Ports	Summary
aol	Generic Proxy	5190/tcp	Fast path=On, TCP idle timeout=3600
dns	DNS Proxy	53 (tcp and udp)	TCP idle timeout=3600, UDP idle timeout=60
finger	Generic Proxy	79/tcp	Fast path=On, TCP idle timeout=1800
ftp	FTP Proxy	21/tcp	Connection type=transparent, Fast path=On, TCP idle timeout=7200

This window is the main window for viewing and creating services. You can perform several tasks directly from this window. Use the toolbar or the right-click menu, shown here, to perform the tasks listed in [Table 19](#).

Figure 82 Tasks available in the Services window

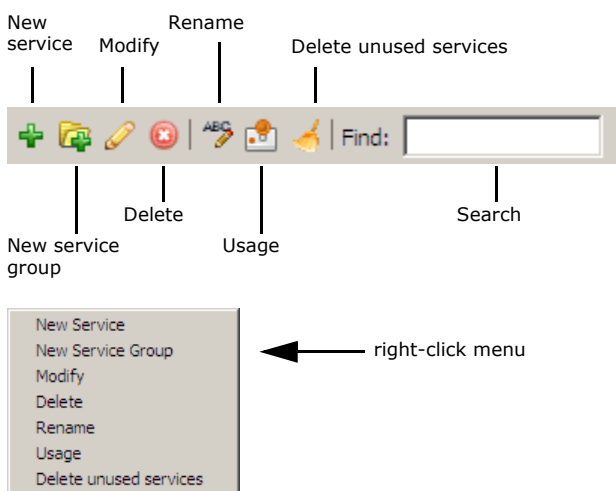


Table 19 Tasks that can be performed from the main Services window


Icon/ Menu item	Task
New Service	Create a new service by clicking New Service . The New Service window appears. See Create and modify services and Configuring server agent properties for more information.
New Service Group	Create a new service group by clicking New Group . The New Service Group window appears. See Create and modify service groups for more information
Modify	Modify a service or service group by double-clicking it, or selecting the item and then clicking Modify . (Read-only administrators can click View to view a service or service group.) <ul style="list-style-type: none"> For services, this opens the Modify Service window. See Create and modify services and Configuring server agent properties for more information. For service groups, this opens the Modify Service Group popup, where you modify the group's description and selected services.
Delete	Delete a service or service group by selecting the item(s) to delete and clicking Delete .
Rename	Rename a service or group by clicking Rename .
Usage	View what rules and rule groups use a service or service group by selecting an item and then clicking Usage .
Delete unused services	Delete services that are not in use by clicking Delete unused services . The Delete unused objects window appears. Select the services that you want to delete and then click OK .
Find	Find a service or service group by entering a character string related to the item you are searching for in the Find field. The search function searches all columns, and filters as you type. For example, if you are searching a service based on the HTTP proxy, typing "http" reduces the list to only the services containing that character string. Clear the Find field to show all options again.

Create and modify services

Use the New/Modify Service window to create or modify services. Once a service is saved, it is available for use in rules. When you create a rule that uses a new service and that rule is enabled, the firewall automatically enables the service's agent in the rule's source burbs which begins managing traffic using this agent. The firewall disables an agent when all rules using that agent are deleted or disabled.

Note: Once a service has been saved, you cannot modify its agent.


Several different actions provide access to a service:

- Select **Policy > Rule Elements > Services**, then click **New** to create a new proxy or filter service.
- Select **Policy > Rule Elements > Services**, then double-click a service (or select it and then click **Modify**) to change an existing service. You can change the service's description, its service properties, or its global agent properties.
 - Use **Rename** to change a service's name.
 - Read-only administrators can click **View** to view a service.
- Select **Policy > Rules**, open a rule and, next to the Service field, click . On the Services popup, click **New > Service**.

Note: You cannot create new servers, or rename or delete existing servers.

The New/Modify Service window appears.

Figure 83 The new/modify service window



To add or modify a service:

- 1 In the **Name** field, type a descriptive name that quickly identifies this service.
 - Valid values include alphanumeric characters, periods (.), dashes (-), underscores (_), and spaces ().
 - The name cannot exceed 256 characters.

- 2 [Optional] In the **Description** field, add any useful information about this service.

- 3 In the **Agent** field, select the agent to use in this service.

Note: This window's fields and property options change based on the agent. Once an agent is selected, only the appropriate fields display.

- 4 Adjust the default service properties as needed.

The sections following this procedure describe the general properties that are available for each service.

Services

Using the main Services window

5 Click **Add** or **OK** to return to the main Services window.

6 Save your changes.

This service is now available for use in a rule.

Create and modify service groups

A service group is a collection of services that have similar security requirements. When your policy requires several services to have identical rules, grouping these services simplifies your policy by reducing the total number of rules. Also, it allows you to change the rule once and update how your organization uses several services, instead of changing each rule individually. The group can contain proxy services, or packet filter services, or servers, but a group cannot contain a mixture of service types.

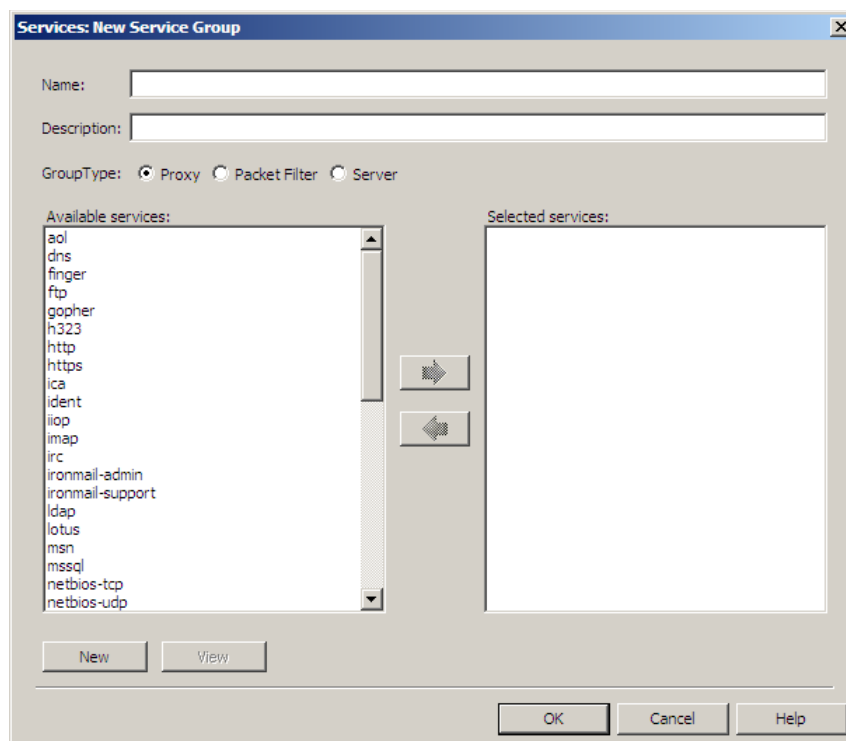
Use the Service Group window to create new service groups or modify existing service groups: select **Policy > Rule Elements > Services**. The Services window appears.

- To create a new service group, click **New Service Group**.
- To modify a service group, select a service group from the list and click **Modify**. (Read-only administrators can click **View** to view a service group.)

Service groups have a folder icon, and are listed as groups in the Agent column.

The Service Group window appears.

Figure 84 Services: New/Modify Service Group window



1 If creating a new service group, enter a name for the service group in the **Name** field.


Note: To rename an existing service group, use Rename on the main Services window.

2 [Optional] In the Description field, enter any information about the service group that may be helpful.


3 At Group Type, select the type of service to appear in the Available services list: **Proxy**, **Packet Filter**, or **Server**.

- To add a new service to the Available services list, click **New** and enter properties in the New Service window.
- To see the properties of an existing service, select a service from the Available services list and click **View**.

4 Add or remove services from the service group:

- To add a service to the service group, select a service in the **Available services** list, and then click the  arrow button.

Select multiple consecutive entries by pressing the **Shift** key as you select the entries. To select multiple non-consecutive entries, press the **Ctrl** key as you select the desired entries.

- To remove a service from the service group, select the service in the **Selected services** list, and then click the  arrow button.

5 When you are done creating or modifying the service group, click **OK**.

6 Save your changes.

The service group is now available for use in a rule.

Configuring proxy agents and services


The following topics are covered in this section:

- [About proxy agents and services](#)
- [Configuring proxy agent properties](#)
- [Configuring proxy service properties](#)
- [Selecting the appropriate proxy agent](#)

About proxy agents and services

A proxy agent is a program that controls communication between clients on one side of a firewall and servers on the other side. The client and server do not communicate directly. Instead, the client and server both “talk” to the proxy agent running on the firewall, which forwards the data back and forth.

Figure 85 Using a proxy agent



The firewall increases a proxy connection's security by receiving each packet, rebuilding it, and then sending it on its way. The traffic's source, or initiator, sends out a request that is routed through the firewall. It inspects the packet, making sure the security policy allows the request. Next the firewall checks if any advanced checks, such as IPS or application defense inspection, are required. Once the firewall is finished handling the request, it rebuilds the packet and sends it to its destination. The firewall also keeps track of what requests were allowed and permits the appropriate responses.

The proxy agents are used to create proxy services. By default, proxy services are disabled. When you use a proxy service in an enabled rule, the firewall automatically enables that service in the corresponding source burb or burbs.

Network applications are typically accessed using one of two lower-level communication protocols: TCP or UDP. TCP is a connection-based protocol that guarantees data is delivered in the same order as sent and ensures address and data integrity. UDP is a connectionless service that delivers data with minimum overhead.

The firewall provides predefined TCP-based proxy services for a variety of Internet services including HTTP, Telnet, FTP, and many others. The firewall also supports proxy services for routing UDP transmissions for applications based on protocols such as SNMP and NTP. Many of these predefined services are based on application-aware proxy agents that can reject packets that do not comply with the protocol's standards. This greatly increases the security and integrity of traffic passed by these proxies. When possible, use the application-aware proxy agents to pass traffic.

The following proxy agents are application-aware: DNS, FTP, H323, HTTP, HTTPS, IIOIP, MS-SQL, Oracle, Ping, RealMedia, RSH, SMTP, SIP, SNMP, SOCKS, SSH, SUN RPC, T120, and Telnet.

Any proxy services that use the Generic Proxy agent are not application-aware. If you must use a service based on the Generic Proxy, increase security for these protocols by restricting the allowed ports and limiting timeout values.

See [Table 20 on page 160](#) for a complete list of proxy services and their descriptions.

Passing traffic transparently and non-transparently

On the Sidewinder, FTP, HTTP, HTTPS, Oracle, and Telnet proxy agents can be configured to be *transparent* or *non-transparent*. For transparent connections, the client is unaware of the firewall. The firewall is implicitly included in the path based on routing. For non-transparent, the client is aware of the firewall and explicitly connects to the firewall. The connection type is determined on the client's side (browser settings or user inputting the firewall's IP address). Proxy services can be configured to allow only transparent connections, only non-transparent connections, or both, depending on which option is indicated in the service's Service Property area.

When using transparent settings, the user appears to connect directly to the desired network's server without connecting to the firewall first. For example, to initiate an outbound Telnet session using a transparent Telnet proxy service, a user would issue the following command from his or her workstation and then connect directly to the external Telnet server:

```
telnet destination
```

With a non-transparent Telnet proxy service, a user must first Telnet to the firewall and specify a destination for the Telnet session. For example, the following shows how an internal user would initiate a Telnet session to a server in an external network using a non-transparent proxy that requires standard password authentication.

```
>telnet firewall_IP_address
```

(connection message from the firewall appears...)

```
>Enter destination: destination_IP_address
```

(authentication prompt from the firewall appears...)

```
>Username: username
```

```
>Password: password
```

(connection message from the destination Telnet server appears...)

```
>login: username
```

```
>Password: password
```

Non-transparent proxy configurations are typically used in networks that use NAT. For example, you would use a non-transparent service if your end users need to access a non-standard port or if there is no direct route between the client and the intended server.

Note: Certain transparent and non-transparent proxy configurations can require users to authenticate before they are allowed to connect. See [Chapter 5, Authentication](#) for more information.

Allowing non-transparent traffic requires configuring end-users' browsers to point to the firewall. To set up browsers to work with the non-transparent proxy option, there are two basic steps:

- Specify the firewall's fully qualified host name or IP address in the browser's proxy line.
- Specify the port number configured in the proxy service's Properties area.

Consult your browser's documentation for defining an HTTP proxy server.

Understanding Fast Path Sessions

By default, the firewall enables a Fast Path Sessions option that improves system performance by lessening the load placed on the system kernel when passing proxy data through the firewall. These sessions involve allowing the kernel to do a raw data transfer instead of copying the data from the kernel to the proxy agent and back. Performance is improved when the Fast Path Sessions option is enabled for protocols that use many small packets, such as Telnet, and for sessions where the proxy can determine that there is no longer any need for data stream inspection (the data channel of an FTP session, the encrypted data from an SSL session, or most data transferred in generic proxies).

Services

Configuring proxy agents and services

In most cases, the Fast Path Sessions option enhances system performance, and in many of these cases the improvement is significant. For this reason, this option rarely needs to be disabled. However, there are a few rare cases where the Fast Path Sessions option may negatively affect performance. Large data transfers on heavily loaded systems, primarily FTP or HTTP traffic, can overload a system. The firewall will also throttle these connections under very heavy load conditions to prevent them from adversely affecting system performance, such as when LAN speeds on both sides of a connection are extremely fast.

The Fast Path Session option is a service property and is configurable on a service-by-service basis. For information on configuring proxy services, see [Configuring proxy service properties](#).

Configuring proxy agent properties

Proxy agent properties are global, meaning the values are shared among all services using that agent. Global properties are related to the agent and not the service. This means that if you have five services using the same agent, such as the HTTP Proxy agent, those five services must all share the same agent properties. If you change a global property value while editing a service, all services using that agent are updated to use the new value.

If an agent has one or more configurable global properties, a Properties button appears next to the Agent field as shown in [Figure 86](#).

Figure 86 Proxy agent properties

The screenshot shows the 'Services: New Service' dialog box. The 'Agent' dropdown menu is set to 'Generic Proxy', and the 'Properties' button next to it is highlighted with a red rectangular box. Below the agent selection, the 'Service properties' section contains two input fields: 'TCP ports' and 'UDP ports', each followed by a three-dot menu button. Underneath these is a 'Timeouts' section with two input fields: 'TCP idle timeout' (set to 7200) and 'UDP idle timeout' (set to 300), both followed by the word 'seconds'. A 'Restore Defaults' button is located to the right of the timeout fields. At the bottom of the dialog, there is a checked checkbox labeled 'Enable fast path sessions' and three buttons: 'Add', 'Close', and 'Help'.

Note: For the Citrix Proxy agent, the UDP ports are a global property. Therefore, if you change the UDP port on one service, all services using the Citrix Proxy agent will be updated with that value.

The following sections explain which agents have global properties and how the values affect that agent's behavior.

Configure expected connections for proxy agents

Certain proxy agents can be configured to enable multiple instances of the same agent in order to load the traffic across the multiple instances. Multiple instantiation of proxy agents is useful for hardware configurations with multiple CPUs or sites that have experienced problems due to an exceedingly large amount of concurrent connections through one of those proxies.

A single proxy instance for any of these agents can generally handle up to 2000 sessions (a session consists of two connections for most protocols). By default, most proxy agents are configured for 4 proxy instances, or about 8000 sessions. This quantity is more than adequate for most sites. However, if your site is consistently recording concurrent sessions that hover around the 8000 range (or if you have experienced problems because the number of connection attempts is significantly higher) for any of these proxies, you may need to increase an agent's number of expected connections in order to enable additional instances for that proxy agent.


The following proxy agents support multiple instantiation:

- Citrix Proxy agent
- FTP Proxy agent
- Generic Proxy agent
- HTTP Proxy agent
- HTTPS Proxy agent
- MS-SQL Proxy agent
- Oracle Proxy agent
- SMTP (Mail) Proxy agent
- SOCKS Proxy agent
- SSH Proxy agent.

Tip: To monitor the number of concurrent connections for the proxy agents listed above, select the Admin Console's dashboard. Click the link titled **Proxy Connections** in the upper-right portion of the dashboard to see a list of all proxy and server services that are currently running and the current number of connections that exist for each.

When you click the **Properties** button next to any of those agents, the following window appears:

Figure 87 Service properties: Expected Connections



Use this window to specify the total number of connections expected for this agent.

For example, if you change this value while creating a new service based on the FTP Proxy agent, then the value changes for all services based on the FTP Proxy agent. If you have two rules using two FTP-based services and the expected connection total is 8000, those rules are expected to support a combined total of 8000 connections.

The default value for all agents is 8000 connections. You can specify expected connection values from 1000–32000.

Configure unique proxy agents

Some proxy agents have unique configuration options. See [Configuring additional proxy agent properties](#) for more information.

Configuring proxy service properties

Proxy service properties apply only to the service, and are not global like agent properties. For example, you can configure multiple HTTP proxy services that listen on different ports but still use the HTTP proxy agent.

Proxy service properties are located in the Service properties area of the New Service and Modify Service windows as shown in [Figure 88](#).

Figure 88 Proxy service properties

Configure common proxy service properties

Most proxy services include some or all of the following properties:

- **TCP/UDP ports** – Select the port or ports on which this service will accept traffic:
 - Enter the port or port range directly, or click to display a list of protocols and their default ports.
 - Do not specify a port number or range that is currently being used by another proxy agent or server agent running on the firewall in the same burb. Use the **Monitor > Service Status** window to see if a different service is already listening on a given port. See [Chapter 12, Service Status](#) for more information.

Note: If you set up your own proxies or reconfigure established proxies, *do not* use ports 9000–9010. These ports are reserved by the firewall for administration purposes.

- **Timeouts** – Set the length of time, in seconds, that the firewall will wait before closing a connection. Return to an agent's default timeout values at any time by clicking **Restore Defaults**.

In most cases, the defaults should be appropriate.

- **TCP idle timeout** – Set the length of time, in seconds, that the TCP connection can remain idle before it is closed.
- **UDP idle timeout** – Set the length of time, in seconds, that the UDP “session” can remain idle before it is closed.
- **Enable fast path sessions** – Leave this option checked unless you are experiencing performance problems. See [Understanding Fast Path Sessions](#) for more information.

Configure connection type

You can configure the following proxies to accept non-transparent connections:

- FTP
- HTTP
- HTTPS
- Oracle
- Telnet

For **Allowed connection types**, select the transparency this service will allow:

- Select **Transparent** to allow only transparent connections.
- Select **Non-Transparent** to allow only non-transparent connections.

- Select **Both** to allow either connection type.

For the HTTP and HTTPS proxy services, if you select **Non-Transparent** or **Both**, make sure the application defense used in a rule with this service specifies which destination ports are allowed. This setting is located on the application defense's Connection tab. See [Chapter 8, Application Defenses](#) for more information.

Note: The SOCKS proxy accepts non-transparent connections only, so its service properties do not include **Allowed connection types**. However, the application defense must still specify which destination ports are allowed.

See [Passing traffic transparently and non-transparently](#) for more information.

Selecting the appropriate proxy agent

The firewall provides a variety of pre-defined proxy services to control connections to popular Internet services using the standard port numbers (see `/etc/services` or www.iana.org/assignments/port-numbers for a list of commonly recognized protocols). These services can be used to quickly set up typical rules. [Table 20](#) shows an alphabetical listing of the proxy services. Determine if these services are appropriate for your site's security policy.

If you determine that your security policy requires proxy services with other properties, the firewall gives you the flexibility to create new services. Each service can be customized and saved under an easily recognizable name. For example, if you want contractors to have shorter timeouts for their FTP sessions than your regular employees, create two services: *FTP contractors* and *FTP standard*. To create additional proxy services, refer to [Create and modify services](#).

See the following sections for additional notes on certain services:

- The proxy services that work together to provide VoIP services such as Microsoft's NetMeeting application require more advanced configuration to interact correctly with the firewall. See [Using the T.120 and H.323 proxy agents together](#) for instructions.
- If you need information on configuring the Session Initiation Protocol service, see "Creating SIP Application Defenses" on page 244.
- If you need to change how the firewall handles FTP server responses, see [Modifying the FTP proxy agent's accepted server responses](#).

The following table lists the default proxy services. Note the following:

- If you selected Standard Internet services during the initial installation, the proxies listed in **bold** are automatically used in the default rule set. This means they are used in enabled rules.
- Agents for the other proxy services will not listen for or manage traffic until they are used in enabled rules.
- Rows containing application-aware proxy agents are shaded.

Table 20 Pre-defined proxy services

Service name	Agent	Type and port	Description
aol	Generic Proxy	TCP 5190	Allows America Online (AOL) members to run their AOL client software and connect directly to America Online. Note: AOL's instant messenger client (AIM) does not limit itself to this port. You cannot grant nor deny AIM access by using this service in a rule.
dns	DNS Proxy	TCP/UDP 53	Allows DNS query traffic and DNS zone file transfers.
finger	Generic Proxy	TCP 79	Allows the UNIX finger command.
ftp	FTP Proxy	TCP 21	Allows transparent or non-transparent access to FTP (File Transfer Protocol) servers. If you require FTP services over HTTP, configure that using an HTTP proxy service and application defense. The HTTP service must be configured as Non-Transparent or Both .
fwregisterp	Cluster Registration Client	TCP 9010	Allows your firewall to join a High Availability (HA) cluster. Note: This proxy is required for internal communication. Do not modify any of its properties unless instructed to do so by Forcepoint support.
gopher	Generic Proxy	TCP 70	Allows communication between Gopher clients and servers.
h323	H323 Proxy	TCP/UDP 1720	Allows audio and video features for H.323 applications, such as Microsoft's NetMeeting application and Cisco® Call Manager. This protocol is commonly used by VoIP-applications. See Using the T.120 and H.323 proxy agents together for more information.

Table 20 Pre-defined proxy services <Comment>(continued)

Service name	Agent	Type and port	Description
http	HTTP Proxy	TCP 80	Allows transparent and non-transparent connections to web servers via HTTP. To allow FTP over HTTP, the HTTP service must be configured as Non-Transparent or Both . To deny FTP over HTTP when the service is non-transparent, clear the GET and PUT checkboxes on the FTP URL Control tab of the HTTP application defense.
https	HTTPS Proxy	TCP 443	Allows transparent and non-transparent connections to web servers via SSL-encrypted HTTP. This proxy can be configured to handle decryption.
ica	Citrix Proxy	TCP 1494 UDP 1604	Allows remote clients to access applications within a Citrix server farm using the Citrix ICA (Independent Computing Architecture) protocol. Locate these Citrix applications either by configuring the client directly, or by pointing them to a master browser. A <i>master browser</i> is a Citrix server that is configured to be responsible for tracking the ICA functions that are available for clients to access, such as applications or other Citrix servers (known as <i>member browsers</i>). <ul style="list-style-type: none"> If you are using Citrix XML Service, to locate the master browser you will need to configure the port that the Citrix server is configured to use in the HTTP proxy service. For the Citrix Proxy agent, the UDP ports are a global property. Therefore, if you change the UDP port on one service, all services using the Citrix Proxy agent will be updated with that value. For information on using the altaddr feature on your Citrix server farm, refer to your Citrix documentation.
ident	Generic Proxy	TCP 113	Allows the UNIX ident command.
iiop	IIO Proxy	TCP 683	Allows the Internet Inter-ORB Protocol (IIOP), the wire protocol used by CORBA (Common Object Request Broker Architecture) applications to interoperate in a heterogeneous network environment. The IIOP proxy allows the firewall administrator to exercise control over the dialogue between the CORBA applications. Note: For more information on CORBA, refer to www.omg.org .
imap	Generic Proxy	TCP 143	Allows the Internet Message Access Protocol, which is used to access e-mail, commonly from a local server.
irc	Generic Proxy	TCP 6667	Allows chat via the Internet Relay Chat (IRC) protocol.
ironmail-admin	HTTPS Proxy	TCP 10443	Allows traffic between an Ironmail® firewall and its management client.
ironmail-support	Generic Proxy	TCP 20022	Allows traffic between Ironmail software and anti-virus updates.
ldap	Generic Proxy	TCP 389	Allows the Lightweight Directory Access Protocol (LDAP).
lotus	Generic Proxy	TCP 1352	Allows the Lotus Notes applications.
msn	Generic Proxy	TCP 569	Allows Microsoft network members to run their MSN client software and connect directly to MSN through the firewall.
mssql	MS-SQL Proxy	TCP 1433	Allows Microsoft servers and clients to pass SQL traffic.
netbios-tcp	Generic Proxy	TCP 139	Allows the generic NetBIOS TCP proxy, which is also known as the NetBIOS Session Service (NBSS). This proxy generally provides access to files and printers. Commonly used with the netbios-udp service.
netbios-udp	Generic Proxy	UDP 137, 138	Allows the generic NetBIOS UDP proxy, which is also known as the NetBIOS Name Service (NBNS). The proxy generally is used for name service resolution in conjunction with the NetBIOS Session Service. Commonly used with the netbios-tcp service.
news	Generic Proxy	TCP 119	Allows access to Usenet News.
ntp	Generic Proxy	UDP 123	Allows clock synchronization via Network Time Protocol (NTP).
oracle	Oracle Proxy	TCP 1521	Allows SQL traffic between Oracle servers and clients.
ping	Ping Proxy	ICMP (na)	Relays ICMP ECHO (ping) requests and ICMP Echo-REPLY messages through the firewall. Note: Enabling the ping proxy does not allow traceroute through the firewall. In addition to security risks, NAT prevents most sites from getting a return from the external network (Internet) because of non-routable addresses. To run traceroute, follow it to the firewall and then initiate a second traceroute from the firewall itself.

Table 20 Pre-defined proxy services <Comment>(continued)

Service name	Agent	Type and port	Description
pop	Generic Proxy	TCP 110	Allows Post Office Protocol (POP) connections.
printer	Generic Proxy	TCP 515	Allows the UNIX <code>lpr</code> command.
realmedia	RealMedia Proxy	TCP/UDP 7070	Allows RealMedia audio and video data packet connections.
rlogin	RSH Proxy	TCP 513	Allows connections to rlogin servers.
rsh	RSH Proxy	TCP 514	Allows RCP (a remote file copy protocol) and RSH (remote shell login).
rtsp	RTSP Proxy	TCP/UDP 554	Allows the RealMedia Player and QuickTime Multimedia Player protocols.
sip	SIP Proxy	UDP 5060	Allows the Session Initiation Protocol (SIP). This protocol is commonly used by VoIP-applications. See "Creating SIP Application Defenses" on page 244.
smtp	Mail Proxy	TCP 25	Allows Simple Mail Transfer Protocol messages through the firewall.
snmp	SNMP Proxy	UDP 161-162	Supports remote management using the SNMP protocol.
socks	SOCKS Proxy	TCP 1080	Allows the SOCKS5 protocol. The only available connection type is non-transparent. When using a SOCKS service in a rule, make sure the associated application defense's Connection tab specifies which destination ports are allowed.
ssh	SSH Proxy	TCP 22	Allows the UNIX Secure Shell command, which provides secure shell access through the firewall to remote systems. See Configuring the SSH proxy agent .
streamworks	Generic Proxy	TCP 1558	Allows Streamworks streaming audio and video.
sunrpc	SunRPC Proxy	TCP/UDP 111	Relays requests between RPC clients and remote servers.
sybase	Generic Proxy	TCP 4000	Allows the Sybase SQL proxy.
syslog	Generic Proxy	UDP 514	Allows the UNIX syslog protocol.
t120	T120 Proxy	TCP 1503	Allows T.120 applications, such as Microsoft's NetMeeting application. This protocol is commonly used by VoIP-applications. See Using the T.120 and H.323 proxy agents together for more information.
telnet	Telnet Proxy	TCP 23	Allows transparent or non-transparent access to Telnet servers.
wais	Generic Proxy	TCP 210	Allows connections between WAIS client software and a database service called WAIS.
whois	Generic Proxy	TCP 43	Allows the UNIX <code>whois</code> command. <code>whois</code> looks up records in the Network Information Center.
wins	Generic Proxy	UDP 42	Allows Microsoft Windows Network Services.
Xwindows	Generic Proxy	TCP 6000	Allows UNIX-based X Windows sessions to pass through the firewall. For instance, an X Windows process running on one terminal could send screen output through the firewall to another window at a different terminal. While redirecting X Windows is a common practice at larger UNIX sites with X Windows environments, X Windows is <i>not</i> a secure application. Using this proxy strictly for sending X Windows traffic through the firewall is not recommended for most sites. However, if the firewall has been placed between two networks, both of which are within your organization (sometimes called "inter-walling"), the Xscreen0 proxy might not pose serious security hazards. This depends on the nature of the site's two networks.
X500	Generic Proxy	TCP 103	Supports the X500 directory server.

Configuring packet filter agents and services

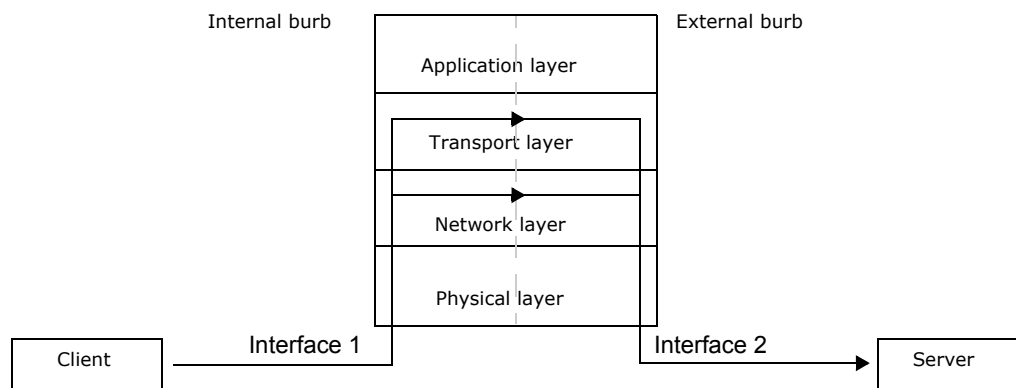
This section covers the following topics:

- [About packet filter agents and services](#)
- [Selecting the appropriate packet filter agent](#)
- [Configuring the TCP/UDP packet filter agent properties](#)
- [Configuring packet filter service properties](#)

About packet filter agents and services

Filter agents are another method for client and servers in different burbs to communicate. They pass traffic at the network layer or the transport layer of the network stack. Filter rules filter incoming packets based on source IP address, destination IP address, and ports. Like proxy rules, filter rules have the option of using network address translation or redirection. Unlike proxy agents, filter agents are not application aware and cannot enforce traffic based on the application protocol. As shown in the following figure, filters inspect traffic at the transport (TCP/UDP) and network (IP) layers. Available agents are the *TCP/UDP Packet Filter agent*, the *ICMP Packet Filter agent*, and the *Other Protocol Packet Filter agent*.

Figure 89 Using a filter



TCP, UDP, and ICMP filters can actively track individual filter sessions using *stateful inspection*. This ensures that only packets valid for a new session or a portion of an existing session are sent on to the final destination.

Filter services are useful in the following situations:

- Traffic that is a protocol other than TCP or UDP, such as AH, ESP, and GRE.
- TCP/UDP protocols where you need a wide port range or maximum performance with minimal security.
- Proprietary traffic that has invalid TCP/UDP headers.

Filter processing can be configured to reject the following source address packets:

- Packets with broadcast source addresses.
- Packets with source addresses on a loopback network that were received on a non-loopback device.

Note: Packets that are rejected for source route information generate a netprobe audit event.

To understand how packet filters work, consider the following topics:

- [How traffic is filtered if stateful packet inspection is enabled](#)
- [How traffic is filtered if stateful packet inspection is disabled](#)
- [Using NAT and redirection for packet filter rules](#)
- [Understanding stateful session failover in an HA cluster](#)

How traffic is filtered if stateful packet inspection is enabled

When the firewall receives TCP, UDP, and ICMP traffic, it starts by checking a filter session record database to determine if an active session record exists for this traffic. A session record indicates that this traffic is in response to a previous successful match to an allow rule. Session records only exist if the matching rule had stateful packet inspection enabled. Stateful packet inspection is only an option for TCP, UDP, and ICMP filter rules.

- **If an active session record exists, the following occurs:**

- a Perform address and port rewriting, if required
- b Perform session processing
- c Forward packet directly to the correct destination interface without any additional processing

- **If no active session record exists, the following occurs:**

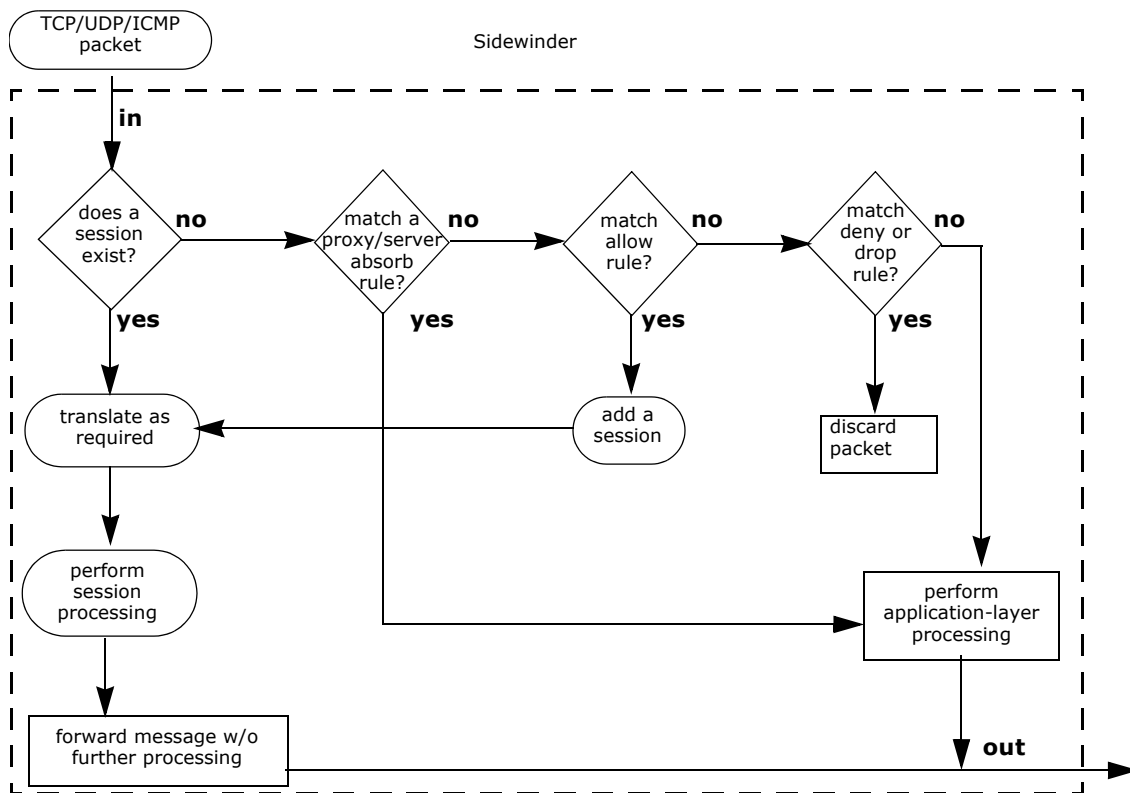
The firewall uses the criteria in [Table 21](#) to check the active filter rules and find a match. The description for how the packet proceeds through the firewall comes after the table. The flowchart in [Figure 90](#) illustrates the complete process.

Table 21 Rule matching criteria with stateful packet inspection enabled

Protocol	Criteria
TCP/UDP	<ul style="list-style-type: none"> • source IP address • destination IP address • ports
ICMP	<ul style="list-style-type: none"> • packet type (echo, message, timestamp) • source IP address • destination IP address

- If a matching allow rule does exist, the following occurs:
 - a Add a session record to the session record database.
 - b Perform Network Address Translation (NAT) if required.
 - c Session processing occurs.
 - d Forward packet directly to the correct destination interface without any additional processing by the firewall.
- If a matching deny rule exists, an RST packet is sent to close the connection. If a drop rule exists, the packet is discarded without further processing.
- If a matching proxy or server rule exists, the packet is sent directly to application-layer processing.
- If no matching filter rule exists, the packet is generally denied. Exceptions:
 - If the packet arrived on a burb that is configured to hide port unreachables, the packet is dropped instead of denied.
 - If a proxy is listening on the packet's port, the proxy handles the packet according to its protocol standards.

Figure 90 Filtering on packets with rules that have stateful packet inspection enabled



How traffic is filtered if stateful packet inspection is disabled

When the firewall receives traffic, it checks the active filter rules for a matching rule. If a rule does not have stateful packet inspection enabled, the firewall checks the criteria in [Table 22](#) to find a match.

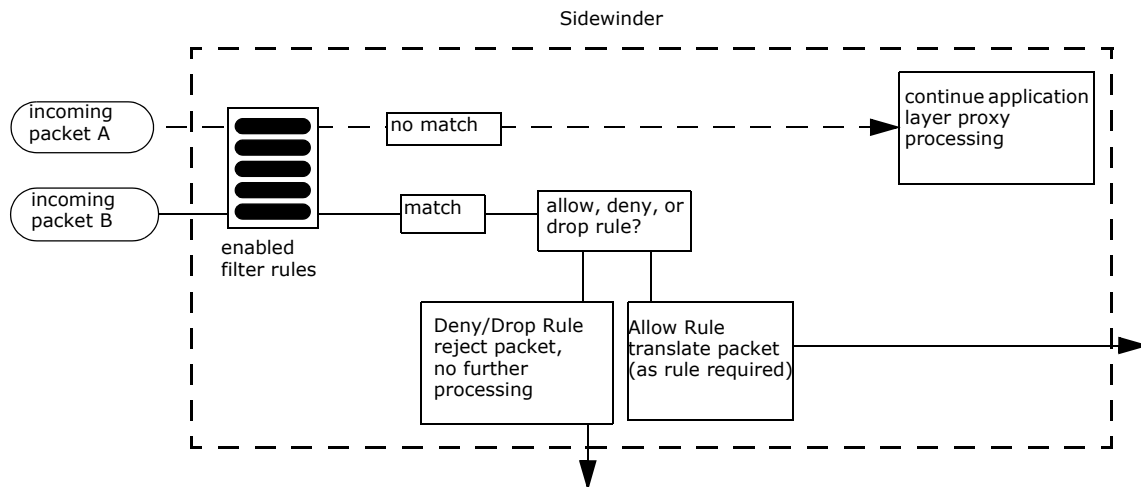
Table 22 Rule matching criteria without stateful packet inspection enabled

Protocol	Criteria
TCP/UDP	<ul style="list-style-type: none"> • source IP address • destination IP address • ports
ICMP	<ul style="list-style-type: none"> • source IP address • destination IP address
Other	<ul style="list-style-type: none"> • source IP address • destination IP address

Using these criteria, the firewall determines if the packet matches any of the active allow or deny/drop rules. The firewall then does one of the following:

- If a rule match is found, the packet source or destination address are translated according to the translation information that is configured for that rule. The packet is then forwarded on for any further firewall processing. The flowchart in [Figure 91](#) illustrates this process.
- If there are no matching rules in the filter database, the firewall sends the packet onto application-layer processing.

Figure 91 Filtering packets when stateful inspection is disabled



Using NAT and redirection for packet filter rules

In general, NAT and redirection are configured the same in filter rules as they are in proxy rules. However, there are some exceptions, particularly in how ports are handled. See the following sections for details.

Limitations of NAT and redirection for filter services

Note the following limitations when setting up rules involving address rewriting for TCP/UDP/ICMP protocols.

- NAT and redirection are not allowed for bi-directional filter rules with stateful packet inspection enabled.
- If stateful inspection is disabled and you want to rewrite an address, the rewritten address must have significant bits value of 32. For example, on an inbound rule the redirect address must be an IP address or hostname network object.

Reserving the port range to use when rewriting source ports

When an outbound packet reaches the firewall and matches a filter rule with NAT configured, the source port and source address will be rewritten and the packet will then be forwarded to its destination.

To facilitate this process, the firewall reserves a range of ports that are to be used exclusively for rewriting source ports. The OS does not allow any processes to bind to a port in this range; configuring proxy services to use ports in this port range will not work.

The default range is set to 9120–9995. If you need to use a port in this range for a different purpose, such as for a new Generic Proxy service, you can adjust the range by doing the following:

- 1 From a command line, run `netstat -an` to view the current port usage. Verify that none of the ports in your selected range are in use.
- 2 Adjust the reserved port range accordingly by editing the **Reserved port range** field in the Global Properties for TCP/UDP Packet Filters window. See [Configure session maximums, port ranges, and intra-burb forwarding for the TCP/UDP Packet Filter agent](#).

Rewriting the address but reserving a packet's specified source port

The firewall enables you to rewrite the source address but maintain the packet's source port. This capability is typically only used when connecting to an application that requires the source port to be a specific value. In some cases, the application requires the source port to be the same value as the port on which the application is listening. This capability is implemented by configuring NAT with **Preserve source port** selected.

The following bullets explain the difference between translating and preserving the source port:

- **Source port is translated** – Each connection uses the same IP address but gets its source port from the reserved port range. The total number of connections can be limited by the number of ports reserved in the Global Properties for TCP/UDP Packet Filters window.
- **Source port is preserved** – Each connection uses the original client source port, but gets its translated IP address one of two ways:
 - If the port range included ports above 1023, this address must be an alias; it cannot be a native IP address. If the port range is below 1024, the address can be a native or *localhost*.
 - From a pool of IP addresses. This requires that there be one or more alias addresses defined for the destination burb's interface and that the NAT field be set to include those addresses. The NAT field can be set to a single IP address or a subnet that includes the alias addresses. The total number of connections is therefore dependent on the number of alias addresses defined for that interface.

Caution: To use this feature with ports above 1023, you must have at least one alias configured for the destination burb's interface or traffic will not pass.

This configuration only applies to uni-directional (source > destination) filter rules with stateful inspection enabled.

By specifying one or more IP aliases, you can have multiple connections because each connection uses the same port number but a different IP address. [Figure 92](#) and [Figure 93](#) illustrate the differences in the two implementations.

Figure 92 NAT with a translated source port

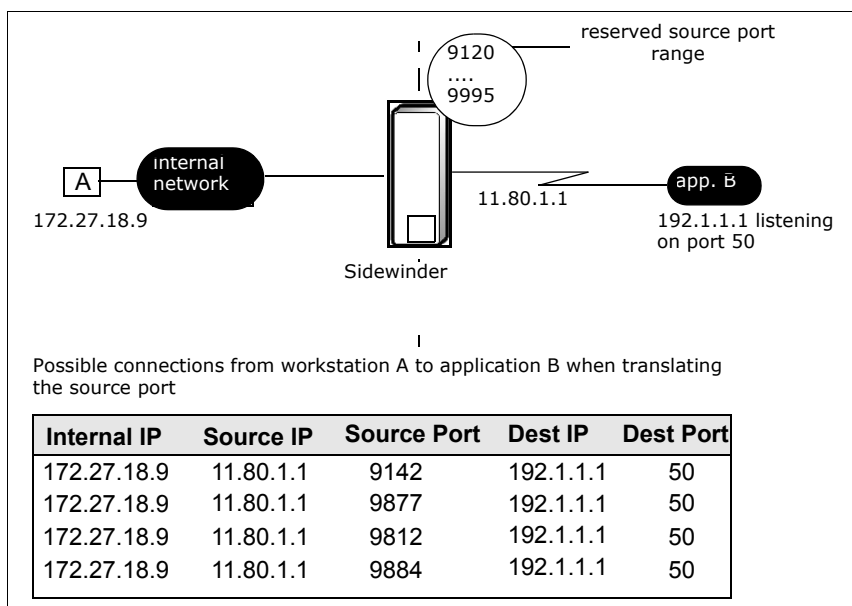



Figure 93 NAT with a preserved source port



Understanding stateful session failover in an HA cluster

When filter session sharing is configured for an HA cluster, the processing firewall sends out multicast messages over the heartbeat interface to notify the other nodes (such as the secondary or standby) of packet filter session activity (such as a new session, closed session, or change in session state). Each time a node receives a message, it updates its local session table accordingly. All sessions received from the primary will have a status of *shared* on the secondary/standby. When HA causes a secondary/standby to take over as the acting primary, the shared sessions on the acting primary become available. When a packet is received for a session, it will be validated against the rules of the processing node. The processing node will then begin sending multicast state-change messages.

Selecting the appropriate packet filter agent

You can create packet filter services based on the following packet filter agents:

- **TCP/UDP Packet Filter** – Use this agent for TCP or UDP traffic on any port(s).
 - Note:** The TCP/UDP packet filter is the only packet filter agent that has configurable agent properties. See [Configuring the TCP/UDP packet filter agent properties](#).
- **FTP Packet Filter** – Use this agent for File Transfer Protocol (FTP) traffic.
 - This agent supports both active and passive FTP by monitoring the control connection and dynamically opening a port for the data connection.
 - If you want to allow FTP over IPv6, you must use this agent. The FTP proxy agent does not currently support IPv6.
 - For more tuning options, see KB article [KB9165](#).
- **ICMP Packet Filter** – Use this agent for Internet Control Message Protocol (ICMP) traffic.
- **Other Protocol Packet Filter** – Use this agent for traffic that is not based on the TCP or UDP protocols. The service provides a list of the Internet protocols you can choose from (see `/etc/protocols` or www.iana.org/assignments/protocol-numbers for a list of commonly recognized protocols).

Note: If you select the **sw-all** protocol, the packet filter service will match traffic for all protocols.

Configuring the TCP/UDP packet filter agent properties


The TCP/UDP packet filter agent has properties that are global, meaning the values are shared among all services using that agent. Its agent properties include the maximum number of TCP sessions and UDP sessions and the reserved port range. Global properties are related to the agent and not the service. This means that if you have five services using the TCP/UDP packet filter agent, those five services must all share the same reserved port range. If you change a global property value while editing a service, all services using that agent are updated to use the new value.

Services

Configuring packet filter agents and services

If an agent has one or more configurable global properties, a **Properties** button appears next to the Agent field as shown in Figure 94.

Figure 94 Packet filter agent properties

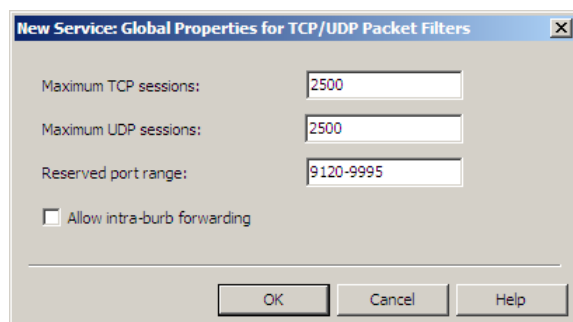


The following section explains the TCP/UDP packet filter agent's global properties and how the values affect its behavior.

Configure session maximums, port ranges, and intra-burb forwarding for the TCP/UDP Packet Filter agent

Click the **Properties** button next to the TCP/UDP Packet Filter agent. The Global Properties for TCP/UDP Packet Filters window appears:

Figure 95 Service properties:



Use this window to set the global properties for the TCP/UDP Packet Filter agent.

- 1 In the **Maximum TCP sessions** field, specify the maximum number of TCP sessions allowed to use the TCP/UDP Packet Filter agent at one time. Valid values are 0–1000000.
- 2 In the **Maximum UDP sessions** field, specify the maximum number of UDP sessions allowed to use the TCP/UDP Packet Filter agent at one time. Valid values are 0–1000000.
- 3 In the **Reserved port range** field, specify the port range that the TCP/UDP Packet Filter agent will reserve for its own use. Valid values are 1024–65533. The default is 9120-9995.
- 4 [Optional] If you want to forward traffic between network interfaces located within the same burb, select **Allow intra-burb forwarding**.

To enforce intra-burb forwarding, create a rule that:

Services

Configuring packet filter agents and services

- has the same source and destination burb.
- uses the TCP/UDP Packet Filter agent as its service.

These values are now set for all services using the TCP/UDP Packet Filter agent.

Configuring packet filter service properties

Packet filter service properties apply only to the service, and are not global like agent properties. For example, you can configure multiple packet filter services for different ports based on the TCP/UDP Packet Filter agent.

Note: Packet filter agents have differing service properties. For example, the TCP/UDP packet filter has UDP-related options while the FTP packet filter does not.

Packet filter service properties are located in the Service properties area of the New Service and Modify Service windows as shown in [Figure 96](#).

Figure 96 Packet filter service properties


The screenshot shows the 'Services: New Service' dialog box. The 'Agent' dropdown is set to 'TCP/UDP Packet Filter'. The 'Service properties' section is highlighted with a red border. It contains the following fields and options:

- TCP ports:** [] ...
- UDP ports:** [] ...
- Stateful packet inspection:**
 - ☒ Enable stateful packet inspection
 - ☒ Enable stateful session failover
 - ☒ Reset TCP connections after connection timeout
 - TCP connection timeout: 15 seconds
 - TCP idle timeout: 7200 seconds
 - UDP idle timeout: 300 seconds
- Advanced Options:**
 - ☐ Require UDP checksums
 - ☐ Restrict source port
 - TCP: []
 - UDP: []
- ☐ Bi-directional

At the bottom of the dialog are 'Add', 'Close', and 'Help' buttons.

Configure packet filter service properties

While each packet filter service has configurable properties, the available properties may differ depending on the type of agent the service is based on. The packet filter service properties are:

- **TCP/UDP ports** – Select the port or port ranges on which this service will accept traffic. Click  to view or search a port list. If you know which port you want to use, enter that port number directly in the field.

Note: Do *not* use ports 9000–9010. These ports are reserved by the firewall for administration purposes.

- **Stateful packet inspection** – Select the **Enable stateful packet inspection** check box. This option must be selected in order to configure the other fields in this area. This option is enabled by default.

To disable stateful packet inspection, clear the **Enable stateful packet inspection** check box.

When enabled, the configurable fields are:

- **Enable stateful session failover:** Select this option to have existing filter sessions transferred to an HA cluster's secondary node during a failover event. This option is enabled by default.

Tip: You may want to disable this option for short-lived connections.

For more information on stateful session sharing, see [Understanding stateful session failover in an HA cluster](#).

- **Reset TCP connections after connection timeout** – When the connection times out, a TCP Reset packet is sent to the client and server.
- **Timeouts** – Set the length of time, in seconds, that the firewall will wait before closing a connection.
 - **TCP connection timeout** – Set the length of time, in seconds, that is allowed for the TCP connection to establish. Valid values are 1–65535.
 - **TCP idle timeout** – Set the length of time, in seconds, that the TCP connection can remain idle before it is closed. Valid values are 0–2147483647.
 - **UDP idle timeout** – Set the length of time, in seconds, that the UDP session can remain idle before it is closed. Valid values are 0–2147483647.
 - **(ICMP Packet Filter only) Response timeout** – Set the length of time, in seconds, that a session will await responses after the final request. Valid values are 1–100000.
- **Require UDP checksum** – Requires the UDP packet to contain a checksum. If this option is enabled and a packet does not contain a UDP checksum, the packet is dropped.
- **Restrict source port** – Specify the port or range of ports (inclusive) from which connections are allowed to be initiated. Note the following:
 - Valid values are 1–65535.
 - To specify “any port,” leave the field blank.
- **Bi-directional** – Allows traffic or session to be initiated from either source or destination addresses. Use this only if your source port and destination port are the same.
- **(ICMP Packet Filter only) Message type** – Select the ICMP message types that you want this service to filter by checking the check box next to each desired message type. Available options are:
 - **echo** – Selecting this matches echo requests and responses used by ping for IPv4 addresses.
 - **info** – Selecting this matches ICMP information requests and responses for IPv4 addresses.
 - **timestamp** – Selecting this matches timestamp requests and responses.
 - **ipv6_echo** – Selecting this matches echo requests and responses used by ping for IPv6 addresses.
 - **ipv6_info** – Selecting this matches ICMP information requests and responses for IPv6 addresses.

Note: ICMP control and error messages generated by TCP/UDP traffic are managed using TCP/UDP rules, as opposed to ICMP rules. For example, if you want to pass “host unreachable” error messages for a specific rule’s undelivered TCP packets through the firewall, you would configure this option on the Packet Filter application defenses instead of using the ICMP service.

- **(Other Protocol Packet Filter only) Protocol** – Expand the drop-down list and select the protocol to use for this service.

Configuring server agents


This section covers the following topics:

- [About server agents](#)
- [Configuring server agent properties](#)
- [Selecting the appropriate server](#)

About server agents

On the Sidewinder, servers provide a variety of system functions, but generally do not pass traffic between burbs. Rules that allow access to a firewall server typically have the same source and destination burbs, as shown in the following figure.

Figure 97 Using a server



Common services include the Admin Console server (used for GUI management), the SSH server (used for command line management), and sendmail. Unlike proxies and filters, you cannot create new server services.

Note: By default, server services are disabled. When you use a server service in an enabled rule, the firewall automatically enables that service in the corresponding source burb or burbs.

Configuring server agent properties

To begin working with server services, select **Policy > Rule Elements > Services**. To access a server service, double-click it, or select it and then click **Modify**. You can change the service's description, its service properties, or its global agent properties.

Unlike proxy and filter services, the firewall has a pre-defined list of services that cannot be deleted or added to. You can modify some of these servers' properties. For a list of all servers, see [Table 24](#).

The Login Console server has no configurable properties.

The following servers use the basic service properties that can be adjusted to suit your policy: `changePW`, `entrelyd`, `fwregisted`, and `telnetd`. They do not have any other configurable properties.

The servers listed in the following table have important configurable properties. The right-hand column gives an overview of what properties are configurable and lists what section to see for detailed configuration information:

Table 23 Servers with advanced configuration properties

Server	Configurable properties
Admin Console	Change the SSL certificate used by the Admin Console client to authenticate to the Admin Console server and the login banner that greets firewall administrators after they log in. See Configuring the Admin Console server .
DHCP Relay	Configure what DHCP servers DHCP requests are forwarded to and other advanced properties. See Configure the DHCP Relay agent .
bgpd	Configure the firewall to participate in Border Gateway Protocol (BGP) dynamic routing. See Configuring BGP (bgpd) .
ISAKMP	Configure the audit level for this server's traffic, negotiation properties, and extended authentication parameters. See Managing the ISAKMP server .
ospfd	Configure the firewall to participate in Open Shortest Path First (OSPF) dynamic routing. See Configuring OSPF (ospfd) .
ospf6d	Configure the firewall to participate in Open Shortest Path First IPv6 (OSPF IPv6) dynamic routing. See OSPF IPv6 on Sidewinder .
pimd	Configure the firewall to participate in Protocol Independent Multicast - Sparse Mode (PIM-SM) multicast routing. See Configuring PIM-SM (pimd) .
ripd	Configure the firewall to participate in Routing Information Protocol (RIP) dynamic routing in a specific burb. See Configuring RIP (ripd) .
ripd-unbound	Configure the firewall to participate in Routing Information Protocol (RIP) dynamic routing in all burbs. See Configuring RIP (ripd) .
sendmail	Edit the sendmail configuration files. You can also run the Reconfigure Mail tool from this service's Property window. See Setting up and reconfiguring mail .
sfadmin	Change the password that is sent by the firewall to the SmartFilter server. You must make the same change to the SmartFilter Admin Console's Plugin Definition Admin Password. Note: SmartFilter is not actively supported in Sidewinder 7.x.
snmpd	Configure communities, trap destinations, and whether or not to send the authentication failure trap. See Setting up the SNMP agent on Sidewinder .
sshd	Generate new host keys, and generate and export new client keys. See Administering Sidewinder using Secure Shell .
ssod	Configure the login and logout page banners displayed to users when they use the web login page to start or end a single sign-on session. See Setting up Passport authentication .


Configure the DHCP Relay agent

Use the DHCP Relay agent to configure your firewall to allow clients to obtain IP addresses from a DHCP server in a different burb.

To configure the DHCP Relay agent:

- 1 Select **Policy > Rule Elements > Services**. The Services window appears.
- 2 From the list of services, select **DHCP Relay** and click **Modify**. The Modify Service window appears.
- 3 Click **Properties**. The DHCP Relay Properties window appears.

Figure 98 DHCP Relay Properties: DHCP Servers tab



- 4 In the DHCP Servers tab, add the servers to which DHCP requests should be forwarded:

- a Click **New**. The New Server Address window appears.
- b Enter the server's address information by doing one of the following:
 - Select **IP address** and type the IP address of the server.
 - Select **Hostname** and type the host name of the server.

Note: If you add a server using its host name, the firewall must be able to resolve the host name to an IP address via DNS.

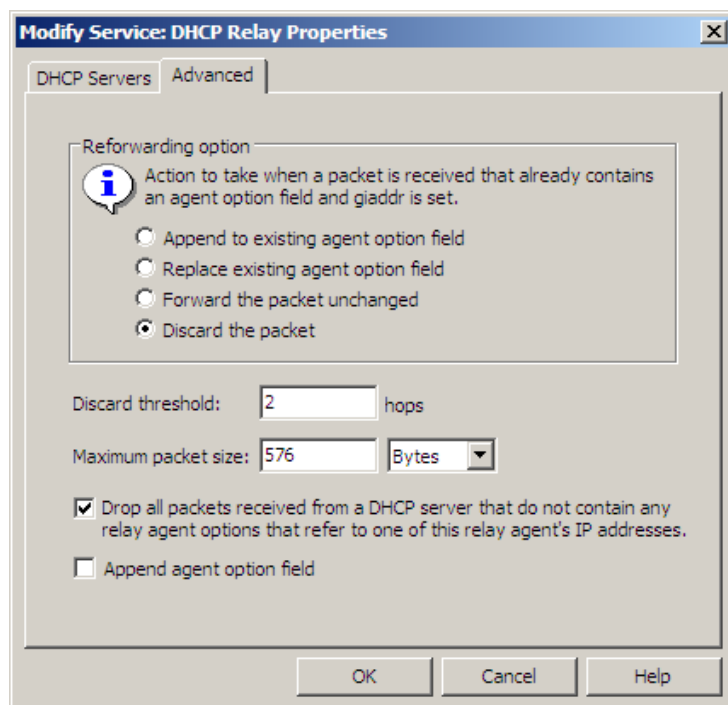
- c Click **Add**. You return to the DHCP Servers tab.

To modify or delete an existing server entry, select the server and click **Modify** or **Delete**.

Note: The DHCP Relay agent forwards DHCP requests to each DHCP server you define on the DHCP Servers tab. If multiple servers respond to a DHCP request, the DHCP Relay agent forwards the first response it receives to the client and ignores the others.

- 5 Click the **Advanced** tab to configure additional DHCP Relay options:

Figure 99 DHCP Relay Properties: Advanced tab



- a In the **Reforwarding option** area, select how DHCP request packets that have already been forwarded by another DHCP relay are processed:
- **Append to existing agent option field** – Appends the firewall's DHCP relay agent option data to DHCP requests and then forwards the requests to the defined DHCP server(s).
 - **Replace existing agent option field** – Replaces the agent option data added to DHCP requests by other DHCP relays with the firewall's DHCP relay information and then forwards the requests to the defined DHCP server(s).
 - **Forward the packet unchanged** – Forwards DHCP requests to the defined server(s) without modifying the agent option data added by other DHCP relays.
 - **Discard the packet** – Discards any DHCP requests that have been forwarded by other DHCP relays.
- b In the **Discard threshold** field, type the maximum number of DHCP relays that DHCP request packets can pass through before being dropped by the firewall. Allowed values are 1–255 hops.
- c In the **Maximum packet size** field, type the maximum size of DHCP request packets that the DHCP Relay agent can create after appending its agent option information. Allowed values are 576–9000 bytes.
- d Select **Drop all packets received from a DHCP server that do not contain any relay agent options that refer to one of this relay agent's IP addresses** to drop packets from DHCP servers that do not correspond to requests forwarded by this firewall.
- e Select **Append agent option field** to append additional DHCP Relay agent information to the agent option field of DHCP request packets, including the printable name of the firewall network interface on which the request was received.
- f Click **OK** until you return to the Services window and save your changes.

6 Create a rule to accept DHCP requests from clients. Include these selections:

- **Service** – Select **DHCP Relay (DHCP Relay Agent)** from the drop-down list.
- **Source Burb** – Select the burb where the clients attempting to obtain IP addresses via DHCP are located.
- **Source Endpoint** – Verify that **<Any>** is selected.
- **Destination Burb** – Select the burb where the clients attempting to obtain IP addresses via DHCP are located.
- **Destination Endpoint** – Select or create an IP Address object with a value of **255.255.255.255**.

Note: The **Source Burb** and **Destination Burb** selections should be the same.

7 Create a rule to allow the DHCP server(s) to respond to DHCP requests. Include these selections:

- **Name** – Type a name for this rule.
- **Service** – Select **DHCP Relay (DHCP Relay Agent)** from the drop-down list.
- **Source Burb** – Select the burb or burbs where the DHCP server(s) are located.
- **Source Endpoint** – Restrict the source as desired, as long as the desired DHCP server(s) is included.
- **Destination Burb** – Select the burb where the clients attempting to obtain IP addresses via DHCP are located.
- **Destination Endpoint** – Select or create an IP Address object with the IP address of the firewall in the burb where the clients attempting to obtain IP addresses via DHCP are located.

8 [Conditional] Create a rule to allow clients to renew their DHCP leases from the DHCP server(s).

Note: Some DHCP clients, such as Windows XP computers, attempt to renew their DHCP address leases by directly connecting to the DHCP server that assigned the address to them. If your network environment requires that this be allowed, complete this step.

a Click **New**. The New Rule window appears.

b Complete the fields as follows:

- **Name** – Type a name for this rule.
- **Service** – Select **DHCP Relay (DHCP Relay Agent)** from the drop-down list.
- **Source Burb** – Select the burb where the clients attempting to renew IP leases via DHCP are located.
- **Source Endpoint** – Restrict the source as desired, as long as the desired DHCP clients are included.
- **Destination Burb** – Select the burb or burbs where the DHCP server(s) are located.
- **Destination Endpoint** – Restrict the destination as desired, as long as the desired DHCP server(s) are included.

c Click **OK** and save your changes.

9 Make sure that the rules you created in [Step 6](#) through [Step 8](#) are enabled and above the Deny All rule.

10 Save your changes.

Selecting the appropriate server

Servers can be classified as belonging to one of the following categories:

- Management – Used for management and administration of the Sidewinder.
- Service – Provides access to a networked service.
- Routing – Provides routing services on the firewall.
- VPN – Used in VPN connections.
- Sidewinder-specific – An inter- or intra-firewall server used in firewall clustering or centralized management.

See the following table for a list of the functions provided by each server.

Table 24 Available servers

Service	Function	Description
Admin Console (Admin Console)	Management	Used when administrators log into the firewall using the Sidewinder Admin Console.
bgpd (BGP Server)	Routing	Used in routing with the Border Gateway Protocol (BGP). See BGP on Sidewinder .
ccmd ccms (Control Center Management Server)	Sidewinder-specific	Used in registration and communication among the Control Center and managed Sidewinders. <ul style="list-style-type: none"> • ccmd is used to send data from the Control Center firewall to the Sidewinder. • ccms is used to send data from the Sidewinder to the Control Center firewall.
changepw (Change Password Server)	Service	Allows external users to use a browser to change their Sidewinder, SafeWord PremierAccess, or LDAP login password. See Setting up users to change their own passwords .
entrelayd (Enterprise Relay Server)	Sidewinder-specific	Used for services that need to communicate with each other in multi-firewall configurations.
fwregisterd (Cluster Registration Server)	Sidewinder-specific	Used for registration and communication among firewalls in High Availability (HA) pairs.
isakmp (ISAKMP Server)	VPN	Used to generate and exchange keys for VPN sessions. See Creating VPN policy .
login (Login Console)	Management	Used when administrators log in at a console attached to the Sidewinder.
ospfd (OSPF Server)	Routing	Used in routing with the Open Shortest Path First (OSPF) protocol. See OSPF on Sidewinder .
ospf6d (OSPF IPv6 Server)	Routing	Used in routing with the Open Shortest Path First IPv6 (OSPF IPv6) protocol. See OSPF IPv6 on Sidewinder .
pimd (XORP Server)	Routing	Used in routing with the Protocol Independent Multicast - Sparse Mode (PIM-SM) protocol. See PIM-SM on Sidewinder .
ripd (RIP Routing Server)	Routing	Used in routing with the Routing Information Protocol (RIP). See Configuring RIP (ripd) .
ripd-unbound (RIP Unbound Server)	Routing	Used in routing with the Routing Information Protocol (RIP). See Configuring RIP (ripd) .
sendmail (Sendmail Server)	Service	Used when running hosted sendmail on a firewall. See Editing sendmail files on Sidewinder .
sfadmin (SmartFilter Admin Console)	Service	Used when communicating with the SmartFilter Administration Console. Note: SmartFilter is not actively supported in Sidewinder 7.x.
sfredirect (SmartFilter Redirect Server)	Service	Used when responding to denied or coached web requests. Note: SmartFilter is not actively supported in Sidewinder 7.x.
snmpd (SNMP Agent)	Service	Used in communication with SNMP management stations.
sshd (SSH Server)	Management	Used when administrators log into the firewall using an SSH client. Often used in troubleshooting and when editing files. See Administering Sidewinder using Secure Shell . The default policy contains a disabled rule allowing internal access to this SSH server. To enable this rule, select Policy > Rules , expand the Administration rule group, and then enable the Secure Shell Server rule. For added security, modify the rule to make it more restrictive.

Table 24 Available servers <Comment>(continued)

Service	Function	Description
ssod (Passport Authenticator)	Service	Used in single sign-on, or out-of-band, authentication and is the basis for the Passport authenticator. See Setting up Passport authentication .
telnetd (Telnet Server)	Management	Used when administrators log into the firewall using a Telnet client. Caution: Telnet sessions are passed in the clear and should only be used within a protected network. For security reasons, always try to use the SSH server for command line sessions.

Configuring additional proxy agent properties

Some proxy agents have unique configuration options. This section covers the following proxy agent configuration procedures:

- [Configuring URL translation on the HTTP proxy agent](#)
- [Using the SSH proxy agent](#)
- [Modifying the FTP proxy agent's accepted server responses](#)
- [Configuring the SMTP proxy agent to strip source routing](#)
- [Using the T.120 and H.323 proxy agents together](#)

Configuring URL translation on the HTTP proxy agent

Use URL translation to configure your firewall to redirect inbound HTTP connections based on application layer data, rather than on transport layer data like conventional redirect rules. By examining the HTTP application layer data, the firewall determines which internal web server inbound requests are destined for even if multiple servers share the same external IP address.

Use URL translation if your network environment matches one or more of the following scenarios:

- You have multiple web sites that resolve via DNS to a single IP on your firewall.
- You have a web site(s) that contains resources that are hosted on different physical servers behind your firewall.

If URL translation is enabled on an internet-facing burb, inbound HTTP requests are handled as follows:

- 1 An inbound HTTP request reaches the firewall.

Note: The TCP connection must be destined for an IP address that is assigned to the firewall.


- 2 The firewall examines the HTTP request's application layer data and compares it to the defined URL translation rules to determine which internal web server the request should be sent to.
- 3 [If **Rewrite URL** is enabled] The firewall rewrites the application data in the HTTP request as configured so that it conforms to the requirements of the internal web server.
- 4 Based on the IP address of the destination web server determined in [Step 2](#), a policy rule match is performed.
- 5 If a policy rule is matched, the connection is redirected to the internal web server.

Configuring URL translation

URL Translation rules are configured on the HTTP Proxy Agent Properties window:

- 1 Select **Policy > Rule Elements > Services**. The Services window appears.
- 2 From the list of services, select **http** and click **Modify**. The Modify Service window appears.
- 3 Click **Properties**. The HTTP Proxy Agent Properties window appears.

Figure 100 HTTP Proxy Agent Properties window



This table lists the configured URL translation rules. The URL translation rules are checked in order and the first rule that matches is used. For this reason, more specific rules should be placed higher in the rules list. Use the up and down arrows to change the rule order. To manage the URL translation rules, use the following buttons:

- **New** – Use this button to create a new URL translation rule.
- **Modify** – Use this button to modify an existing URL translation rule. You can also double-click a rule to modify it.
- **Delete** – Use this button to delete the selected URL translation rule.
- **Duplicate** – Use this button to make a copy of the selected URL translation rule.
- **Rename** – Use this button to rename the selected URL translation rule.

When you create URL translation rules, refer to the following guidelines:

- Order your rules so that the most specific rules are placed first.
- Avoid using file names in the **Path Prefix** fields.
- Avoid adding trailing slashes to paths you specify in the **Path Prefix** fields.

Path prefix matches are exact, so a trailing slash can cause unwanted behavior. For example, specifying `/directory_name/` in the **Path Prefix** does not match the request `GET /directory_name` because the trailing slash is missing.

Note: Performing URL translation and conventional redirection for the same firewall IP address is not supported.

To configure inbound HTTP access to an internal web server using URL Translation, click **New**. The New URL Translation Rule window appears.

Figure 101 New URL Translation Rule window

- 1 In the **Name** field, type a descriptive name for this rule.
- 2 [Optional] In the **Description** field, enter any useful information about this rule.
- 3 In the **Client Source** area, choose the burb or burbs where the clients that generate the inbound HTTP requests are located by doing one of the following:
 - Select **Burbs** and then select the appropriate burb or burbs from the list.
 - Select **Burbgroups** and then select the appropriate burb group or groups from the list.
- 4 In the **Original URL** area, configure the HTTP matching parameters by doing one of the following:
 - Select **Matching URL** and type the URL that this rule should match.
To specify a custom port, add the port to the end of the URL. Example: *http://example.net:3128*.
The **Host**, **Ports**, and **Path Prefix** fields are automatically populated based on the URL you enter.
 - Select **Matching URL attributes** and complete the **Host**, **Ports**, and **Path Prefix** fields with the data used to match inbound HTTP requests.
- 5 In the **New Server Destination** area, select or create an IP address object that corresponds to the internal web server that connections matching this rule should be redirected to.

- 6 [Optional] Select **Rewrite URL** if you need to translate the inbound HTTP request so that it matches the host name and path structure of the internal web server.

Note: The new URL information replaces only the original URL information you entered in [Step 6](#). Path information beyond the original URL path prefix in the HTTP request is unaffected.

Do one of the following:

- Select **New URL** and type the URL that should replace the original URL.

To specify a different port, clear the **Maintain original port** check box and add the port to the end of the URL. Example: *http://example.net:3128*.

The **Host**, **Ports**, and **Path Prefix** fields below are automatically populated based on the URL you enter.

- Select **New URL attributes** and complete the **Host**, **Ports**, and **Path Prefix** fields with the data to replace the original URL attributes.

Note: Sidewinder does not modify hyperlinks in HTML files, so web servers that the firewall performs URL translation for should employ relative links whenever possible. The firewall does translate the **Location** header in 3xx redirection server status codes.

- 7 Click **Add**. You return to the HTTP Proxy Agent Properties window.

- 8 Click **OK** and save your changes.

- 9 Select **Policy > Rule Elements > Services** and create a new HTTP proxy service to accept inbound connections that require URL translation. Configure the following fields:

- **Name** – Enter a descriptive name for this service.
- **Agent** – Select **HTTP Proxy**.
- **TCP ports** – Enter all of the ports you specified in [Step 4](#).
- **Allowed connection types** – Select **Non-Transparent**.

- 10 Select **Policy > Rules** and create a policy rule to authorize the inbound connection.

Note: URL translation rules only determine the internal IP address to redirect the inbound HTTP requests to. Policy rules are needed to authorize inbound connections based on the information provided by the URL translation rules.

Make the following selections:

- **Service** – Select the service you created in [Step 9](#).
- **Source Burb** – Select the burb or burbs where the clients are located. This selection should match your selection in [Step 3](#).
- **Source Endpoint** – Verify that **<Any>** is selected or restrict as desired.
- **Destination Burb** – Select the burb where the destination web server is located.
- **Destination Endpoint** – Select the IP address object that corresponds to the destination web server. This object should match your selection in [Step 5](#).
- **Destination Redirect** – Verify that **<None>** is selected.

Make sure that the rule you created in is enabled and above the Deny All rule, and then save your changes.

Using the SSH proxy agent

You can configure the SSH proxy to decrypt SSH traffic, perform content inspection, and then re-encrypt the traffic before sending it to its destination.

To decrypt and re-encrypt the SSH traffic, the proxy acts like a server when communicating with the client, and acts like a client when communicating with the server. Therefore, it must maintain two databases:

- A known hosts database to store SSH server keys
- A database of SSH server keys to present to clients

Both the known hosts database and the server keys are managed on the SSH proxy agent.

To learn how to configure the SSH proxy agent, refer to the following sections:

- [Understanding the SSH known host keys trust relationship](#)
- [Configuring the SSH proxy agent](#)
- [Manage known host keys](#)
- “Creating SSH Application Defenses” on page 247

Understanding the SSH known host keys trust relationship

The SSH protocol relies upon users to decide if the server host keys that are presented to them are valid. Because the firewall acts like a client when it communicates with SSH servers, server host keys are stored in the firewall's SSH known host keys database. To distinguish between server host keys that have been administrator-approved and those that have not, the firewall classifies each host key by trust level. The trust level configured for each SSH known host key represents your level of confidence that the host key belongs to the host (IP address) that it claims to belong to. There are two trust levels:

- **Strong** – SSH host keys are considered strong if they have been imported into the SSH known hosts database by administrators or promoted to strong trust level by administrators.
- **Weak** – SSH host keys are considered weak if they are accepted by users without administrator intervention during the initiation of an SSH session.

When you configure the SSH Application Defense for a new SSH proxy rule, you can decide what SSH host key trust level to require in order to allow the SSH connection to take place. For example:

- Enforce **Strict** key checking policy for rules that allow access to critical network security devices.

Host keys with strong trust level must already exist in the known hosts database for the security devices that the rule allows access to. These host keys must also pass cryptographic checks for authenticity.

- Enforce **Medium** key checking policy for rules that allow access to non-critical hosts.

Host keys with strong or weak trust level are allowed. If a host key is not present in the known hosts database, the client can accept it, which adds the host key to the known hosts database.

- Allow **Relaxed** key checking policy for rules not related to business operations, such as a rule allowing access to an employee's personal computer at home.

Host keys with strong or weak trust level are allowed. If a host key is not present in the known hosts database, the client can accept it, which adds the host key to the known hosts database. If a server's host key has changed, the client can accept it, which replaces the old key in the known hosts database.

By tailoring each rule's key checking policy to the security risk involved, you can ensure that SSH host keys from critical servers receive administrator verification, while less critical SSH servers can be accessed without administrator intervention. For more information, see “Creating SSH Application Defenses” on page 247.

Strong host key scenario


Consider the following scenario in which an SSH client needs to connect to a network security device through the firewall's SSH proxy. The network security device is critical for the integrity of the network, so the administrator chooses to enforce strict key checking policy. As a result, the administrator needs to make sure that there is a strong known host key for the network security device in the firewall's known hosts database. The following configuration steps are necessary to allow the connection to take place:

- 1 Create an SSH Application Defense that enforces **Strict** key checking policy. See “Creating SSH Application Defenses” on page 247 for details.

Note: For the connection to be allowed, a strong host key must already be present in the SSH known host keys database for the network security device.

- 2 Create an SSH proxy rule to allow the SSH client to connect to the network security device.
- 3 Import the network security device’s SSH host key into the firewall’s SSH known host keys database, assigning it a strong trust level.

Figure 102 Example strong SSH known host key scenario



The figure above shows what happens when the SSH client initiates an SSH session to the network security device through the firewall’s SSH proxy agent:

- a The client initiates an SSH connection to the network security device. The firewall, acting like an SSH server, accepts the client’s connection.
- b The firewall sends its SSH host key to the client.
- c The firewall, acting like an SSH client, initiates an SSH connection to the network security device. The network security device accepts the firewall’s connection.
- d The network security device sends the firewall its SSH host key.

The firewall examines the SSH host key from the network security device and allows the connection. Because the administrator imported a strong SSH host key for the network security device into the firewall’s SSH known hosts database, the requirements of strict key checking policy are met.

Weak host key scenario


Consider the following scenario in which an employee wants to connect to their home computer through the firewall's SSH proxy. The employee's home computer is not critical for the integrity of the network, so the administrator chooses to enforce relaxed key checking policy. As a result, the administrator does not need to import or approve the SSH host key that belongs to the employee's home computer. The following configuration steps are necessary to allow the connection to take place:

- 1 Create an SSH Application Defense that enforces **Relaxed** key checking policy. See "Creating SSH Application Defenses" on page 247 for details.

Note: Host keys with strong or weak trust level are allowed. If a host key is not present in the known hosts database, the client can accept it, which adds the host key to the known hosts database. If a server's host key has changed, the client can accept it, which replaces the old key in the known hosts database.

- 2 Create an SSH proxy rule to allow the SSH client to connect to the employee's home computer.

Figure 103 Example weak SSH known host key scenario



The figure above shows what happens when the SSH client initiates an SSH session to the employee's home computer through the firewall's SSH proxy agent:

- a The client initiates an SSH connection to the employee's home computer. The firewall, acting like an SSH server, accepts the client's connection.
- b The firewall sends its SSH host key to the client.
- c The firewall, acting like an SSH client, initiates an SSH connection to the employee's home computer. The employee's home computer accepts the firewall's connection.
- d The employee's home computer sends the firewall its SSH host key.
- e The firewall sends the SSH host key presented by the employee's home computer to the client for approval.


The firewall allows the connection if the user approves the SSH host key presented by the employee's home computer. Since the administrator configured relaxed key checking policy for the SSH Application Defense, the user has the ability to approve any SSH host key.

Configuring the SSH proxy agent

To configure the SSH proxy agent properties:

- 1 In the Admin Console, select **Policy > Rule Elements > Services**.
- 2 In the list of services, select **ssh** and then click **Modify**. The Modify Service window appears.
- 3 Click **Properties**. The SSH Proxy Properties: SSH Known Hosts tab appears.

Figure 104 SSH proxy agent properties



The SSH Proxy Properties window has two tabs:

- **SSH Known Hosts** – Use this tab to manage the database of known host keys.

Note: To configure this tab, you must have an SSH proxy rule configured, enabled, and positioned above the Deny All rule on the **Policy > Rules** window.

- **SSH Server Keys** – Use this tab to manage SSH keys that the proxy presents to SSH clients. See “Managing VPN certificates” on page 633 for more information.

Manage known host keys

Perform the following tasks to manage known host keys:

- Add a known host key by clicking **New** and entering the appropriate information in the pop-up window. See [Creating or modifying an SSH known host key](#).
- Modify a known host key by selecting it in the list and clicking **Modify**.

You can modify the following fields:


- **Trust Level**
- **IP address**
- **Port**
- **Key type**
- **Key value**

Note: You can also change the trust level by selecting a known host key from the list and clicking **Set trust level to Strong** in the toolbar.

- Delete a known host key by selecting it in the list and clicking **Delete**.

Creating or modifying an SSH known host key

Figure 105 New SSH Known Host window



To create or modify an SSH known host key:

- 1 In the **Trust level** drop-down menu, select **Strong** or **Weak**. See [Understanding the SSH known host keys trust relationship](#).
- 2 In the **IP address** field, type the IP address of the host that the new known host key corresponds to.
- 3 If necessary, change the port specified in the **Port** field to match the port that the host's SSH server is listening on.
- 4 From the **Key type** drop-down list, select the appropriate key type.
- 5 Enter the host key data by doing one of the following:
 - Paste the key data in the **Key value** field.
 - Retrieve the key from the remote host by clicking **Retrieve key**.
 - Import the key by clicking **Import from file** and then browsing to the appropriate key file.
- 6 Click **Add**. You return to the SSH Proxy Properties window and the new host key is added to the list of host keys.

Note: When you accept a host key presented by a server while connecting to that server through the SSH proxy, it is added to the SSH Known Hosts list. Accepted keys automatically have a weak trust level.


Modifying the FTP proxy agent's accepted server responses

By default, the firewall restricts which FTP server responses it will accept. Accepted FTP server response codes range from 100 to 599. To alter which codes are accepted or to turn off server response checking, do the following:

Note: Only experienced administrators should edit configuration files.

- 1 In the Admin Console, go to **Maintenance > File Editor**, and then click **Start File Editor** in the right pane. The File Editor window appears.
- 2 From the **File** menu, select **Open**. The Open File window appears.

Figure 106 File Editor: Open File window



- 3 In the **Source** area, select **Firewall File**.
- 4 In the **File** field, type **/secureos/etc/proxy/pftp.conf**, then click **OK**. The pftp.conf file opens in the File Editor.
- 5 If you want to turn off server response checking, find the following line:

```
validate_server_response[yes]
```


and change `[yes]` to `[no]`.

- 6 If you want to limit which FTP server responses the firewall accepts, edit the values in the following lines:

```
min_server_response_code[100]
max_server_response_code[599]
```

Valid values are between 000 and 999.

Figure 107 Example configuration file



- 7 Save your changes and close the File Editor.
- 8 Restart the FTP proxy agent to make your changes active:

- a Select **Monitor > Service Status**. The Service Status window appears.
- b To restart the FTP proxy agent, right-click **ftp** in the Service list and then select **Restart**.

The FTP proxy has now been restarted and is using the updated configuration file.


Configuring the SMTP proxy agent to strip source routing

Source routing allows the sender of a piece of mail to specify the intermediate hosts that the message should be sent to in order to reach its final destination. This feature is not supported by the SMTP proxy because it poses a security risk and has been deprecated by RFC 2821.

By default, the SMTP proxy blocks RCPT and MAIL commands that include mailbox addresses with source routing. To configure the SMTP proxy to remove source routing information from messages before delivering them:

- 1 In the Admin Console, go to **Maintenance > File Editor**, and then click **Start File Editor** in the right pane. The File Editor window appears.
- 2 From the **File** menu, select **Open**. The Open File window appears.

Figure 108 File Editor: Open File window

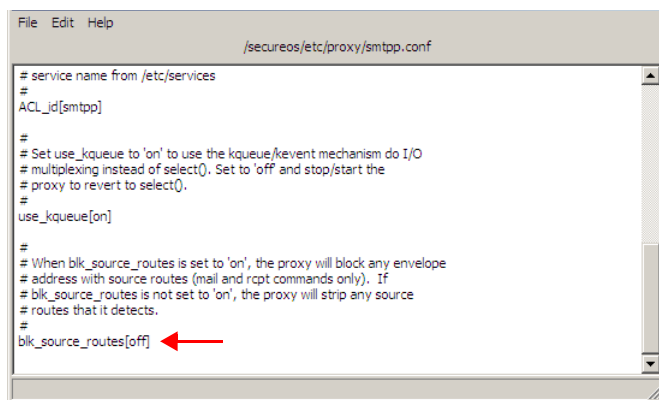


- 3 In the **Source** area, select **Firewall File**.
- 4 In the **File** field, type **/secureos/etc/proxy/smtp.conf**, then click **OK**. The smtp.conf file opens in the File Editor.
- 5 Find the following line:

```
blk_source_routes[on]
```

and change the text in the square brackets from **on** to **off**.

Figure 109 Example configuration change



- 6 Save your changes and close the File Editor.

7 Restart the SMTP proxy agent to make your changes active:

- a Select **Monitor > Service Status**. The Service Status window appears.
- b To restart the SMTP proxy agent, right-click **smtp** in the Service list and then select **Restart**.

The SMTP proxy agent now strips source routing information from mail messages.

Using the T.120 and H.323 proxy agents together

The T.120 and H.323 proxy agents can be configured to work together, allowing you to make use of both the data-sharing and audio/video features of data conferencing products, such as Microsoft NetMeeting, in a single conference. This section provides an overview of each agent and its role in data conferencing. It also provides information on configuring the two agents to work together to enable the complete realm of NetMeeting features.

About the T.120 proxy agent

The T.120 proxy agent provides support for applications built using the International Telecommunication Union (ITU) T.120 recommendations. The T.120 recommendations are most prevalent in data conferencing applications. T.120 defines several standardized data conferencing services including application sharing, text chat, shared whiteboard, and multipoint file transfer.

Microsoft's NetMeeting is a popular example of a T.120 enabled application. The T.120 proxy agent enables you to use all of the standard T.120 data conferencing services, and provides you with a means to control which services are accessible. The T.120 proxy agent also provides support for the Microsoft NetMeeting chat and application sharing, which are non-standard T.120 application services.

Note: The audio, video, ILS, and ULS features of NetMeeting are not supported by the T.120 agent. These features are supported in the H.323 agent. To use this functionality, enable the default NetMeeting rule. This will ensure that services using both agents remain synchronized with one another. See [Synchronizing T.120 and H.323 for use with NetMeeting](#) for more information.

When configured, the T.120 proxy agent is transparent to the participants of the data conference. The T.120 proxy agent comes into play when a conference participant attempts to join an existing conference or attempts to invite another participant that resides in a different burb. The T.120 proxy agent intercepts and mediates the session between the pair of conference host machines. These host machines are referred to as *nodes* in T.120 parlance.

T.120 conferences are arranged into a hierarchy of nodes. The placement of the firewall with respect to the nodes in the conference affects how many sessions are created through the proxy agent and the communication path of the conference data. When a first conference participant joins a conference in a different burb, a T.120 session is created between the participant's node and the contacted node. If a second conference participant attempts to contact the new conference node, a separate session is created.

The preconfigured NetMeeting rule, when enabled, will apply to each participant's respective node IP address. If the second participant contacts the first participant and asks to join the conference, the same proxy session will be used. The NetMeeting rule that applies to the first participant's node will also apply to this session.

The T.120 proxy is configured to use port 1503 by default. This can be changed as described in [Create and modify services](#).

About the H.323 proxy agent

H.323 is an International Telecommunications Union (ITU) standard that provides support for audio and video conferencing across a shared medium such as the Internet. The H.323 proxy agent provides standard functions such as filtering on source and destination hosts and burbs, and NAT and redirection. The H.323 proxy agent is a protocol-aware, application layer agent that examines H.323 packets for correctness and adherence to site security policy. In addition to the standard filtering mentioned above, the H.323 agent provides a mechanism for allowing or disallowing certain codecs (audio or video encoding schemes) within the H.323 protocol.

Microsoft NetMeeting is a popular implementation of the H.323 protocol. The H.323 proxy agent enables you to use the audio and video features of data conferencing products like NetMeeting.

Note: The standard data conferencing features, as well as the chat and application sharing features of NetMeeting, are not supported by the H.323 agent. These features are supported in the T.120 agent. To use this functionality, enable the default NetMeeting rule. This will ensure that services using both agents remain synchronized with one another. See [Synchronizing T.120 and H.323 for use with NetMeeting](#) for more information.

The H.323 proxy agent can function between two endpoints (a single client implementation such as NetMeeting), or between one or more endpoints and a Multi-point Control Unit (MCU). The MCU enables two or more endpoints to simultaneously participate in a call. Each endpoint sends its audio and video signals through the firewall to the MCU. The MCU then combines the audio signals and selects one or more video signals to return to each endpoint.

Note: The H.323 agent does not recognize any configuration difference between an endpoint and an MCU.

The H.323 proxy agent must examine the contents of the protocol packets for encoded addresses and port numbers. Therefore, any sort of encryption of H.323 sessions is not possible in conjunction with the H.323 proxy agent. When implementing the H.323 protocol, you must disable NetMeeting's security features, or the security features of any other endpoint or MCU you may be using. Additionally, you must not route H.323 traffic through a VPN.

Also, any calls originating from the outside network and destined for a host on the internal network may be configured to use the netmaps feature. (For information on using netmaps, see "About the Network Objects: Netmap window" on page 70.) This provides a form of redirection that allows you to hide a group of addresses behind the firewall while still allowing the inbound caller to reach the proper destination machine.

About using a gatekeeper with the H.323 proxy agent

The H.323 proxy agent can also function between endpoints and a gatekeeper. A *gatekeeper* sits between source and destination endpoints and typically provides services such as authentication, authorization, alias resolution, billing and call routing. The RAS (Registration, Admission, and Status) protocol is used between the endpoints and the gatekeeper. RAS uses UDP port 1719.

If endpoints are configured to make use of the services of a gatekeeper, the firewall must be configured to properly handle this traffic. The preconfigured VoIP H.323 rule allows both conferencing services and RAS services to be provided by an H.323 proxy service. The conferencing services include audio/video and data, as in the NetMeeting rule previously discussed. When the endpoints are configured to use a gatekeeper, use an H.323 rule rather than the default NetMeeting rule.

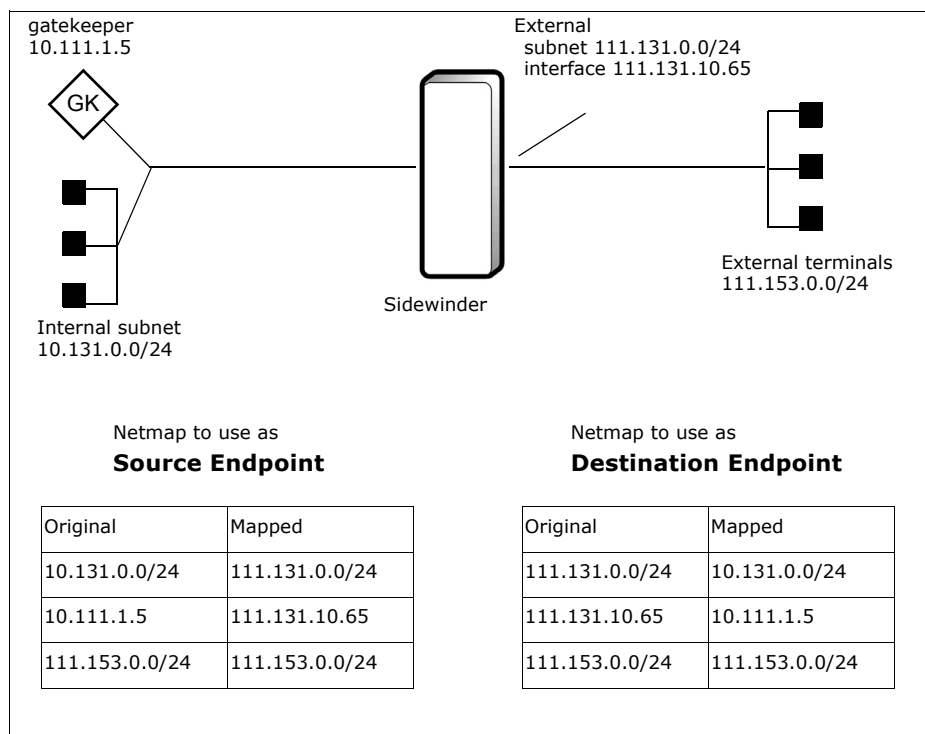
A gatekeeper can operate in one of two modes: *direct* and *routed*. The gatekeeper's mode is important when configuring the VoIP H.323 rule on the firewall. In direct mode, the gatekeeper grants permission for the call, but the call setup and call data are passed directly from endpoint to endpoint. In routed mode, the gatekeeper grants permission for the call and handles the call setup. Call data is then passed directly from endpoint to endpoint. The firewall policy must allow for the proper communication paths.

To appropriately restrict access for the H.323 proxy rule, configure network objects that describe the hosts receiving calls and sending calls. Also configure a network object for the gatekeeper. The source and destination of the H.323 rule should contain the endpoints and the gatekeeper as appropriate for the mode of operation configured on the gatekeeper. This may include adding netmaps to add all call endpoints and the gatekeeper to a single rule, and making changes to the H.323 configuration file to support your gatekeeper environment.

If the gatekeeper is on the internal network, configure a netmap to allow hosts on the outside network to communicate with the gatekeeper as well as with endpoints on the internal network. The netmap needs to include the gatekeeper, the hosts allowed to initiate calls, and the hosts allowed to receive calls. The internal gatekeeper and internal hosts permitted to send and receive calls must be mapped to external address. If the internal hosts are exchanging calls with terminals on the Internet, then the mapped addresses must be publicly routable.

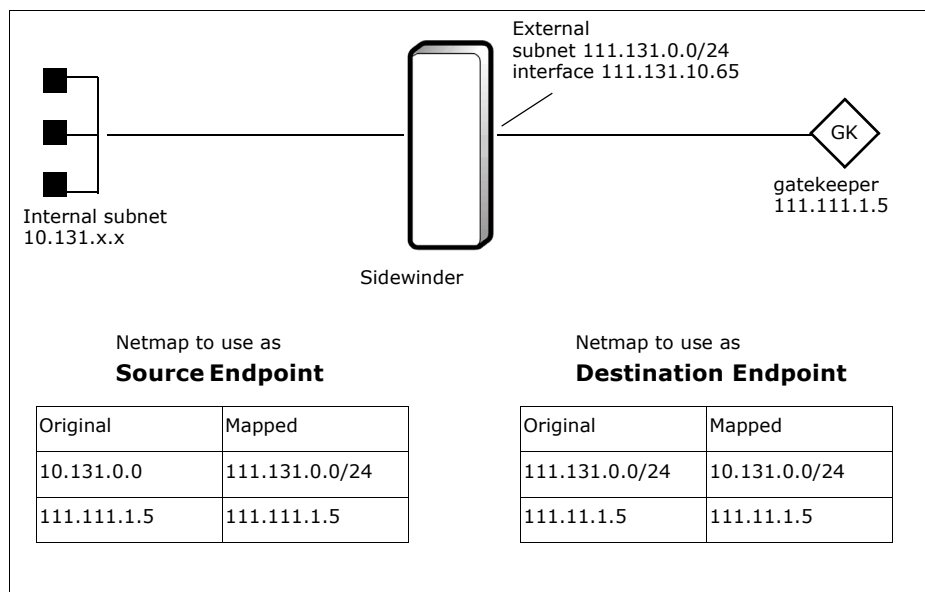
If the gatekeeper is not on the same subnet as the hosts permitted to receive incoming calls, then the netmap must include a mapping of the gatekeeper to itself so the gatekeeper is a recognized destination.

Figure 110 Rule source and destination netmaps when gatekeeper is internal



If the gatekeeper is on the external network, then a connection may be made from the gatekeeper to any internal host permitted to receive calls. Configure a netmap that includes the internal terminals and a mapping of the gatekeeper to itself, so that the gatekeeper is a recognized destination. If external terminals will be allowed to initiate or receive calls, they should also be added to the netmap.

Figure 111 Rule source and destination netmaps when gatekeeper is external



If the gatekeeper is in a burb completely separate from the call endpoints, you need to adjust the H.323 configuration file. To edit the file so the firewall recognizes that the gatekeeper is in a burb separate from the call endpoints, do the following:

1 Using a file editor, open `/secureos/etc/proxy/h323p.conf`.

2 Locate the following lines:

```
gatekeeper_alone[NO]
```

3 Change `[NO]` to `[YES]`.

4 Save your changes and exit the file.

5 Restart the H.323 Proxy agent:

a Select **Monitor > Service Status**.

b Select **h323**.

c Click **Restart**.

The firewall adjusts its routing accordingly.

In general, gatekeepers pass the IP address of the call-initiator endpoint to the call-receiver endpoint. This allows the systems to verify both ends of the connection. However, some gatekeepers pass their own IP address instead of the call initiator's address. When the firewall cannot verify the other endpoint, it ignores the connection and generates the following audit message:

```
H.245 connect received from unknown_ip_addr while expecting one from known_ip_addr.  
Unexpected connect ignored.
```

If your gatekeeper does not pass the call initiator's IP address, you need to adjust the H.323 configuration file. To edit the file so the firewall allows connections where the initiating IP address cannot be verified, do the following:

Caution: Making this change decreases security. Do not edit this value unless it is required for your gatekeeper configuration.

1 Using a file editor, open `/secureos/etc/proxy/h323p.conf`.

2 Locate the following line:

```
accept_anonymous_endpoint [NO]
```

3 Change `[NO]` to `[YES]`.

4 Save your changes and exit the file.

5 Restart the H.323 Proxy agent:

a Select **Monitor > Service Status**.

b Select **h323**.

c Click **Restart**.

The firewall now accepts H.245 connections from unknown IP addresses.

Synchronizing T.120 and H.323 for use with NetMeeting

The T.120 and H.323 proxy agents can work together, allowing you to make use of both the data-sharing and audio/video features of NetMeeting in a single conference as follows:

- The T.120 proxy agent enables you to use all of the standard T.120 data conferencing services and provides you with a means to control which services are accessible. The T.120 proxy agent also provides support for the Microsoft NetMeeting chat and application sharing, which are non-standard T.120 application services.
- The H.323 proxy agent provides support for the audio and video features of NetMeeting.

To make use of both the data-sharing and audio/video features of NetMeeting in a single conference, you must ensure that both the T.120 and H.323 proxy services are enabled in the same burbs. This is necessary because for a single NetMeeting session, part of the traffic (the H.323 portion) is routed through the H.323 proxy, and part of the traffic (the T.120 portion) is routed through the T.120 proxy. If the H.323 and T.120 proxy configurations are out of synchronization, it is likely that NetMeeting conferences will not function correctly or completely (for example, audio and video work, but data-sharing does not work).

Services

Configuring additional proxy agent properties

To prevent the two services from becoming out of synchronization, enable the preconfigured NetMeeting rule. The NetMeeting rule allows access to both the T.120 and H.323 proxy services (using the preconfigured NetMeeting Service Group), and allows access to all available NetMeeting features.

You can modify the default NetMeeting rule or create your own rules to allow only a portion of NetMeeting's features, such as the chat and whiteboard features. These properties are configured via the Multimedia Application Defense. For information on configuring Application Defenses for H.323/T.120, see "Creating T.120 Application Defenses" on page 235.

To appropriately restrict access for the NetMeeting proxy rule, configure network objects or other rule elements. For example, if you want to allow only administrators access to all NetMeeting features, create and specify a network object within a rule that contains the IP addresses for all of your administrators.

Services

Configuring additional proxy agent properties

8 Application Defenses

Contents

[Understanding Application Defenses](#)
[Creating HTTP or HTTPS Application Defenses](#)
[Creating Mail \(Sendmail\) Application Defenses](#)
[Creating Mail \(SMTP proxy\) Defenses](#)
[Creating Citrix Application Defenses](#)
[Creating FTP Application Defenses](#)
[Creating ILOP Application Defenses](#)
[Creating T.120 Application Defenses](#)
[Creating H.323 Application Defenses](#)
[Creating Oracle Application Defenses](#)
[Creating MS SQL Application Defenses](#)
[Creating SOCKS Application Defenses](#)
[Creating SNMP Application Defenses](#)
[Creating SIP Application Defenses](#)
[Creating SSH Application Defenses](#)
[Creating Packet Filter Application Defenses](#)
[Configuring Application Defense groups](#)

Understanding Application Defenses

Forcepoint Sidewinder policy is applied primarily by rules, which are made up of many elements. The table below shows the progression of a rule's creation using these elements and their corresponding chapters in this guide.

You are here in the Policy section	Use this chapter to...
Chapter 3, Policy Configuration Overview	understand the policy creation process.
Chapter 4, Network Objects and Time Periods	create or modify any network objects or time periods that will be used by rules.
Chapter 5, Authentication	create or modify authenticators that will be used by rules.
Chapter 6, Content Inspection	configure content inspection methods that will be used by rules.
Chapter 7, Services	create or modify services or service groups that will be used by rules.
Chapter 8, Application Defenses	create or modify Application Defenses that will be used by rules.
Chapter 9, Rules	create rules using the elements you created in the previous chapters in the policy section.

Use Application Defenses to configure advanced properties for rules. You can refine rules for specific applications that use proxies and filter agents. You can also configure key services such as anti-virus/anti-spyware, SSL decryption, and web services management.

- You configure Application Defenses in the appropriate Application Defense window.
- An Application Defense is selected in the Rules window. Certain services have related Application Defenses that you can apply to the rule you are creating.

To view the Application Defenses windows, select **Policy > Application Defenses > Defenses**, and then select the type of Application Defense you want to view from the tree. A window similar to the following appears:

Figure 112 Application Defenses window (HTTP)

Name	Description	Type	HTTP URL control	FTP URL control	HTTP req
default	default http application defense.	Combined	Off	Off	Off
SmartFilter Redirect	SmartFilter redirect for explanations.	Server	On	Off	On

Buttons: New, Modify, Rename, Delete, Usage, Duplicate

Name: default Type: Combined Description: default http application defense.

Enforcements: HTTP URL control | FTP URL control | HTTP request | HTTP reply | MIME/Virus/Spyware | Content control | SmartFilter | Connection

HTTP client and server enforcements

- ☐ HTTP URL control
- ☐ FTP URL control
- ☐ HTTP request
- ☐ HTTP reply
- ☐ MIME/Virus/Spyware
- ☐ Content control
- ☐ SmartFilter

☐ Relax protocol enforcements: Client

The top pane of each Application Defense window consists of a table that lists all of the Application Defenses (by row) that are currently configured for the category selected in the tree.

- The Application Defenses that are displayed in the table will vary depending on the defense category you select from the tree.
- The table columns display the attributes for the selected defense. The columns will vary by application defense.
- Basic default defenses (such as **default**) are pre-configured for each category of Application Defense.

You can perform the following actions in any of the Application Defense windows:

- **Create a new Application Defense** – To create a new Application Defense:
 - a Select the appropriate type of defense in the tree, then click **New**. The New Application Defense window appears.
 - b Type a name for your application defense. If you are creating an HTTP or HTTPS Application Defense, select a type.
 - c Click **OK** and modify the properties in the lower portion of the window.
- **Duplicate an existing Application Defense** – To duplicate an existing Application Defense:
 - a Select the appropriate defense from the table, then click **Duplicate**. The New/Duplicate Application Defense window appears.
 - b Type a name for your application defense. (If you are duplicating an HTTP or HTTPS Application Defense, you cannot select a type.)
 - c Click **OK** and modify the properties in the lower portion of the window.
- **Modify an existing Application Defense** – Select the defense that you want to modify from the table. The configuration information is displayed in the bottom portion of the window.
To modify the Application Defense in a pop-up window format, click **Modify**. (Read-only administrators can click **View** to view an Application Defense in a pop-up window.)
- **Rename an existing Application Defense** – Select the appropriate Application Defense from the table and click **Rename**, then type a new name in the Rename window.
- **Delete an existing Application Defense** – Select the appropriate Application Defense from the table and click **Delete**.
Note: You cannot delete an Application Defense if it is being used in a rule or a group. If the Application Defense is used in a rule, a pop-up window will appear informing you which rules are currently using this defense. Before you can delete the defense, you will need to modify each of the rules to remove the specified defense from those rules.
- **View the rules in which an Application Defense/Group is currently used** – Select the appropriate defense (or group) and click **Usage**. A pop-up window appears listing the rule and group names that are currently using the specified defense. Click **Close** when you are finished viewing the rule list.

The bottom portion of each window (or pop-up, if you clicked **Modify**) displays the actual configuration information for the selected Application Defense. The information will vary depending on the Application Defense category you select. The following fields remain constant among all Application Defense windows:

- **Name** – This field contains the name of the Application Defense that you are viewing. If you need to rename an Application Defense, click **Rename** and type a new name.
- **[HTTP/HTTPS only] Type** – Use this field to specify whether a defense will be used to protect a server, client, or both. For more information about the Type field, see [Creating HTTP or HTTPS Application Defenses](#).
- **Description** – Use this field to provide information about the Application Defense to help you more easily identify it.

For information on configuring a specific Application Defense, see the following:

- [Creating HTTP or HTTPS Application Defenses](#)
- [Creating Mail \(Sendmail\) Application Defenses](#)
- [Creating Mail \(SMTP proxy\) Defenses](#)
- [Creating Citrix Application Defenses](#)
- [Creating FTP Application Defenses](#)
- [Creating IIOF Application Defenses](#)
- [Creating T.120 Application Defenses](#)
- [Creating H.323 Application Defenses](#)
- [Creating Oracle Application Defenses](#)
- [Creating MS SQL Application Defenses](#)
- [Creating SOCKS Application Defenses](#)
- [Creating SNMP Application Defenses](#)
- [Creating SIP Application Defenses](#)
- [Creating SSH Application Defenses](#)
- [Creating Packet Filter Application Defenses](#)

Note: For information on configuring Application Defense groups, see [Configuring Application Defense groups](#).

Creating HTTP or HTTPS Application Defenses

The HTTP/HTTPS Application Defenses allow you to configure advanced parameters for HTTP or HTTPS and SSO proxy rules. To create HTTP or HTTPS Application Defenses, select **Policy > Application Defenses > Defenses** and then select **HTTP** or **HTTPS**. One of the following windows appears.

Figure 113 Application Defense (default): HTTP and HTTPS

The figure displays two screenshots of the Application Defense configuration interface. The left screenshot shows the 'default' HTTP application defense configuration, and the right screenshot shows the 'default' HTTPS application defense configuration.

Left Screenshot (HTTP Application Defense):

Name	Description	Type	HTTP URL control	FTP URL control	HTTP req
default	default http application defense.	Combined	Off	Off	Off
SmartFilter Redirect	SmartFilter redirect for explanations.	Server	On	Off	On

Enforcements: HTTP URL control | FTP URL control | HTTP request | HTTP reply | MIME/Virus/Spyware | Content control | SmartFilter | Connection

HTTP client and server enforcements

- ☐ HTTP URL control
- ☐ FTP URL control
- ☐ HTTP request
- ☐ HTTP reply
- ☐ MIME/Virus/Spyware
- ☐ Content control
- ☐ SmartFilter

☐ Relax protocol enforcements: Client

Right Screenshot (HTTPS Application Defense):

Name	Description	Type	Decrypt	URL control	HTTP request
default	default https application defense.	Server	Off	Not Applicable	Not Applicable
Passport	Single sign-on application defense.	Server	On	On	On
SmartFilter Admin	SmartFilter administration application defense.	Server	On	On	On

Enforcements: HTTP URL control | HTTP request | HTTP reply | MIME/Virus/Spyware | Content control | SmartFilter | Connection

☒ Do not decrypt HTTP traffic:
 HTTPS server enforcements: SSL connections will be validated
☐ SmartFilter

☐ Decrypt HTTP traffic:
 (Used to decrypt inbound HTTP traffic for a protected web server, traffic is decrypted and relayed.
 NOTE: Redirection must be specified on proxy rules using this Application Defense.)

HTTPS server enforcements

- ☐ HTTP URL control
- ☐ HTTP request
- ☐ HTTP reply
- ☐ MIME/Virus/Spyware
- ☐ Content control
- ☐ SmartFilter

Firewall certificate: Select One --> (To add certificates, see Certificate Management screen.)

SSL settings

Relax protocol enforcements: Client

Rewrite Microsoft's OWA HTTP

Configuring the HTTP/HTTPS: Enforcements tab

Use the Enforcements tab to select the feature enforcement tabs that you want to make available for configuration, as well as relax enforcement of HTTP proxy standards. If you are configuring an HTTPS Application Defense, you can also configure SSL decryption properties in the Enforcements tab.

In the **Type** drop-down list, you can specify whether this defense will be used to protect a server, client, or both:

- **Combined** – [HTTP only] This option allows you to create an Application Defense that can protect both an HTTP client (outbound) and an HTTP server (inbound) behind the Sidewinder. When you select this option, all of the configuration options for this defense will appear. However, some of the options that you configure will only apply to the client or server. (For example, HTTP Request properties do not apply to the client. Therefore, if you select **Combined**, HTTP Request properties that you configure will only apply to the server.)
- **Client** – This option allows you to create an Application Defense that protects a client behind the Sidewinder. Options that do not apply for client protection (such as HTTP Requests) will not be available for configuration.
- **Server** – This option allows you to create an Application Defense that protects a server behind the Sidewinder. Options that do not apply for server protection (such as Content Control options other than SOAP) will not be available for configuration.

To enable enforcement of HTTP proxy standards in a manner that allows traffic from systems that do not adhere to strict RFC standards for the HTTP proxy, select the **Relax Protocol Enforcements** option. Enabling relaxed mode allows the following RFC infractions:

- Media types in Content-Type: headers in a relaxed form, where the subtype is not required
- Empty headers
- Duplicated responses from the server where the response is the same but the version is different
- Query strings containing arbitrary data

Caution: Each listed infraction introduces an element of risk into your security policy, particularly if enabled on server-side rules. Use this mode only when necessary, and implement on a rule-by-rule basis.

Select the **Relax Protocol Enforcements** option if the above infractions are acceptable or required in your network. When you enable this option, you will also need to specify whether the protocol enforcements will be relaxed when receiving HTTP traffic from clients, servers, or both by selecting one of the following options from the drop-down list:

- **Client** – Select this option to relax protocol enforcements only when receiving HTTP traffic from clients.
- **Server** – Select this option to relax protocol enforcements only when receiving HTTP traffic from servers.
- **Client and Server** – Select this option to relax protocol enforcements when receiving HTTP traffic from both clients and servers.

Enabling HTTP/HTTPS configuration tabs

To configure an HTTP or HTTPS tab, enable the service on the Enforcements tab. You cannot configure a tab unless it is enabled.

- The Connection tab for HTTP and HTTPS does not need to be enabled before you can configure it.
- If you are configuring an **HTTPS** defense and you select the **Decrypt HTTP Traffic** check box, you can enable any of the tabs below. If you select the **Do Not Decrypt HTTP Traffic** check box, you can enable only the SmartFilter tab.

Note: SmartFilter is not actively supported in Sidewinder 7.x.

The following tabs can be enabled:

- **HTTP URL Control** – Use the HTTP URL Control tab to configure filtering on the URL contained in the HTTP request. To enable URL filtering, select this check box. To configure HTTP URL filtering properties, select the HTTP URL Control tab and see [Configuring the HTTP/HTTPS: HTTP URL Control tab](#).
- **FTP URL Control** – Use the FTP URL Control tab to configure filtering on the URL contained in an HTTP request for FTP traffic. To enable FTP URL filtering, select this check box. To configure FTP URL filtering properties, select the FTP URL Control tab and see [Configuring the HTTP: FTP URL Control tab](#).
- **HTTP Request** – Use the HTTP Request tab to configure header filtering on HTTP requests. To enable HTTP header filtering for HTTP requests, select this check box. To configure HTTP header request properties, select the HTTP Request tab and see [Configuring the HTTP/HTTPS: HTTP Request tab](#).
- **HTTP Reply** – Use the HTTP Reply tab to configure header filtering on HTTP replies. To enable HTTP header filtering for HTTP replies, select this check box. To configure HTTP header reply properties, select the HTTP Reply tab and see [Configuring the HTTP/HTTPS: HTTP Reply tab](#).
- **MIME/Virus/Spyware** – Use the MIME/Virus/Spyware tab to configure MIME (Multi-Purpose Internet Mail Extensions) and anti-virus/spyware filtering, and infected file handling. To enable filtering for MIME/virus/spyware, select this check box. To configure MIME/virus/spyware properties, select the MIME/Virus/Spyware tab and see [Configuring the HTTP/HTTPS: MIME/Virus/Spyware tab](#).
- **Content Control** – Use the Content Control tab to configure filtering for web content types including ActiveX, Java, scripting languages, and SOAP. (For HTTPS, you can only configure SOAP filtering.) To enable content filtering, select this check box. To configure content control properties, select the Content Control tab and see [Configuring the HTTP/HTTPS: Content Control tab](#).
- **SmartFilter** – Use the SmartFilter tab to enable filtering of web traffic using SmartFilter.

Note: Do not alter the **SmartFilter Redirect** Application Defense. This Application Defense is used on the rule that enables communication with the SmartFilter server.

For information on configuring the SmartFilter tab, see [Configuring the HTTP/HTTPS: SmartFilter tab](#).

Configuring SSL decryption properties [HTTPS server only]

The Sidewinder can perform SSL decryption services at the firewall level on a per-rule basis, increasing the security of your data transactions.

To use SSL decryption services on the Sidewinder, you must have the following features licensed:

- **Strong Cryptography** – This feature is included with the basic Sidewinder license.
- **SSL Decryption** – This feature is an add-on module. If it is purchased after the Sidewinder's initial activation, you will need to relicense your firewall to activate this feature. For licensing information, see [Activating the Sidewinder license on page 586](#).

To configure decryption properties for an HTTPS Application Defense, follow the steps below.

Note: Proxy rules that use HTTPS Application Defenses with the **Decrypt HTTP Traffic** option enabled must have redirection configured.

1 Select from the following:

- To enable SSL decryption for an Application Defense, select **Decrypt HTTP Traffic**. Remember to verify that the SSL Decryption and Strong Cryptography features are licensed.
- To allow HTTP traffic to pass through without being decrypted, select **Do Not Decrypt HTTP Traffic**. SSL connections will be validated when this option is selected. If you select this option, you can select the SmartFilter check box to enable web filtering and enable the SmartFilter tab for configuration.

Note: SmartFilter is not actively supported in Sidewinder 7.x.

- 2** [Conditional] If you are configuring an HTTPS defense to allow clientless VPN sessions to access a Microsoft Exchange® Server, select the **Rewrite Microsoft OWA HTTP** check box.
- 3** Select the appropriate firewall certificate from the **Firewall Certificate** drop-down list. This is the certificate that is used to authenticate the Sidewinder to the remote HTTPS/SSL client. For information on configuring firewall certificates, see [About Certificate/Key Management on page 619](#).
- 4** Click **SSL Settings** to configure SSL properties. Configurable properties include specifying the accepted SSL/TLS versions and the minimum cryptography strength.
- 5** Save your changes.

Configuring the HTTP/HTTPS: HTTP URL Control tab

Use the HTTP URL Control tab to configure URL control properties, such as which HTTP operations will be allowed and which URLs will be explicitly denied or allowed.

Figure 114 HTTP/HTTPS: HTTP URL Control tab

Enforcements | HTTP URL control | FTP URL control | HTTP request | HTTP reply | MIME/Virus/Spyware | Content control | SmartFilter | Connection

Allow selected HTTP commands:

- ☒ BASELINE-CONTROL - Creates snapshot of version-controlled V
- ☒ CHECKIN - Checks in a version-controlled WebDAV resource
- ☒ CHECKOUT - Checks out a version-controlled WebDAV resource
- ☒ CONNECT - Allows proxies to act as tunnels for SSL
- ☒ COPY - Copy file to server from location identified in request
- ☒ DELETE - Delete file on server
- ☒ GET - Get file from server

Select All Deselect All

☐ Enforce strict URLs ☐ Allow unicode

Maximum URL length: 1024

☒ Require HTTP version in request

Allow selected HTTP versions:

- ☒ 1.0 ☒ 1.1

☒ Deny / ☐ Allow specified URL matches:

Match type	Match parameter	String
------------	-----------------	--------

New Modify Delete

To configure the HTTP URL Control tab:

- 1 In the **Allow Selected HTTP Commands** area, select the commands (operations) that you want to allow users to issue by clicking in the corresponding check box(es).
To select all of the commands, click **Select All**. To clear all of the commands, click **Deselect All**. A description of each command is provided within the window.
- 2 To disallow special characters in a query, select the **Enforce Strict URLs** check box. If you select this option, URLs with certain special characters will be disallowed under certain circumstances (such as RFC violation). For example: quote ("), back quote (`), brackets ([], { }, < >), pipe (|), back slash (\), and caret (^).
- 3 To allow international multi-byte characters in a query, select the **Allow Unicode** check box.
- 4 [Server or Combined only] In the **Maximum URL Length** field, specify the maximum length allowed for a URL. The default value is 1024 characters. Valid values are 1–10000.
- 5 To require that the HTTP version be included in all requests, select the **Require HTTP Version in Request** check box.
- 6 [Conditional] If you selected **Require HTTP Version in Request** in the previous step, specify the HTTP versions that you want to allow in the **Allow Selected HTTP Versions** area: version 1.0 and 1.1 are available.
- 7 In the **Deny / Allow Specified URL Matches** table, you can specify strings that can be matched to parts of the URL. Select one of the following options to control enforcement behavior:
 - **Deny** – If the string is found in a particular URL, the request is explicitly denied. The table lists the match strings that are currently denied.
 - **Allow** – If the string is found in a particular URL, the request is allowed. The table lists the match strings that are currently allowed.

Tip: URLs that do not contain a string listed in the table are denied.

To add a match string to the list, click **New**. To modify a match string in the list, select it and click **Modify**. To remove a match string from the list, select it and then click **Delete**.

Configuring the HTTP: FTP URL Control tab

Use this tab to control access to FTP servers through HTTP proxies. Access to FTP servers is allowed by default.

Figure 115 HTTP: FTP URL Control tab

The screenshot shows a configuration window with a tabbed interface. The tabs are: Enforcements, HTTP URL control, FTP URL control, HTTP request, HTTP reply, MIME/Virus/Spyware, Content control, SmartFilter, and Connection. The 'FTP URL control' tab is selected. Inside this tab, there is a section titled 'Allow selected FTP commands:' with a list box containing two items: 'GET - Get file from server' and 'PUT - Put file on server'. Both items are checked. Below the list box are two buttons: 'Select All' and 'Deselect All'. To the right of the list box, there is a section titled 'Data connection type:' with three radio buttons: 'Active', 'Passive', and 'Both'. The 'Both' radio button is selected.

You can perform the following actions:

- **Select the type of commands you will allow** – You can allow FTP traffic to upload and/or download files and directories from an FTP server.
 - Select **GET** to allow files to be downloaded. Clear this option to deny downloaded files.
 - Select **PUT** to allow files to be uploaded. Clear this option to deny uploaded files.
 - Use the **Select All** and **Deselect All** buttons to select or clear both options at once.
- **Select the data connection type** – Select which commands the firewall sends to the FTP server to initiate the data exchange:
 - **Active** – Select this option to tell the FTP server which port to send data to.
 - **Passive** – Select this option to allow the FTP server to specify which port to send data to.
 - **Both** – Select this to make both options available. The passive option is tried first. This is the default selection.

Configuring the HTTP/HTTPS: HTTP Request tab

Use the HTTP Request tab to configure header filtering for HTTP requests. This tab is only available if you selected **Server** or **Combined** in the **Type** field.

Figure 116 HTTP/HTTPS: HTTP Request tab

Enforcements | HTTP URL control | FTP URL control | **HTTP request** | HTTP reply | MIME/Virus/Spyware | Content control | SmartFilter | Connection

☐ Allow / ☒ Deny Selected HTTP request header filter types:

☐ None
☐ Standard
☐ Paranoid
☒ Custom

Accept
Accept-Charset
Accept-Encoding
Accept-Language
Allow
Authorization
Cache-Control
Connection
Content-Encoding
Content-Language
Content-Length
Content-Location
Content-MD5

Select All Deselect All

Denied Header Action

☒ Block entire page
☐ Allow page through without denied headers

Denied header values:

Header	Value
--------	-------

New Modify Delete

☐ Deny binary data

Note: The fields in this tab will be disabled unless you select the **HTTP Request** check box on the Enforcements tab.

To configure the HTTP Request tab:

- 1 Select the type of HTTP header filtering you want to allow or deny in the **Selected HTTP Request Header Filter Types** area:

Note: The **X-*** filter type is a wildcard filter that will allow or deny all X-xxx request headers (commonly found in user-defined headers). If you create an **Allow** list and do not include the **X-*** filter type, most HTTP traffic will be denied.

- **None** – Select this option if you want to clear all HTTP request header filter types in the list. (You can also clear all of the types by clicking **Deselect All**.)
- **Standard** – Select this option if you want to automatically select all of the header types contained in the list. (You can also select all header types by clicking **Select All**.)
- **Paranoid** – Select this option if you want to exclude all options not defined in the RFC.
- **Custom** – Select this option if you want to manually select which HTTP header types you will allow or deny.

Note: Header types that are not in the list are handled the same as unselected header types.

- 2 In the **Filter Option** field, determine whether you want to allow or deny the header types you select, as follows:
 - **Allow** – Select this option to *allow* all header types that are selected in the HTTP Request Header Filter Types window. All other types will be denied.
 - **Deny** – Select this option to *deny* all header types that are selected in the HTTP Request Header Filter Types window. All other types will be allowed.
- 3 In the **Denied Header Action** area, select one of the following options:
 - **Block Entire Page** – Select this option to block the entire page when an HTTP header is denied.
 - **Allow Page Through Without Denied Headers** – Select this option to mask the denied HTTP header, but still allow the page to be viewed. (A denied HTTP header will be overwritten with Xs.)
- 4 In the **Denied header values** area, you can create a list of headers and matching values that you want blocked. If a specified header appears in a request or response, and it contains the specified value, it is dropped from the message.
 - Full header names must be used.

- Regular expressions are not supported.
- Values are matched in a case-insensitive manner, and are used exactly as specified.

Click **New** to create a new header and value. Click **Modify** to change an existing header.

Note: For more information on HTTP message headers, refer to RFC 2616 which can be found at www.ietf.org/rfc.html.

5 To block headers that contain binary data, select **Deny binary data**.

Every header is scanned to detect binary data. This prevents attacks that put binary data in requests.

- Binary data means ASCII codes 0x00 to 0x1f and 0x7f hexadecimal.
- This does not affect escaped characters that convert to legal ASCII characters. For example, %41 in a header would convert to the letter A in ASCII.

Note: This feature reduces your firewall's performance.

Configuring the HTTP/HTTPS: HTTP Reply tab

Use the HTTP Reply tab to configure header filtering for HTTP replies. Follow the steps below.

Figure 117 HTTP/HTTPS: HTTP Reply tab

The screenshot shows the 'HTTP Reply' tab within a configuration window. The 'Enforcements' tab is selected, and the 'HTTP reply' sub-tab is active. The 'Allow / Deny Selected HTTP reply header filter types:' section has the 'Deny' radio button selected. A list of header types is shown, including 'Accept-Ranges', 'Age', 'Allow', 'Cache-Control', 'Connection', 'Content-Encoding', 'Content-Language', 'Content-Length', 'Content-Location', 'Content-MD5', 'Content-Range', and 'Content-Type'. The 'Deny Header Action' section has the 'Block entire page' radio button selected. The 'Denied header values:' section is empty, with 'New', 'Modify', and 'Delete' buttons below it. The 'Deny binary data' checkbox is unchecked.

Note: The fields in this tab will be disabled unless you select the **HTTP Reply** check box on the Enforcements tab. Also, this tab is not available for HTTPS if you select **Client** in the **Type** field.

To configure the HTTP Reply tab:

- 1 Select the type of HTTP header filtering you want to allow or deny in the **Selected HTTP Reply Header Filter Types** area. The following options are available:

Note: The **X-*** filter type is a wildcard filter that will allow or deny all X-xxx reply headers (commonly found in user-defined headers). If you create an **Allow** list and do not include the **X-*** filter type, most HTTP traffic will be denied.

- **None** – Select this option if you want to clear all HTTP reply header filter types in the list. (You can also clear all of the types by clicking **Deselect All**.)
- **Standard** – Select this option if you want to automatically select all of the header types contained in the list. (You can also select all header types by clicking **Select All**.)
- **Paranoid** – Select this option if you want to exclude all options not defined in the RFC.
- **Custom** – Select this option if you want to manually configure which HTTP reply header types you will allow or deny.

Note: Header types that are not in the list are handled the same as unselected header types.

- 2 In the **Filter Option** field, determine whether you want to allow or deny the header types you select, as follows:

- **Allow** – Select this option to *allow* all header types that are selected in the HTTP Reply Header Filter Types window. All other types will be denied.
- **Deny** – Select this option to *deny* all header types that are selected in the HTTP Reply Header Filter Types window. All other types will be allowed.

- 3 In the **Denied Header Action** area, select one of the following options:

- **Block Entire Page** – Select this option to block the entire page when an HTTP reply header is denied.
- **Allow Page Through Without Denied Headers** – Select this option to mask the denied HTTP reply header, but still allow the page to be viewed. (A denied HTTP reply header will be scrubbed.)

- 4 In the **Denied header values** area, you can create a list of headers and matching values that you want blocked. If a specified header appears in a request or response, and it contains the specified value, it is dropped from the message.

- Full header names must be used.
- Regular expressions are not supported.
- Values are matched in a case-insensitive manner, and are used exactly as specified.

Click **New** to create a new header and value. Click **Modify** to change an existing header.

Note: For more information on HTTP message headers, refer to RFC 2616 which can be found at www.ietf.org/rfc.html.

- 5 To block headers that contain binary data, select **Deny binary data**.

Every header is scanned to detect binary data. This prevents attacks that put binary data in requests.

- Binary data means ASCII codes 0x00 to 0x1f and 0x7f hexadecimal.
- This does not affect escaped characters that convert to legal ASCII characters. For example, **%41** in a header would convert to the letter **A** in ASCII.

Note: This feature reduces your firewall's performance.

Configuring the HTTP/HTTPS: MIME/Virus/Spyware tab

Use the MIME/Virus/Spyware tab to configure filtering for MIME, virus, and spyware scanning services. The tab contains a rule table that displays any MIME/Virus/Spyware filtering rules that have been created. The tab also contains various virus scanning and handling configuration options.

Figure 118 HTTP/HTTPS: MIME/Virus/Spyware tab

Pos	MIME Type	Extensions	Action
1	DEFAULT	DEFAULT	allow

Scanner Behavior

☒ Reject all files if scanning is unavailable

☐ Use heuristic scanning (scan for unknown viruses)

Infected File Handling

☒ Discard infected files

☐ Repair infected files

Maximum Scan Size

Scan file size limit (KB): 32768

☐ Files over the scan limit will be allowed through unscanned

☒ Files over the scan limit will be rejected

Security Alert: If you want to perform virus and spyware scanning, you must create the appropriate MIME rules with **Virus/Spyware Scan** selected in the **Action** field. Rules that are configured only to allow or deny traffic based on rule criteria will not perform virus and spyware scanning. (See [Step 1](#) for information on configuring MIME/Virus/Spyware filter rules.)

- The fields in the MIME/Virus/Spyware tab will be disabled unless you select the **MIME/Virus/Spyware** check box on the Enforcements tab.
- For HTTP defenses, MIME/Virus/Spyware scanning services are not available if you select **Server** in the Type field.
- For HTTPS defenses, MIME/Virus/Spyware scanning services are not available if you select **Client** in the Type field.
- The MIME type tells the browser or server what type of information it is receiving.
- Virus and spyware scanning is performed on data sent from the client if the request method is either PUT or POST, and the appropriate file type is specified for scanning in the MIME/Virus/Spyware filtering rules table.

Note: You must license scanning services before the MIME/Virus/Spyware filter rules you create will scan HTTP/HTTPS traffic. See [Configuring virus scanning services on page 134](#).

To configure the MIME/Virus/Spyware tab:

- 1 Configure the appropriate MIME/Virus/Spyware filter rules in the **MIME/Virus/Spyware Filter Rules** table:
 - **Create a new filter rule** – To create a new filter rule, click **New**. See [About the MIME Rule Edit window](#).
 - **Modify an existing filter rule** – To modify an existing filter rule, select the rule you want to modify, and click **Modify**. See [About the MIME Rule Edit window](#). (If you are modifying the default MIME filtering rule, see [Configuring the Default filtering rule action](#).)
 - **Delete a filter rule** – To delete an existing filter rule, select the rule you want to delete and click **Delete**.
- 2 To reject all files in the event that scanning is not available, select the **Reject all files if scanning is unavailable** check box. If you select this option, the connection will be dropped if scanning is unavailable (for example, due to out-of-date virus data, an expired license, or a configuration error).
- 3 To scan files for viruses for which virus signatures do not exist, select the **Use heuristic scanning (scan for unknown viruses)** check box.

Note: Enabling this option may reduce virus scanning performance.
- 4 Determine how infected files will be handled in the **Infected File Handling** area as follows:
 - To discard infected files, select **Discard infected files**.

- To remove the virus from the file and then continue processing the file, select **Repair infected files**.

5 Configure the **Maximum Scan Size** area.

- In the **Scan file size limit (KB)** field, specify the maximum file size that will be allowed in KB.
- Determine how files larger than the **Scan file size limit** will be handled by selecting one of the following:
 - **Files over the scan limit will be allowed through unscanned**
 - **Files over the scan limit will be rejected**

About the MIME Rule Edit window

Use this window to add or modify MIME/Virus/Spyware filtering rules.

- Rules that are configured with an allow or deny action will allow or deny traffic based on the rule criteria that is defined for those rules. Allow and deny rules do not perform virus scanning. To perform virus scanning for traffic that matches a rule before it is allowed, you must specify **Virus/Spyware Scan** in the rule's **Action** field.
- Rules that specify both a MIME type/subtype and file extensions will allow or deny any traffic that matches either the MIME Type **or** a File Extension type. That is, the traffic does not need to match both criteria to match the rule.

Figure 119 Mime Rule Edit window

To add or modify MIME/Virus/Spyware filtering rules:

- 1 In the **MIME Type** drop-down list, select the MIME type for which you want to filter. If you select the asterisk (*) option, the filter rule will ignore this field when determining a match.
- 2 In the **MIME Subtype** drop-down list, select a subtype for the MIME type that you selected in the previous step (the available options will vary depending on the MIME type you selected). If you select the asterisk (*) option, the filter rule will ignore this field when determining a match.
- 3 In the **File Extensions** area, specify the type of file extensions that you want to filter:
 - **Ignore Extensions (*)** – Select this option to ignore extensions when determining a match.
 - **Archive Extensions** – Select this option to specify basic archive extensions (such as *.tar*, *.zip*, etc.) for the specified MIME types/sub-type.

- **Standard Extensions** – Select this option to specify the standard file extensions associated with the selected MIME type/subtype. For example, if you select text in the **MIME Type** field, and HTML in the **MIME Subtype** field, the *.htm* and *.html* file extensions will appear in the standard list.
- **Custom** – Select this option to create a custom list of file extensions for the selected MIME type/subtype.
 - To add a file extension to the list, click **New** and type the extension (*without* the leading period) that you want to add.
 - To delete a file extension, select the extension you want to delete and click **Delete**.
 - You can use the **Reset** button to clear all extensions from the list, or to select a different file extension list (Archive or Standard).

4 In the **Action** area, select one of the following options:

- **Allow** – Select this option if you want to explicitly allow the file extensions and/or MIME type that you specified in this window. (Virus scanning will not be performed.)
- **Deny** – Select this option if you want to explicitly deny the file extensions and/or MIME type that you specified in this window. (Virus scanning will not be performed.)
- **Virus/Spyware Scan** – Select this option if you want to perform virus scanning on the file extensions and/or MIME type that you specified in this window. If no viruses are detected, the file will be allowed through the system.

Configuring the Default filtering rule action

The Default filter rule is a catch-all rule designed to occupy the last position in your rule table.

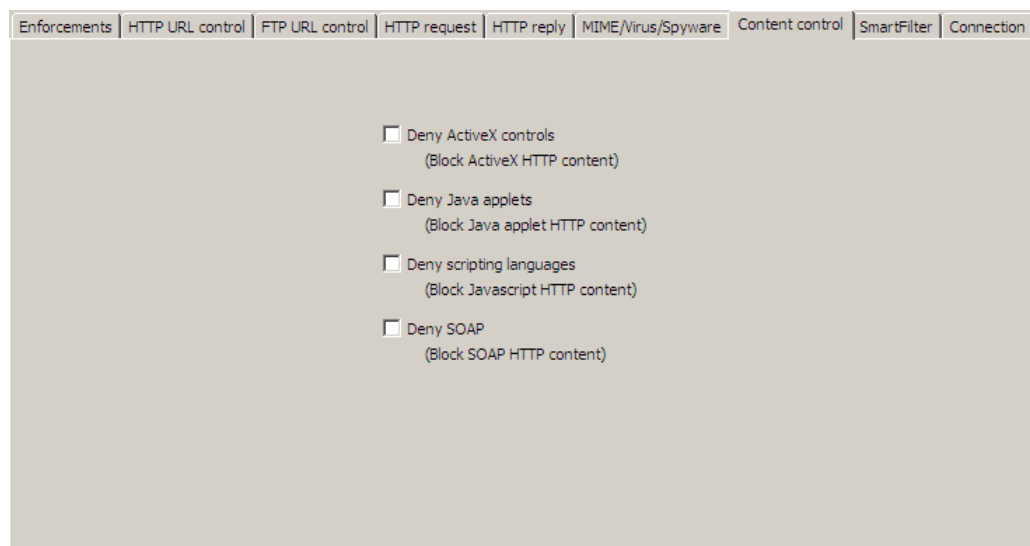
To modify the default action for the default MIME filtering rule:

- 1 Select the default rule in the table and click **Modify**. The MIME Default Action window appears.
- 2 Select the appropriate action for this rule and then click **OK**.
 - **Allow** – The default rule is initially configured to *allow* all data that does not match other filter rules. If you leave the default rule as an allow rule, you must create filter rules that require virus scanning or explicitly deny any MIME types that you do not want to allow, and place them in front of the default allow rule.
 - **Deny** – If you prefer the default rule to *deny* all data that did not match a filter rule, you must create the appropriate virus scan and allow rules and place them in front of the default deny rule.
 - **Virus/Spyware Scan** – If you want to perform virus and spyware scanning for traffic that does not match any allow or deny filter rules you create, select this option. You will then need to create the appropriate allow and deny rules that will not require scanning.
- 3 Save your changes.

Configuring the HTTP/HTTPS: Content Control tab

Use the Content Control tab to configure filtering to deny certain types of embedded objects. Follow the steps below.

Figure 120 HTTP/HTTPS: Content Control tab



Note: If you are configuring an HTTP or HTTPS defense for type Server, you will only be allowed to select the **Deny SOAP** option. If you are configuring an HTTP defense for type Client, the **Deny SOAP** option is not available.

To configure the Content Control tab:


- 1 Select the **Deny ActiveX Controls** check box to scrub ActiveX embedded objects from the web content.
- 2 Select the **Deny Java Applets** check box to scrub Java Applet objects from the web content.
- 3 Select the **Deny Scripting Languages** check box to scrub scripting languages from the web content.
- 4 Select the **Deny SOAP** check box to scrub SOAP embedded objects from the web content. In some cases, selecting this option can cause the entire page to be denied if it contains SOAP embedded objects.

Configuring the HTTP/HTTPS: SmartFilter tab

When SmartFilter is configured, use this window to determine whether requests will be rejected if the SmartFilter server is unavailable.

Note: SmartFilter is not actively supported in Sidewinder 7.x.

Figure 121 HTTP/HTTPS: SmartFilter tab



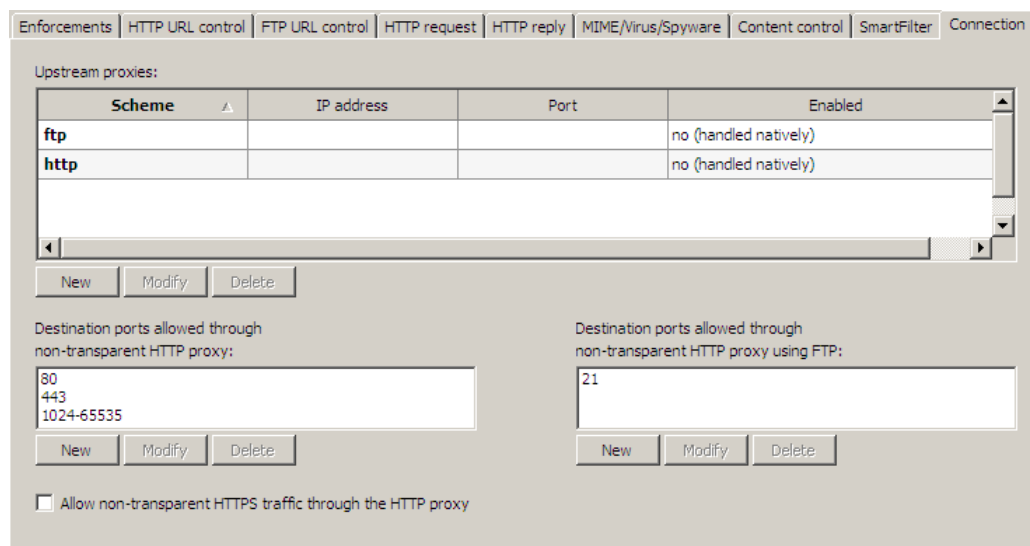
Select the **Reject all requests if SmartFilter is unavailable** check box to reject any requests that occur when the SmartFilter server on the firewall is unavailable.

For more information about configuring SmartFilter, see [Configuring SmartFilter for Sidewinder on page 147](#).

Configuring the HTTP/HTTPS: Connection tab

Use the HTTP/HTTPS Connection tab to configure whether to send traffic to an upstream proxy, and to define ports that non-transparent proxies can send traffic to.

Figure 122 HTTP/HTTPS: Connection tab



To configure the Connection tab:

1 To forward requests to upstream proxies:

a In the Upstream proxies area, click **New** and define the upstream proxy:

- **Scheme** – Enter the scheme of the requests to be forwarded. A scheme is the protocol identifier in the URI naming structure, for example, *gopher*.
- **IP address** – Enter the IP address of the upstream proxy where the request is being sent.
- **Port** – Specify the port of the upstream proxy where the request is being sent.
- **Enabled** – Select this check box to allow the defined scheme to be forwarded.

b Click **OK** and save your changes.

- HTTP and FTP traffic by default is handled locally by the Sidewinder. To forward HTTP or FTP requests to an upstream proxy, select the scheme and click **Modify**, then define the upstream proxy.
- HTTP requests can be transparent or non-transparent. If you allow transparent HTTP connections when using this option, the URL will be rewritten to contain an IP address rather than a hostname. If you allow transparent connections, you must first ensure that the upstream proxy server will accept an IP address.
- The HTTP scheme handles both HTTP and HTTPS, if non-transparent HTTPS is allowed through the proxy.
- Non-HTTP requests must be non-transparent so that the protocol can be identified. The HTTP service must be set to allow Non-Transparent or Both connection types.
- The connection request must match existing HTTP rules.
- An upstream proxy must be available.

2 To define allowable destination ports for non-transparent proxies: In the **Destination ports allowed through non-transparent HTTP proxy** area, click **New**. Specify a port, a port range, or select from pre-defined ports on the Edit a Port window.

- Pre-defined ports are 80, 443, 1024–65535.
- To modify an existing port entry, select the entry and click **Modify**.
- To delete an existing port entry, select the entry and click **Delete**.

Note: This table identifies the destinations the non-transparent proxy is allowed to send traffic to. If no destinations are identified, proxy connection will be denied. (HTTP and FTP connections will still be processed.)

3 [HTTP only] To allow non-transparent HTTPS traffic through the HTTP proxy, select the **Allow non-transparent HTTPS traffic through the HTTP proxy** check box. (The service must allow non-transparent connections.)

4 [HTTP only] To define allowable destination ports for FTP traffic through non-transparent proxies: In the **Destination ports allowed through non-transparent HTTP proxy using FTP** area, click **New**. Specify a port, a port range, or select from pre-defined ports on the Edit a Port window.

- The pre-defined port is 21.
- To modify an existing port entry, select the entry and click **Modify**.
- To delete an existing port entry, select the entry and click **Delete**.

Creating Mail (Sendmail) Application Defenses

Mail (Sendmail) Application Defenses are used in Sendmail rules. To configure Mail (Sendmail) Application Defenses, select **Policy > Application Defenses > Defenses > Mail (Sendmail)**.

Note: You must have Secure Split SMTP mail servers configured to use mail filtering.

Configuring the Mail (Sendmail): Control tab

Use this tab to configure filtering for sendmail services.

Figure 123 Mail (Sendmail): Control tab

The screenshot shows the 'Control' tab of the Mail (Sendmail) configuration interface. At the top, there are four tabs: 'Control', 'Size', 'Keyword Search', and 'MIME/Virus/Spyware'. The 'Control' tab is active. Below the tabs, there are two main sections: 'Enable Mail Services' and 'Enable Mail Filters'. In the 'Enable Mail Services' section, the 'Anti-Relay' checkbox is checked. Below this, a note states: 'NOTE: Detailed mail service settings may be edited in the sendmail configuration files.' In the 'Enable Mail Filters' section, there are three unchecked checkboxes: 'Size', 'Keyword Search', and 'MIME/Virus/Spyware'. At the bottom, there is a 'Rejected Mail Handling' section with two radio buttons: 'Discard' (selected) and 'Return To Sender'. A note below this section states: '(Applies to Keyword Search and MIME/Virus/Spyware Filters only)'.

Note: The Anti-Relay feature prevents your mailhost from being used by a hacker as a relay point for spam to other sites. This option is automatically enabled for all mail defenses and cannot be disabled.


To configure a Mail (Sendmail) Application Defense:

- 1 To enable (or disable) a particular type of filtering, you must select the appropriate check box in the **Enable Mail Filters** area. Once you enable a mail filter, you can configure it by selecting the appropriate tab. You cannot configure a mail filter unless you have selected it in this tab. The following filters can be enabled:
 - **Size** – The Size filter allows you to specify the maximum size for mail messages. To configure the Size filter once it has been enabled, select the Size tab. See [Configuring the Mail \(Sendmail\): Size tab](#).
 - **Keyword Search** – The Keyword Search filter allows you to filter mail messages based on the presence of defined key words (character strings). To configure the Keyword Search filter once it has been enabled, select the Keyword Search tab. See [Configuring the Mail \(Sendmail\): Keyword Search tab](#).
 - **MIME/Virus/Spyware** – The MIME/Virus/Spyware filter allows you to configure MIME, virus, and spyware filtering for e-mail messages. To configure the filter once it has been enabled, select the MIME/Virus/Spyware tab. See [Configuring the Mail \(Sendmail\): MIME/Virus/Spyware tab](#).
- 2 To specify how mail messages that are rejected should be handled, select one of the following options in the **Rejected Mail Handling** field:
 - **Discard** – Select this option if you want to discard rejected mail messages without notifying the sender.
 - **Return To Sender** – Select this option if you want to send a rejection notice to the sender.

Note: If a message is denied by the MIME/Virus/Spyware filter rules (configured in the MIME/Virus/Spyware tab), that message will be discarded without sending a rejection notice regardless of which option you select here.

Configuring the Mail (Sendmail): Size tab

Use this tab to configure size restrictions for a Mail (Sendmail) defense.

Figure 124 Mail (Sendmail): Size tab

The screenshot shows a configuration window for Mail (Sendmail) with four tabs: Control, Size, Keyword Search, and MIME/Virus/Spyware. The 'Size' tab is selected. Inside the window, there is a label 'Maximum Message Size (KB):' followed by a text input field containing the number '1024'.

The Size filter checks e-mail messages for the number of bytes the message contains, including the message header. A message is rejected if it is greater than or equal to the threshold size you specify when you configure a filter.

To configure the Size filter, in the **Maximum Message Size** field specify the maximum message size (in KB) that will be allowed to pass through the firewall. The default is 1024 KB. Valid values are 1–2147483647 KB.

Configuring the Mail (Sendmail): Keyword Search tab

Use this tab to configure the Sidewinder to perform a search for specified character set(s), or key words, within an e-mail message. The search scans the message's header and body sections.

- If the mail body contains MIME encoded attachments, the encoded attachments are scanned.
- If the filter finds a specific number of key word matches, the message is rejected.
- If the filter does not match a specific number of key words, it passes the message onto the next filter or to the intended recipient.

Figure 125 Mail (Sendmail): Keyword Search tab

The screenshot shows the 'Keyword Search' tab in the Mail (Sendmail) configuration window. The window has four tabs: 'Control', 'Size', 'Keyword Search' (selected), and 'MIME/Virus/Spyware'. On the left, there are two settings: 'Minimum Number of Phrase Matches Required for Rejection of Message:' with a text box containing '5', and 'Total Number of Phrase Matches to Verify Before Rejection:' with two radio buttons, 'Minimum' (selected) and 'All'. On the right, there is a 'Phrase List' table with three columns: 'Before', 'Phrase Text', and 'After'. The table is currently empty. Below the table are three buttons: 'New', 'Modify', and 'Delete'.

Select your key words carefully. For best results:

- Use spaces before and after each defined phrase.
- Create a comprehensive list of phrases instead of relying on wildcard-like searching.
- Note that key word searching is most reliable on MIME attachments with ASCII content-types. If dealing with non-ASCII types of attachments, false positives are likely if the length of the key words are short and the attachments are large.

Following these guidelines can decrease the chance of mistakenly rejecting a legitimate message.

To configure character sets to search for:

- 1 In the **Minimum Number of Phrase Matches Required for Rejection of Message** field, specify the number of key word matches that must be found in a message before it is rejected.
- 2 In the **Total Number of Phrase Matches to Verify Before Rejection** field, specify whether the filter will search the entire message for key words, or whether it will stop searching for key words if the minimum number of matches is met:
 - **Minimum** – Select this option if you want the filter to stop searching and fail the message if the minimum number of key word matches is met. This is based on the number that you enter in the previous step. The filter will reject a mail message once the minimum number of key words are matched.
 - **All** – Select this option if you want the filter to continue searching the message for key words after the minimum number of key word matches is met, for auditing purposes. After searching the entire message for key word matches, the message is rejected.
- 3 The Phrase List table provides the list of phrases that will be filtered for this Application Defense. The table contains three columns:
 - **Before** – This column indicates whether a space is required immediately before the specified phrase to match the filter. An asterisk (*) indicates that the phrase will not match unless there is a space immediately in front of the phrase.


- **Phrase Text** – This column lists each phrase for which the filter will search.
- **After** – This column indicates whether a space is required immediately after the specified phrase to match the filter. An asterisk (*) indicates that the phrase will not match unless there is a space immediately following the phrase.

To add a phrase, click **New**. To modify a phrase, highlight the appropriate row and click **Modify**. The Keyword Search: Phrase Edit window appears.

Configuring the Keyword Search: Phrase Edit window

Use this window to add or modify character strings (known as “key words”).

Figure 126 Keyword Search: Phrase Edit window



To configure a keyword search:

- 1 In the **Text** field, type the text you want to filter. The keyword search is not case sensitive. The character string must consist of at least two characters. You can include any printable character, as well as spaces.

Note: Some special characters, such as a space, will be displayed in the Key Word list using their hexadecimal equivalents.

You can also define a key word entry that consists partly or entirely of binary characters. The binary characters you want to search for are entered into the Key Word list using their hexadecimal equivalents. Each character must be preceded with a back slash (\). This distinguishes the character from a regular character. You can specify several characters in a row, but each character must be preceded by a back slash. You can also intermingle the binary characters with regular characters. For example, the following are valid entries in the Key Word list:

- `\ac\80\fe`
- `\ff\00\fb\40secrets`
- `password\df\01\04`

Valid hexadecimal characters are allowed immediately following a back slash. To use the back slash character as part of a key word entry, you must type a double back-slash (\\).

Note: The exception is \0a (the new line character). The filter will not detect a key word that contains this character unless it is the first character in the key word entry or unless the character is preceded by \0d (the line feed) character (e.g., \0d\0a).

- 2 If you want to require that there be white space directly in front of and/or after a key word, select the **Require whitespace immediately before phrase** and/or **Require whitespace immediately after phrase** check boxes. This prevents the filter from misidentifying character strings that innocently appear as part of another word.

For example, if you require whitespace before and after the key word “for,” words like “forest,” “formula,” “information,” and “uniform” will be allowed to pass through the filter, while the word “for” would not. If you do not require whitespace before and after the key word “for,” the “for” string within the word would match the filter and cause the message to be rejected (if the specified number of matches are found).

- 3 To add the new or modified key word, click **OK**.

Configuring the Mail (Sendmail): MIME/Virus/Spyware tab

Use the MIME/Virus/Spyware tab to configure MIME, virus, and spyware filtering services. The tab contains a rule table that displays any MIME/Virus/Spyware filtering rules that have been created. It also contains various virus/spyware scanning and handling configuration options.

Note: You must license and configure additional services before the MIME/Virus/Spyware filter rules you create will scan mail messages. See [Configuring virus scanning services on page 134](#).

Figure 127 Mail (Sendmail): MIME/Virus/Spyware tab

The screenshot shows the 'MIME/Virus/Spyware' tab in a configuration window. The window has tabs for 'Control', 'Size', 'Keyword Search', and 'MIME/Virus/Spyware'. The 'MIME/Virus/Spyware' tab is active, displaying a table of filter rules and configuration options on the right.

MIME/Virus/Spyware Filter Rules:

Pos	MIME Type	Extensions	Action
1	DEFAULT	DEFAULT	allow

Below the table are 'New', 'Modify', and 'Delete' buttons.

Scanner Behavior:

- ☒ Reject all files if scanning is unavailable
- ☐ Use heuristic scanning (scan for unknown viruses)

Infected File Handling:

- ☒ Discard infected files
- ☐ Repair infected files

Maximum Scan Size:

Scan file size limit (KB):

- ☐ Files over the scan limit will be allowed through unscanned
- ☒ Files over the scan limit will be rejected

SMTP Scanning:

- ☒ Full scan of entire mail message
- ☐ Discard entire message if denied or infected files are found

To configure MIME/Virus/Spyware properties for an Application Defense, verify that the Control tab's **MIME/Virus/Spyware** check box is selected and then follow the steps below.

Security Alert: If you want to perform virus and spyware scanning, you must create the appropriate MIME rules with **Virus/Spyware Scan** selected in the **Action** field. Rules that are configured only to allow or deny traffic based on rule criteria will not perform virus and spyware scanning. (See [Step 1](#) for information on configuring MIME/Virus/Spyware filter rules.)

- 1 Configure the appropriate MIME/Virus/Spyware filter rules in the **MIME/Virus/Spyware Filter Rules** table, as follows:
 - **Create a new filter rule** – To create a new filter rule, click **New** and see [About the MIME Rule Edit window](#).
 - **Modify an existing filter rule** – To modify an existing filter rule, select the rule you want to modify, and click **Modify**. See [About the MIME Rule Edit window](#). (If you are modifying the default MIME filtering rule, see [Configuring the Default filtering rule action](#).)
 - **Delete a filter rule** – To delete an existing filter rule, select the rule you want to delete and click **Delete**. You will be prompted to confirm your decision.
- 2 To reject all files in the event that scanning is not available, select the **Reject all files if scanning is unavailable** check box. If you select this option, the connection will be dropped if scanning is unavailable (for example, due to out-of-date virus data, an expired license, or a configuration error).
- 3 To scan files for viruses for which virus signatures do not exist, select the **Use heuristic scanning (scan for unknown viruses)** check box.

Note: Enabling this option may reduce virus scanning performance.
- 4 Determine how infected files will be handled in the **Infected File Handling** area as follows:
 - To discard infected files, select **Discard infected files**.
 - To remove the virus from the file and then continue processing the file, select **Repair infected files**.
- 5 Configure the **Maximum Scan Size** area.
 - a In the **Scan file size limit (KB)** field, specify the maximum file size that will be allowed in KB.
 - b Determine how files larger than the **Scan file size limit** will be handled by selecting one of the following:
 - **Files over the scan limit will be allowed through unscanned**
 - **Files over the scan limit will be rejected**
- 6 Configure the **SMTP Scanning** area.
 - Select **Full scan of entire mail message** if you want to perform scanning on the entire mail message (that is, the message with all of its MIME types is scanned as a single entity). A mail message is scanned only if one or more of its extensions match the MIME type/subtype settings on a filter rule with Virus/Spyware Scan selected.
 - Select **Discard message if denied or infected files are found** if you want to discard mail once a MIME/Virus/Spyware filter rule denies its attachment(s). If you select this option, files will either be discarded silently (sender is not notified) or returned to sender, as specified by the **Rejected Mail Handling** option selected on the Mail (Sendmail) Control tab.
 - If **Discard** is selected, the entire message is discarded if it contains a denied attachment.
 - If **Return To Sender** is selected, the message is sent on without the denied attachment.

About the MIME Rule Edit window

Use this window to add or modify MIME/Virus/Spyware filtering rules.

- Rules that are configured with an allow or deny action will allow or deny traffic based on the rule criteria that is defined for those rules. Allow and deny rules do not perform virus scanning. To perform virus scanning for traffic that matches a rule before it is allowed, you must specify **Virus/Spyware Scan** in the rule's **Action** field.
- Rules that specify both a MIME type/subtype and file extensions will allow or deny any traffic that matches either the MIME Type **or** a File Extension type. That is, the traffic does not need to match both criteria to match the rule.

Figure 128 Mime Rule Edit window

MIME/Virus/Spyware: MIME Rule Edit

MIME Type: *

MIME Subtype: *

File Extensions:

- ☒ Ignore Extensions (*)
- ☐ Archive Extensions
- ☐ Standard Extensions
- ☐ Custom

New Delete Reset

Action:

- ☒ Allow
- ☐ Deny
- ☐ Virus/Spyware Scan

OK Cancel Help

To add or modify MIME/Virus/Spyware filtering rules:

- 1 In the **MIME Type** drop-down list, select the MIME type for which you want to filter. If you select the asterisk (*) option, the filter rule will ignore this field when determining a match.
- 2 In the **MIME Subtype** drop-down list, select a subtype for the MIME type that you selected in the previous step (the available options will vary depending on the MIME type you selected). If you select the asterisk (*) option, the filter rule will ignore this field when determining a match.
- 3 In the **File Extensions** area, specify the type of file extensions that you want to filter:
 - **Ignore Extensions (*)** – Select this option to ignore extensions when determining a match.
 - **Archive Extensions** – Select this option to specify basic archive extensions (such as *.tar*, *.zip*, etc.) for the specified MIME types/sub-type.
 - **Standard Extensions** – Select this option to specify the standard file extensions associated with the selected MIME type/subtype. For example, if you select text in the **MIME Type** field, and HTML in the **MIME Subtype** field, the *.htm* and *.html* file extensions will appear in the standard list.
 - **Custom** – Select this option to create a custom list of file extensions for the selected MIME type/subtype.
 - To add a file extension to the list, click **New** and type the extension (*without* the leading period) that you want to add.
 - To delete a file extension, select the extension you want to delete and click **Delete**.
 - You can use the **Reset** button to clear all extensions from the list, or to select a different file extension list (Archive or Standard).
- 4 In the **Action** area, select one of the following options:
 - **Allow** – Select this option if you want to explicitly allow the file extensions and/or MIME type that you specified in the previous steps. (Virus scanning will not be performed.)
 - **Deny** – Select this option if you want to explicitly deny the file extensions and/or MIME type that you specified in the previous steps. (Virus scanning will not be performed.)
 - **Virus/Spyware Scan** – Select this option if you want to perform virus scanning on the file extensions and/or MIME type that you specified in the previous steps. If no viruses are detected, the file will be allowed through the system.

Configuring the Default filtering rule action

The Default filter rule is a catch-all rule designed to occupy the last position in your rule table.

To modify the default action for the default MIME filtering rule:

- 1 Select the default rule in the table and click **Modify**. The MIME Default Action window appears.
- 2 Select the appropriate action for this rule and then click **OK**.
 - **Allow** – The default rule is initially configured to *allow* all data that does not match other filter rules. If you leave the default rule as an allow rule, you must create filter rules that require virus scanning or explicitly deny any MIME types that you do not want to allow, and place them in front of the default allow rule.
 - **Deny** – If you prefer the default rule to *deny* all data that did not match a filter rule, you must create the appropriate virus scan and allow rules and place them in front of the default deny rule.
 - **Virus/Spyware Scan** – If you want to perform virus and spyware scanning for traffic that does not match any allow or deny filter rules you create, select this option. You will then need to create the appropriate allow and deny rules that will not require scanning.

Creating Mail (SMTP proxy) Defenses

Use the Mail (SMTP proxy) Application Defense to filter mail using the SMTP proxy and to conceal your internal mail infrastructure.

To configure Mail (SMTP proxy) Application Defenses, select **Policy > Application Defenses > Defenses > Mail (SMTP proxy)**.

Configuring the Mail (SMTP proxy): General tab

Use the General tab to hide your internal mail infrastructure and configure message destination and size options.

Figure 129 Mail (SMTP proxy): General tab

The screenshot shows the 'General' tab of the 'Mail (SMTP proxy)' configuration window. The 'Name' field is set to 'default' and the 'Description' field is 'default smtp application defense.'. There is a checkbox for 'Disable application defense filtering'. The 'General' tab is selected, with 'Commands' and 'Header filters' also visible. Under the 'General' section, there are checkboxes for 'Replace server's greeting with:', 'Replace server's FQDN with:', and 'Replace client's FQDN with:'. The 'Replace server's greeting with:' checkbox is checked, and the text 'Service ready' is entered in the adjacent field. There are also checkboxes for 'Pass server's reply text' and 'Verify client's FQDN'. The 'Max PDU size' is set to 1000 Bytes. In the 'Mail messages' section, there are radio buttons for 'Allow mail to any destination' (selected) and 'Only allow mail to defined destinations'. There are checkboxes for 'Limit message size:' (set to 0 Bytes), 'Limit number of recipients:' (set to 100), 'Banned mailbox characters:', and 'Add received header'. At the bottom, there are 'New', 'Modify', and 'Delete' buttons.

To configure the General tab:

- 1 [Optional] Select **Disable application defense filtering** to configure the SMTP proxy to ignore all options on this Application Defense, causing it to behave like a transport layer relay.
- 2 To modify the server's greeting text, select **Replace server's greeting with** and do one of the following:
 - To replace the server's greeting, type a replacement greeting in the field.
 - To remove the server's greeting, clear the field.

The default is to replace the greeting text with *Service ready*.

- 3 To replace the fully qualified domain name (FQDN) of an internal mail transfer agent (MTA), select one of the following options:
 - **Replace server's FQDN with** – Select this option and type an FQDN to replace the SMTP server's FQDN. This feature is commonly used with inbound redirect rules to hide an internal email server's domain name.
 - **Replace client's FQDN with** – Select this option and type an FQDN to replace the SMTP client's FQDN. This feature is commonly used with outbound NAT rules to hide an internal email server's domain name.

Note: In SMTP connections, the MTA sending the message is considered the client while the MTA receiving the message is considered the server.

- 4 To allow human-readable reply text to pass from the server to the client, select **Pass server's reply text**.

Note: Enabling this feature on outbound SMTP rules may reveal private network information.

- 5 To configure the allowed length of SMTP commands and responses, type a value in the **Max PDU size** field. Allowed values are 512 bytes to 64 kilobytes.

Note: This limit does not apply to data or authorization commands.

- 6 To require the client's IP address to match the domain specified in the client's HELO or EHLO command, select **Verify client's FQDN**. If enabled and the client's domain and IP address do not match, a 554 reply code is sent to the client.
- 7 In the Mail messages area, configure destination-based mail filtering.

Note: The SMTP proxy blocks messages that contain source routing information by default. To configure the proxy to allow these messages while stripping the source routing information, see [Configuring the SMTP proxy agent to strip source routing on page 188](#).

- **Allow mail to any destination** – Select this option to allow mail to any destination.
- **Only allow mail to defined destinations** – Select this option to specify the domains, IP address, and IP ranges to which the firewall will forward mail. The Sidewinder allows mail based on the contents of its RCPT TO: field; if the domain name portion of the RCPT TO: field matches a character string in the domain address list, the mail is allowed to pass.

To create or change a definition, click **New** or **Modify**. The Allowed SMTP Destination window appears. For information, see [About the Allowed SMTP Destination window](#).

To delete a definition, select the definition and click **Delete**.

- 8 To restrict the allowed size for mail messages, select **Limit message size** and type a value. Mail that exceeds the specified limit is rejected. Allowed values are 1 byte to 2 gigabytes.
- 9 To limit the number of recipients allowed per mail message, select **Limit number of recipients** and type a value. Allowed values are 1–100000 recipients.
- 10 To ban non-printable or potentially dangerous characters in mailbox addresses, type the desired characters in the **Banned mailbox characters** field. This field has no delimiters.

Note: Adding commonly used characters in this field is not recommended. For example, entering the character **o** blocks mail to all **.com** domains.

- 11 To configure the SMTP proxy to add an informational header to the beginning of messages it receives, select **Add received header**. This header advertises that the Sidewinder handled the message.

Note: This feature is intended to be used for troubleshooting or internal auditing purposes. It is not recommended to enable this feature on outbound SMTP rules because doing so may expose private network information.

About the Allowed SMTP Destination window

Use this window to allow a new mail destination or modify an existing mail destination.

Figure 130 Allowed SMTP Destination window

General: Allowed SMTP Destination

Allowed mail destination:

☒ Fully qualified domain name

☐ Include subdomains

Domain:

☐ IP address

IP address:

☐ IP range

Beginning of IP address range:

End of IP address range:

Match the entry to the destination's expected format in the *RCPT TO:* field. Identify an allowed SMTP destination by specifying one of the following:

- **Fully qualified domain name** – Select this option to specify a fully qualified domain name (FQDN).
 - In the **Domain** field, enter an FQDN, such as *example.com*.
 - To include the specified FQDN's subdomains, select **Include subdomains**. For example, if you allow mail to *example.com* and select this option, messages sent to *mail.example.com* are also allowed.

Tip: This is the most reliable option, as most destinations in the *RCPT TO:* field are formatted as a domain name.

- **IP address** – Select this option to specify a single IP address. In the **IP address** field, enter the destination as a valid IP address. To find the IP address for a host name, type the name and click **DNS Lookup**.
- **IP range** – Select this option to specify an address range. In the **Beginning of IP address range** and **End of IP address range** fields, specify the range of addresses that are allowed. To find the IP address for a host name, type the name and click **DNS Lookup**.

Configuring the Mail (SMTP proxy): Commands tab

Use the Commands tab to specify which SMTP commands are allowed.

Figure 131 Mail (SMTP proxy): Commands tab

The screenshot shows the 'Commands' tab of the 'Mail (SMTP proxy)' configuration window. At the top, there are fields for 'Name' (default) and 'Description' (default smtp application defense.), and a checkbox for 'Disable application defense filtering'. Below these are three tabs: 'General', 'Commands' (selected), and 'Header filters'. The 'Commands' tab is divided into two main sections: 'Allowed extensions' and 'Relayed commands'. The 'Allowed extensions' section contains a list of SMTP extensions with checkboxes: AUTH (checked), CHUNKING (checked), ETRN (checked), EXPN (unchecked), SIZE (checked), STARTTLS (unchecked), and VRFY (unchecked). Below this list are 'Select All' and 'Deselect All' buttons. The 'Relayed commands' section contains a table with columns 'Use', 'Command', 'Extension', and 'Description'. The table has four rows: 'onex', 'x-exps', 'x-link2state', and 'xexch50'. Each row has a 'Use' checkbox, which is currently unchecked for all. Below the table are 'New' and 'Delete' buttons.

Use	Command	Extension	Description
<input type="checkbox"/>	onex	onex	
<input type="checkbox"/>	x-exps	x-exps	
<input type="checkbox"/>	x-link2state	x-link2state	
<input type="checkbox"/>	xexch50	xexch50	

To configure the Commands tab:

- 1 [Optional] Select **Disable application defense filtering** to configure the SMTP proxy to ignore all options on this Application Defense, causing it to behave like a transport layer relay.

- 2 In the **Allowed extensions** area, select the SMTP extensions to allow.

Note: If you allow *starttls* and a session includes that command, the Sidewinder will no longer perform any command filtering for the rest of that session.

- 3 In the **Relayed commands** area, select the SMTP commands to relay.

To create a new command, click **New** and define the command that the SMTP proxy will relay.

Note: When a command selected in this list is encountered in a session, the Sidewinder will no longer perform any command filtering for the rest of that session.

About the Commands: Relayed command window

Use this window to define new commands that can be relayed.

Figure 132 Relayed command window

- 1 In the **Command** field, type the name of the command you want to add.
- 2 If necessary, complete the **Extension** field based on the following conditions:
 - If the command you are adding is defined by an SMTP extension, you must specify the extension name or SMTP clients will be unaware that the extension is supported.
 - If the command you are adding is not defined by an extension, leave this field blank.
- 3 If desired, type a description in the **Description** field.
- 4 Click **OK** and save your changes.

Configuring the Mail (SMTP proxy): Header filters tab

Use the Header filters tab to configure which mail headers are allowed.

Note: The SMTP proxy allows a maximum of 1000 headers per mail message.

Figure 133 Mail (SMTP proxy): Header filters tab

Name: Description:

☐ Disable application defense filtering

General | **Header filters**

☒ Allow all headers
☐ Allow selected headers only
☐ Strip selected headers

☒ Allow all header values
☐ Block messages with selected header-value pairs

Use	Header	Description
<input type="checkbox"/>	Bcc	Indicates the recipients of a message
<input type="checkbox"/>	Cc	Indicates the secondary recipients
<input type="checkbox"/>	Comments	A comment added to a message
<input type="checkbox"/>	Content-Description	A text description to be associated with the content
<input type="checkbox"/>	Content-Disposition	Control whether an entity should be displayed inline or as an attachment
<input type="checkbox"/>	Content-ID	A way to uniquely identify a message
<input type="checkbox"/>	Content-Language	Indicate the language used in the content
<input type="checkbox"/>	Content-MD5	Add a MD5 integrity check to the content
<input type="checkbox"/>	Content-Transfer-Encoding	Indicates which binary-to-text encoding is used
<input type="checkbox"/>	Content-Type	Describes the data contained in the message
<input type="checkbox"/>	Date	The date and time an email was sent
<input type="checkbox"/>	From	Indicates the sender of the message

New Delete

Use	Header	Value
-----	--------	-------

New Delete

To configure the Header filters tab:

- 1 [Optional] Select **Disable application defense filtering** to configure the SMTP proxy to ignore all options on this Application Defense, causing it to behave like a transport layer relay.
- 2 Configure mail header filtering by doing one of the following:
 - To perform no header filtering, select **Allow all headers**.
 - To allow only specific headers, select **Allow selected headers only** and then select the appropriate headers from the list.
 - To remove specific headers, select **Strip selected headers** and then select the appropriate headers from the list.

To add additional headers, click **New** and enter a name and description for a new header in the pop-up window. Only headers added in this manner can be deleted.

- 3 Configure message blocking based on header-value pairs:
 - To perform no message blocking based on header values, select **Allow all header values**.
 - To block messages with specific header values, select **Block messages with selected header-value pairs** and then click **New** to add new header values. See [About the Header filters: Header value window](#) for information on creating new header values.

Note: Header matches are case-insensitive.

About the Header filters: Header value window

Use the Header value window to define new header-value combinations:

- 1 Select the desired header from the **Header** drop-down list or type the name of the header.

Note: The Header drop-down list is populated from the headers defined on the Header filters tab.

- 2 Type the appropriate value in the **Value** field.

Matches made based on this value are case insensitive, and do not need to be full length matches. For example, entering *example* in this field would match *testexampledomain.net*.

- 3 Click **OK**. You return to the Header filters tab.

Creating Citrix Application Defenses

Use a Citrix Application Defense to configure advanced ICA proxy parameters.

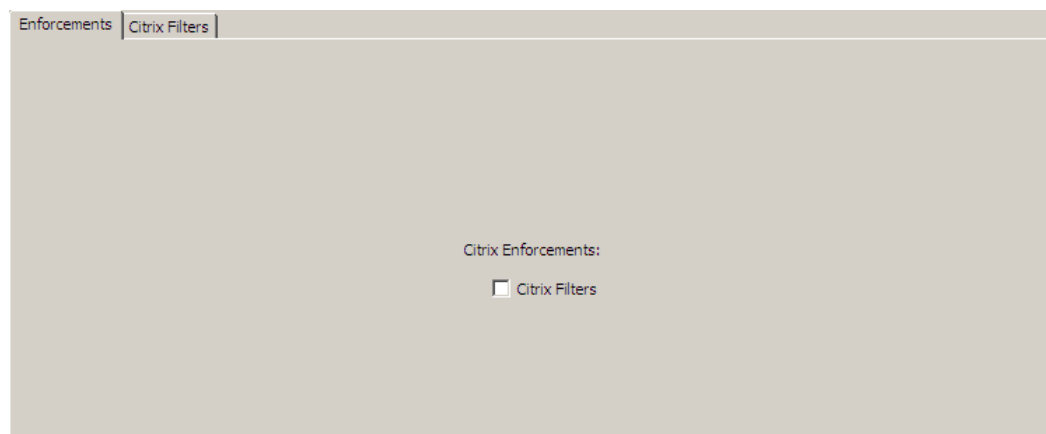
To configure Citrix Application Defenses, select **Policy > Application Defenses > Defenses > Citrix**.

Configuring the Citrix: Enforcements tab

Use the Enforcements tab to enable or disable Citrix filtering. The **Citrix Filters** check box must be selected in order to select and enforce values in the Citrix Filters tab.

To disable Citrix filtering, clear the **Citrix Filters** check box.

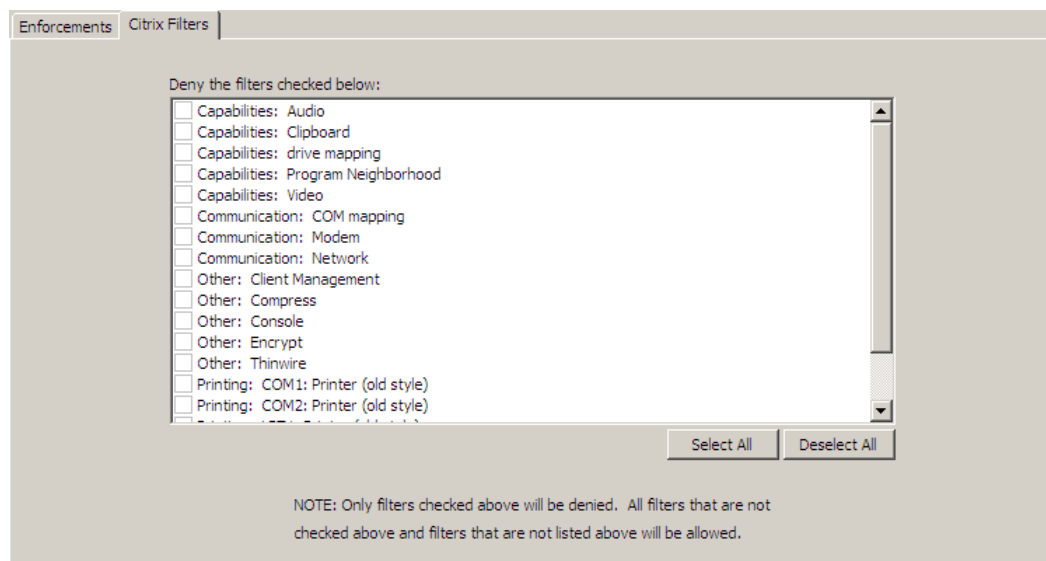
Figure 134 Citrix: Enforcements tab



Configuring the Citrix: Filters tab

Use the Citrix Filters tab to configure filtering properties for Citrix.

Figure 135 Citrix: Filters tab



To configure filters in Citrix, select the items that you want to deny. Each entry in the list represents a type of application or communication channel supported by Citrix. A check box will appear in front of types that will be denied. Clear the check boxes for the items you want to allow in Citrix.

To deny all of the types listed, click **Select All**. To allow everything (no filter restrictions), click **Deselect All**.

Creating FTP Application Defenses

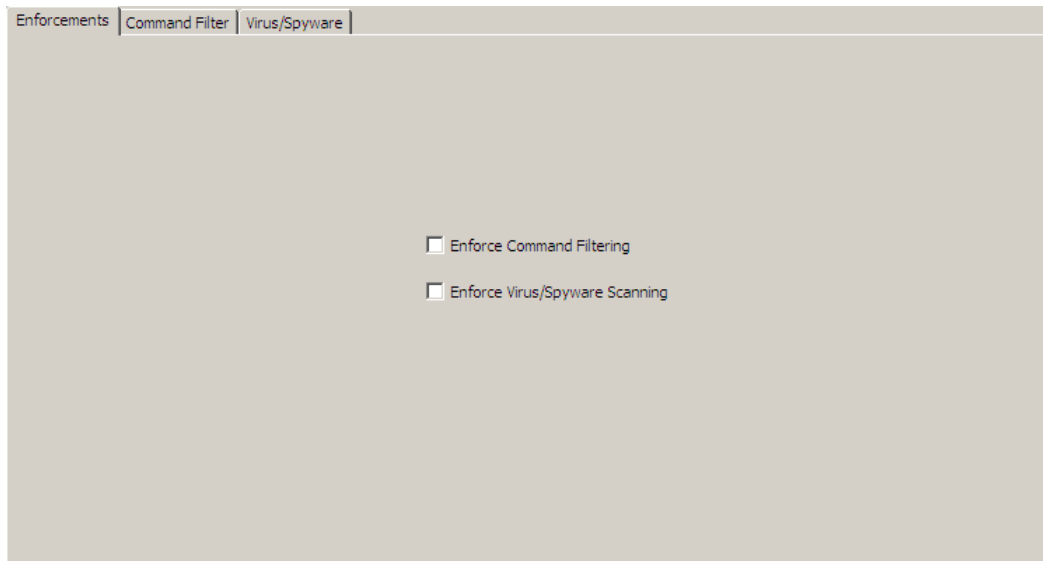
Use an FTP Application Defense to configure FTP permissions and the scanning of FTP files.

To configure FTP Application Defenses, select **Policy > Application Defenses > Defenses > FTP**.

Configuring the FTP: Enforcements tab

To enable or disable FTP feature enforcement tabs, you must first select the appropriate check box in the Enforcements tab. When you select the check box for a feature, that tab becomes enabled.

Figure 136 FTP: Enforcements tab



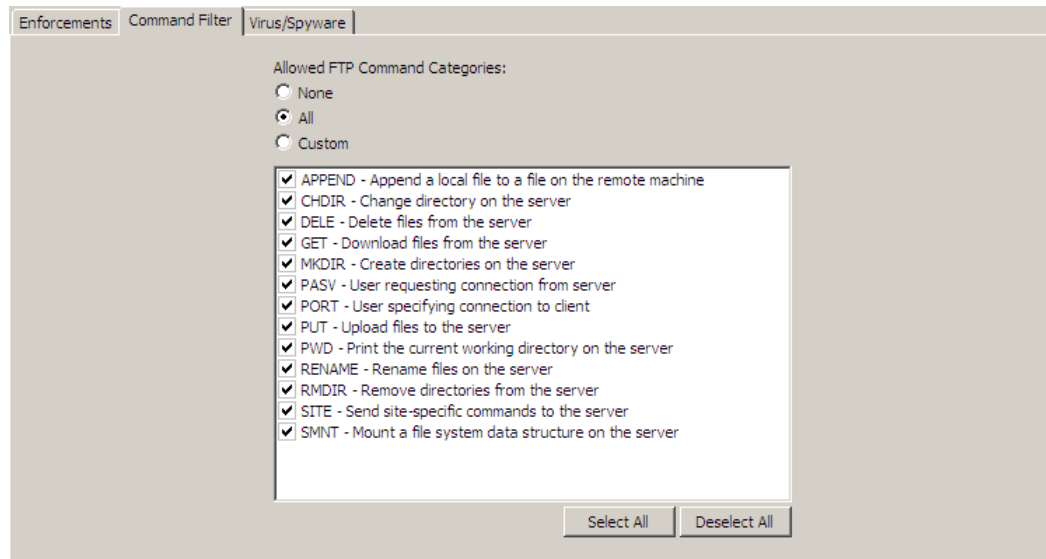
The following tabs can be enabled:

- **Enforce Command Filtering** – Use the FTP Command Filter tab to specify the categories of FTP commands that you want to allow your users to issue.
- **Enforce Virus/Spyware Scanning** – Use the Virus/Spyware tab to set the filtering parameters, such as infected file handling, which commands to scan, and which extensions to allow or deny.

Configuring the FTP: Command Filter tab

Use this tab to specify the categories of FTP commands that you want to allow your users to issue. The available FTP commands, as well as a description of each, are included in the **Allowed FTP Command Categories** area. For example, selecting *GET* allows the FTP commands necessary to download files from a server.

Figure 137 FTP: Command Filter



Select one of the following options:

- **None** – Select this option if you do not want to allow any FTP commands. (None of the check boxes will be selected.)
- **All** – Select this option if you want to allow all of the categories of FTP commands that are displayed. (All of the check boxes will be selected.)
- **Custom** – Select this option if you want to allow only certain FTP commands. To select the categories of FTP commands that will be allowed, click the appropriate check box. A check mark appears in front of commands that are allowed.

Note: If you select **None** or **All** and then make modifications to the commands, the **Custom** option will automatically become selected.

Configuring the FTP: Virus/Spyware tab

Use this tab to configure virus and spyware scanning services. The tab contains a rule table that displays any virus and spyware filtering rules that have been created. The tab also contains various virus and spyware scanning and handling configuration options.

Figure 138 FTP: Virus/Spyware tab

Note: You must license and configure scanning services before the Virus/Spyware filter rules you create will scan FTP traffic. See [Configuring virus scanning services on page 134](#).

To configure the Virus/Spyware tab:

- 1 Configure the appropriate virus and spyware filter rules in the **Virus/Spyware Filter Rules** table, as follows:
 - **Create a new filter rule** – To create a new filter rule, click **New**. See [Configuring Virus/Spyware filtering rules](#).
 - **Modify an existing filter rule** – To modify an existing filter rule, select the rule you want to modify, and click **Modify**. See [Configuring Virus/Spyware filtering rules](#). (If you are modifying the default filtering rule, see [Configuring the Default filtering rule action](#).)
 - **Delete a filter rule** – To delete an existing filter rule, select the rule you want to delete and click **Delete**.
- 2 To reject all files in the event that scanning is not available, select the **Reject all files if scanning is unavailable** check box. If you select this option, the connection will be dropped if scanning is unavailable (for example, due to out-of-date virus data, an expired license, or a configuration error).
- 3 To scan files for viruses for which virus signatures do not exist, select the **Use heuristic scanning (scan for unknown viruses)** check box.

Note: Enabling this option may reduce virus scanning performance.
- 4 Determine how infected files will be handled in the **Infected File Handling** area as follows:
 - To discard infected files, select **Discard infected files**.
 - To remove the virus from the file and then continue processing the file, select **Repair infected files**.

5 Configure the **Maximum Scan Size** area.

- a** In the **Scan file size limit (KB)** field, specify the maximum file size that will be allowed in KB.
- b** Determine how files larger than the **Scan file size limit** will be handled by selecting one of the following:
 - **Files over the scan limit will be allowed through unscanned**
 - **Files over the scan limit will be rejected**


6 Determine which commands to scan by selecting one of the following options in the **Apply Filter Rules to FTP** area:

- **Uploads (PUT)** – Scan all files going to the FTP server.
- **Downloads (GET)** – Scan all files coming from the FTP server.
- **Uploads and Downloads (PUT, GET)** – Scan all files going to (put) and coming from (get) the FTP server.

Configuring Virus/Spyware filtering rules

Use this window to add or modify virus/spyware filtering rules.

Figure 139 Virus/Spyware: FTP Edit window



Note: Rules that are configured with an allow or deny action will allow or deny traffic based on the rule criteria that is defined for those rules. Allow and deny rules do not perform virus and spyware scanning. To perform virus and spyware scanning for traffic that matches a rule before it is allowed, you must specify **Virus/Spyware Scan** in the rule's **Action** field.

By default, a single *allow* rule is contained in the filter rule table. If you choose to leave the default allow rule as the last rule in your table (that is, all traffic that isn't explicitly denied will be allowed), you will need to configure the appropriate *virus/spyware scan* and/or *deny* rules and place them in front of the default allow rule. If you configure the default rule action to deny (that is, all traffic that is not explicitly allowed will be denied) you will need to configure the appropriate virus/spyware scan and/or allow rules and place them in front of the default deny rule.

To create Virus/Spyware filter rules:

1 In the **Action** area, select one of the following options:

- **Allow** – Select this option if you want to explicitly allow the file extensions that you will specify in the next step. (Virus and spyware scanning will not be performed.)
- **Deny** – Select this option if you want to explicitly deny the file extensions that you will specify in the next step. (Virus and spyware scanning will not be performed.)
- **Virus/Spyware Scan** – Select this option if you want to perform virus and spyware scanning on the file extensions that you will specify in the next step. If no viruses or spyware are detected, the file will be allowed through the system.

2 In the **File Extensions** area, specify the type of file extensions that you want to filter:


- **Perform action on all file extensions** – Select this option to perform the action specified in [Step 1](#) on all file extension.
- **Choose from predefined categories** – Select this option to perform the action specified in [Step 1](#) on file extensions associated with a particular category, such as image, audio, video, etc.
- To choose the file extension, select the appropriate category from the Category drop-down list. Check the desired extensions.
- **Custom List** – Select this option to create a custom list of file extensions.
- To add a file extension to the list, click **New** and type the extension (*without* the leading period) that you want to add. The file extension is added to the **Custom** file extension list.
- To delete a file extension, select the extension you want to delete and click **Delete**.
- You can use the **Clear** button to clear all extensions from the list.

3 Click **OK** to save the rule.

Configuring the Default filtering rule action

Use this window to modify the default action for the default virus/spyware filtering rule. The default filter rule is a catch-all rule designed to occupy the last position in your rule table.

Figure 140 FTP: MIME Default Action window



To configure the MIME default action:

- 1 Select the default rule in the table and click **Modify**. The Default Action window appears.
- 2 Select the appropriate action for this rule and then click **OK**.
 - **Allow** – The default rule is initially configured to *allow* all data that does not match other filter rules. If you leave the default rule as an allow rule, you must create filter rules that require virus scanning or explicitly deny any extensions that you do not want to allow, and place them in front of the default allow rule.
 - **Deny** – If you prefer the default rule to *deny* all data that did not match a filter rule, you must create the appropriate virus scan and allow rules and place them in front of the default deny rule.
 - **Virus/Spyware Scan** – If you want to perform virus and spyware scanning for traffic that does not match any allow or deny filter rules you create, select this option. You will then need to create the appropriate allow and deny rules that will not require scanning.

Creating IIOP Application Defenses


IIOP (Internet Inter-ORB Protocol) is a protocol that makes it possible for distributed programs written in different programming languages to communicate over the Internet.

To configure IIOP Application Defenses, select **Policy Configuration > Application Defenses > Defenses > IIOP**.

Configuring the IIOP: Filter tab

Use this tab to configure filtering properties for the Internet Inter-ORB Protocol (IIOP) proxy.

Figure 141 IIOP: IIOP Filter tab

The screenshot shows a configuration window titled "IIOP Filter". It contains two unchecked checkboxes: "Allow Bi-directional GIOP" and "Validate Content Format". Below these is a text field labeled "Maximum message size (PDU):" with the value "72000" entered.

To configure the Filter tab:

- **Allow Bi-directional GIOP** – Select this option to enable support for bi-directional 1.2 GIOP (General Inter-ORB Protocol).
- **Validate Content Format** – Select this option to filter the message encapsulated in the GIOP PDU (protocol data unit), and verify that the header content, message direction, and message length are valid for the GIOP message type identified in the GIOP header.

Note: The data in the GIOP header portion of the PDU is always validated.

- **Maximum message size (PDU)** – Enter the largest message allowed through the proxy. The default is 72000.

Creating T.120 Application Defenses


T.120 is a standard for real-time data conferencing. The T.120 Application Defense allows you to use T.120 applications such as Microsoft's NetMeeting application.

To configure T120 Application Defenses, select **Policy > Application Defenses > Defenses > T120**.

Configuring the T.120: General tab

Use this tab to enable the T.120 filter, which is commonly used to allow Microsoft's NetMeeting. You cannot configure the T.120 properties unless you have selected the check box.

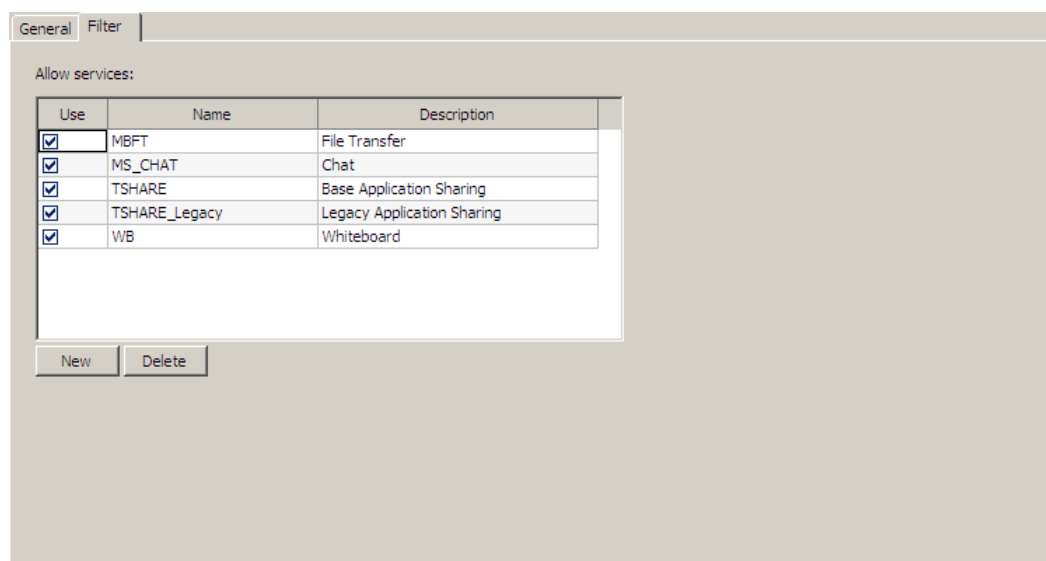
Figure 142 T.120: General tab



Configuring the T.120: Filter tab

Use this tab to specify which T.120 services you will allow your users to access. One of the more common T.120 applications is Microsoft's NetMeeting.

Figure 143 T.120: Filter tab



You can perform any of the following actions:

- **Enable or disable a service** – Select or clear the check box next to a service to enable or disable it.
- **Add an allowed service** – To add an allowed service, click **New**. In the pop-up window, enter a name and description for the new service and then click **OK**.
- **Delete a service** – If the service that you want to delete is enabled, clear the check box next to it and then save your changes. When saving is complete, select the service and then click **Delete**.

The following services are included by default:

- Whiteboard (T.126)
- File Transfer (T.127)
- Base Application Sharing (T.128)
- Legacy Application Sharing (T.128)
- Chat (Microsoft specific)

Note: These services cannot be deleted.


Creating H.323 Application Defenses

H.323 is a standard that provides support for audio and video conferencing across a shared medium such as the Internet. To configure H.323 Application Defenses, select **Policy > Application Defenses > Defenses > H.323**.

Configuring the H.323: General tab

Use this tab to enable the H.323 Filter.

Figure 144 H.323: General tab

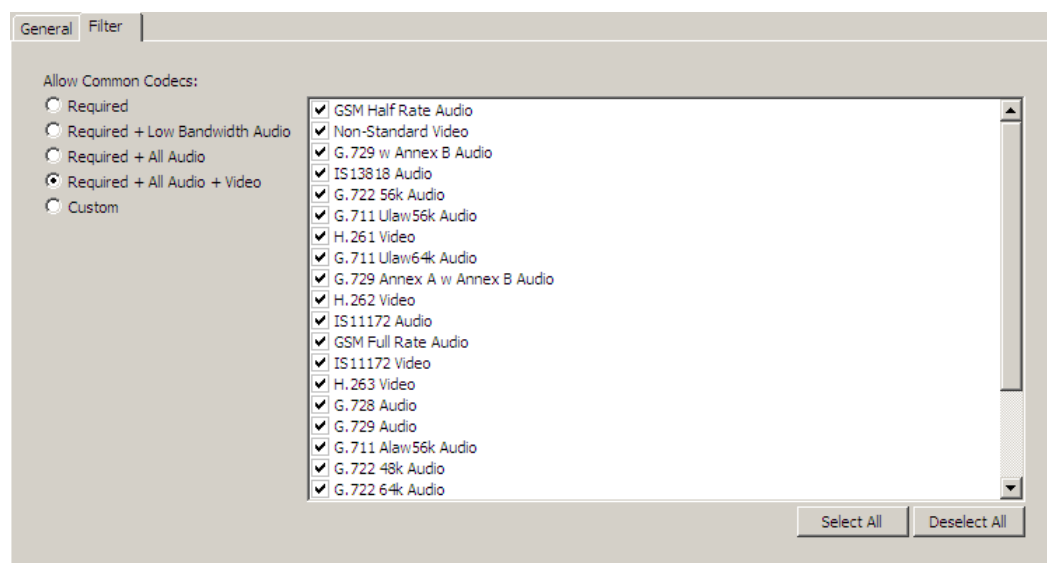
The screenshot shows a configuration window with two tabs: 'General' and 'Filter'. The 'General' tab is active. Inside the window, there is a checkbox labeled 'Enforce Permission Checking for H.323' which is currently unchecked. Below this checkbox, there is a label 'Maximum Call Duration:' followed by a text input field containing the value '86400' and the unit 'seconds'.

- 1 To enable H.323 configuration, select the **Enforce Permission Checking for H.323** check box. You cannot configure the H.323 properties unless you have selected the check box.
- 2 Enter a maximum call duration in seconds. The default is 86400 seconds.

Configuring the H.323: Filter tab

Use this tab to select H.323 codecs you will allow your users to access.

Figure 145 H.323: Filter tab



You can select from the following options:

- **Required** – Select this option to allow only the codecs required by H.323 for compliance.
- **Required + Low Bandwidth Audio** – Select this option to allow the required H.323 codecs as well as low bandwidth options.
- **Required + All Audio** – Select this option to allow all H.323 codecs except the codecs that allow video.
- **Required + All Audio + Video** – Select this option to allow all available H.323 codecs.
- **Custom** – Select this option to specify which codecs you want to allow. To allow a codec, select the appropriate check box. A check mark appears in the corresponding check box when a codec is allowed.
- **Select All** – Click this button to select all of the H.323 codecs (all codecs will be selected).
- **Deselect All** – Click this button to clear all of the H.323 codecs.

Note: If you select an option other than **Custom** and then make modifications to the selected codecs, the **Custom** option will automatically become selected.

The following list provide an example of codecs commonly used by Microsoft's NetMeeting:

- **G.711** – The G.711 codec options can transmit audio at 48, 56, and 64 kB per second (kBps). Select this codec for audio that is being passed using high speed connections.
- **G.723** – The G.723 codec options determine which format and algorithm will be used for sending and receiving voice communications over a network. This codec transmits audio at 5.3 and 6.3 kBps, which will reduce bandwidth usage.
- **H.261** – The H.261 codec will transmit video images at 64 kBps (VHS quality). Select this codec for video that is being passed using high speed connections.
- **H.263** – The H.263 codec determines which format and algorithm will be used to send and receive video images over a network. This codec supports common interchange format (CIF), quarter common interchange format (QCIF), and sub-quarter common interchange format (SQCIF) picture formats. It is also a good match for Internet transmission over low-bit-rate connections (for example, a 28.8 kBps modem).

Creating Oracle Application Defenses


Use an Oracle Application Defense to configure continuous session monitoring to prevent spoofing and tunneling attacks while sessions are in progress for the SQL proxy.

To configure Oracle Application Defenses, select **Policy > Application Defenses > Defenses > Oracle**.

Configuring the Oracle: Enforcements tab

Use this tab to enable or disable Oracle service name checking. Service name checking allows you to restrict access to the SQL server by specifying which service names will be explicitly allowed. If service name checking is enabled, only sessions that match a service name specified in the Service Name (SID) tab will be allowed.

Figure 146 Oracle: Enforcements tab



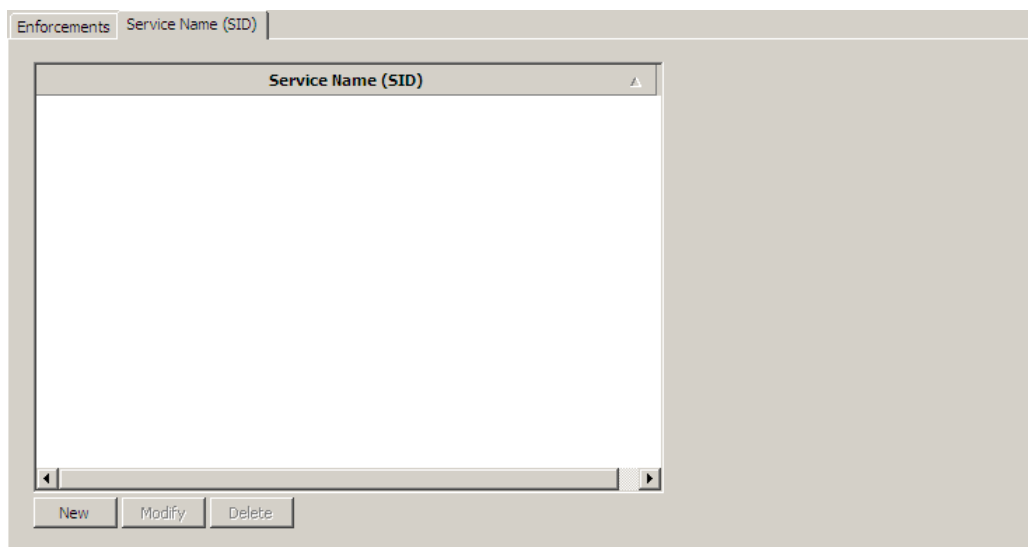
You cannot configure service name checking on the Service Name (SID) tab unless the **Enforce Service Name Checking** check box is selected. When this check box is selected, the values you configure in the Service Name (SID) tab will be enforced.

To disable service name checking, clear the **Enforce Service Name Checking** check box.

Configuring the Oracle: Service Name (SID) tab

Use this tab to configure which service names will be allowed access to the SQL server. If you do not specify any service names, service names will not be used in determining whether a session is allowed or denied.

Figure 147 Oracle: Service Name (SID) tab



You can perform the following actions:

- To configure a service name, click **New**. In the **Service Name (SID)** field, type the service name you want to add. The service name you enter must be an exact match (including capitalization) of the full service name that is in the Oracle tnsnames.ora file.
- To modify a service name, select the service name you want to modify, and click **Modify**. In the **Service Name (SID)** field, modify the service name.
- To delete a service name, select the appropriate service name and click **Delete**.

Creating MS SQL Application Defenses

The MS SQL Application Defense is not currently available. It is reserved for future features.

Creating SOCKS Application Defenses


Use the SOCKS Application Defense to configure advanced properties for the SOCKS proxy.

To configure SOCKS Application Defenses, select **Policy > Application Defenses > Defenses > SOCKS**.

Configuring the SOCKS: SOCKS 5 Filter tab

Use this tab to configure the type of SOCKS traffic that will be allowed when using the SOCKS5 proxy.

Figure 148 SOCKS: SOCKS 5 Filter tab




The following options are available:

- **Allow TCP SOCKS traffic** – Select this option to allow TCP traffic.
- **Allow UDP SOCKS traffic** – Select this option to allow UDP traffic.
- **Allow Both** – Select this option to allow both TCP and UDP traffic.
- **Enforce SOCKS 4 Filtering** – Select this option if you want to support SOCKS at version 4. (If this check box is not selected, you will not be able to pass traffic using SOCKS 4.)

Configuring the SOCKS: Connection tab

Use this tab to configure which ports will be open for the SOCKS proxy.

Figure 149 SOCKS: Connection tab



To define allowable destination ports for non-transparent proxies, click **New**, then specify a port, a port range, or select from pre-defined ports on the Edit a Port window.

To modify a destination port, select it in the list and click **Modify** and make your changes in the pop-up window.

To delete a destination port, select it in the list and click **Delete**.

Note: This table identifies which ports the SOCKS proxy is allowed to send traffic to. If no ports are identified, the proxy connection will be denied.

Creating SNMP Application Defenses


Use the SNMP Application Defense to configure advanced properties for the SNMP proxy.

To configure SNMP Application Defenses, select **Policy > Application Defenses > Defenses > SNMP**.

Configuring the SNMP: Filter tab

Use this tab to specify the SNMP version you want to configure.

Figure 150 SNMP: SNMP Filter tab

The screenshot shows the 'SNMP Filter' tab in a configuration window. At the top, there are two tabs: 'SNMP Filter' and 'SNMP v1'. The main area contains three radio button options: 'Allow SNMP v1 filtering' (which is selected), 'Allow SNMP v2c traffic', and 'Allow SNMP v1 and v2c traffic'. Below these options is a text input field labeled 'Maximum message size (PDU):' with the value '535' entered.

The options that you can configure within the subsequent SNMP tabs will vary depending on which option you select. The following options are available:

- **Allow SNMP v1 filtering** – Select this option to allow SNMP v1 traffic and configure object ID (OID) filtering. For information on configuring OID filtering for SNMP v1 traffic, see [Configuring the SNMP: v1 tab](#).
- **Allow SNMP v2c traffic** – Select this option to allow SNMP v2c traffic. OID filtering is not available for SNMP v2c traffic.
- **Allow SNMP v1 and v2c traffic** – Select this option to allow SNMP v1 and v2c traffic. OID filtering is not available when both SNMP v1 and v2c are allowed.

To set a maximum message size, type the maximum protocol data unit (PDU) allowed for a message in the **Maximum message size (PDU)** field. The default is 535.

Configuring the SNMP: v1 tab

Use this tab to configure Object ID (OID) filtering for SNMP v1 traffic.

Figure 151 SNMP: SNMP v1 tab

The screenshot shows the 'SNMP v1' configuration tab. On the left, under 'Options', there are three checkboxes: 'Allow Read Requests' (checked), 'Allow Write Requests' (unchecked), and 'Allow Notify Events' (unchecked). On the right, there is a checkbox for 'Enable OIDs Filtering' which is unchecked. Below this is an 'Actions' section with two radio buttons: 'Allow' (selected) and 'Deny' (unselected). A large empty table is positioned below the actions, and at the bottom are 'New', 'Modify', and 'Delete' buttons.

Note: Filtering is not available for SNMP v2c. If you selected **Allow SNMP v2c Traffic** or **Allow SNMP v1 and v2c Traffic** on the SNMP Filter tab, you cannot configure any options on this tab.

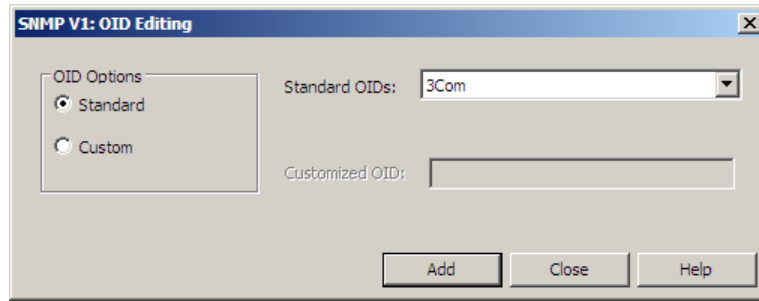
To configure the SNMP v1 tab:

- 1 In the **Options** area, determine the types of requests and events that the SNMP proxy will filter:
 - **Allow Read Requests** – Select this option to allow the **Get** and **Get Next** requests. (If you select SNMP v2c, this is automatically allowed.)
 - **Allow Write Requests** – Select this option to allow the **Set** request. (If you select SNMP v2c, this is automatically allowed.)
 - **Allow Notify Events** – Select this option to allow v1 traps. (If you select SNMP v2c, this is automatically allowed.)
- Note:** Additional SNMP requests are not supported in SNMP v1.
- 2 Select the **Enable OIDs Filtering** check box to configure object IDs (OIDs) for the SNMP proxy. OIDs are a unique, numeric representation of a device within the SNMP network.
- 3 In the **Actions** field, determine whether the list of OIDs that you define will be allowed or denied:
 - **Allow** – Select this option to allow only the OIDs that you specify in the table. All other OIDs will be denied.
 - **Deny** – Select this option to deny only the OIDs that you specify in the table. All other OIDs will be allowed.
- 4 To manage OIDs:
 - To add an OID to the table, click **New**. See [Configuring the SNMP v1: OID Editing window](#).
 - To modify an existing OID, select that ID and click **Modify**. See [Configuring the SNMP v1: OID Editing window](#).
 - To delete an existing OID, select that ID and click **Delete**.

Configuring the SNMP v1: OID Editing window

Use this window to add a new object ID (OID). You can select from the list of standard OIDs, or you can create your own OID using the custom option.

Figure 152 SNMP v1: OID Editing window



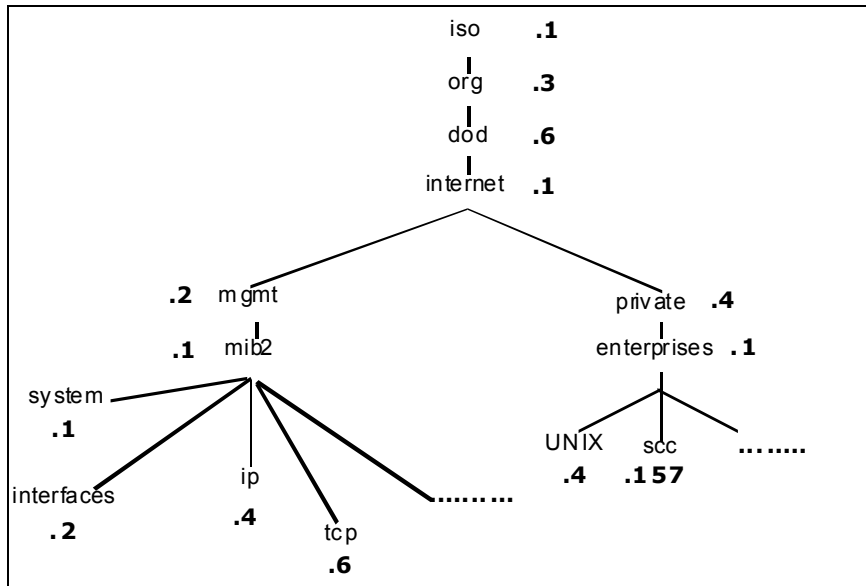
To add a new object ID:

- 1 In the **OID Options** area, select whether the OID will be **Standard** (pre-defined) or **Custom** (you determine and enter the OID manually).
 - If you select **Standard**, select the appropriate OID from the **Standard OIDs** drop-down list.
 - If you select **Custom**, type the OID number in the **Customized OID** field using the standard OID structure. The numbering scheme for each object is determined by the object's management information base (MIB) location, as shown in the figure below.

For example, the object ID for the SCC node in the private enterprise portion of the network would be **.1.3.6.1.4.1.1573**.

Note: The object ID will always begin with the pattern **.1.3.6.1**. For assistance on obtaining object IDs, visit the Internet assigned numbers authority web site at www.iana.org/assignments/enterprise-numbers or contact the appropriate vendor.

Figure 153 Example of OID numbering scheme



- 2 Click **Add** to add the OID to the table. Repeat these steps for each OID you want to add or modify.
- 3 Click **Close** to return to the SNMP v1 tab.

Creating SIP Application Defenses

Use the SIP Application Defense to configure media filtering, call duration, and peer types for the Session Initiation Protocol (SIP) proxy.

The Session Initiation Protocol (SIP) is defined by Internet Engineering Task Force (IETF) RFC 3261. SIP is an application-layer control (signaling) protocol for creating, modifying, and terminating sessions with one or more participants. These sessions include Internet telephone calls, multimedia distribution, and multimedia conferences. The SIP proxy agent provides standard functions such as filtering on source and destination hosts and burbs, and NAT and redirection. The SIP proxy is a protocol-aware, application-layer agent that examines SIP packets for correctness and adherence to site security policy. The SIP agent may be configured to prevent audio and/or video connections from being established via SIP.

SIP is used to locate a user agent and negotiate a multimedia session between user agents. A user agent is a device that terminates one side of a call (for example, the calling or answering phone). Once a session is negotiated, RTP is used to exchange the multimedia information between the user agents. The SIP proxy agent only examines the SIP traffic that negotiates the multimedia session. The RTP traffic itself is passed unexamined through the proxy. This traffic will make use of the Fast Path Sessions capability if the option is enabled for this service.

To configure SIP Application Defenses, select **Policy > Application Defenses > Defenses > SIP**.

Configuring the SIP: General tab

Use this tab to enable media filtering, to set the call duration, and to configure the types of peers that may participate in a SIP call.

Figure 154 SIP: General tab

The screenshot shows the 'General' tab of the SIP configuration window. At the top, there are two tabs: 'General' and 'Media Filters'. Below the tabs, there is a checkbox labeled 'Enforce Media Filtering'. Underneath this checkbox is a text field for 'Maximum Call Duration' with the value '86400' and the unit 'seconds'. Below the text field is a section titled 'Peer Types' which contains two radio button options. The first option is selected and is labeled 'The SIP peers must be user agents (e.g., phones)'. Below this option is the text 'This option is the most restrictive.' The second option is labeled 'The SIP peers can be routers (intermediaries that negotiate calls on behalf of user agents)'. Below this option is the text 'This option is the least restrictive.' At the bottom of the 'Peer Types' section is an information icon (a blue circle with a white 'i') followed by the text 'Some routers masquerade as user agents. See the Sidewinder product documentation for more information.'


You can perform the following actions:

- **Enable media filtering** – Select **Enforce Media Filtering** to enable SIP filtering. Use the Media Filters tab to select the desired filters.
- **Set the duration of calls** – Use the **Maximum Call Duration** field to enter the maximum number of seconds a call can last.
- **Configure peer types** – Select whether SIP calls may be negotiated by intermediaries.
 - Select **The SIP peers must be user agents** to require that all calls be negotiated by the SIP user agents of a call. The source and destination of each SIP message must be the SIP user agents (for example, SIP phones). Some SIP routers and gateways can masquerade as SIP user agents.
 - Select **The SIP peers can be routers** to allow SIP devices to negotiate calls on behalf of other SIP user agents. In this case, the source and destination of SIP messages processed by the proxy may differ from the SIP user agents that are participating in the call.

Configuring the SIP: Media Filters tab

Use this tab to configure media filters for an SIP session.

Figure 155 SIP: Media Filters tab



- Select **Audio** to allow audio streams via SIP.
- Select **Video** to allow video streams via SIP.

Use the **Select All** and **Deselect All** buttons to select or clear both options at once.

Creating SSH Application Defenses

Use the SSH Application Defense to configure advanced properties for SSH proxy rules. To configure your Sidewinder to pass SSH traffic using the SSH proxy, perform the following tasks:

- 1 Configure the appropriate policy rule using the SSH proxy as the rule's service.
- 2 Configure the SSH proxy agent properties. See [Configuring the SSH proxy agent on page 185](#).
- 3 Create an SSH Application Defense and apply it to the rule you created in [Step 1](#).

To configure SSH Application Defenses, select **Policy > Application Defenses > Defenses > SSH**.

Configuring the SSH: Channels tab

Use this tab to configure channel filtering for SSH connections.

Figure 156 SSH: Channels tab

To configure what content is allowed through SSH connections:

- 1 Configure what administration traffic is allowed:
 - Select **Allow remote shell execution** to allow terminal access to remote hosts.
 - Select **Allow remote command execution (includes SCP)** to allow commands to be sent to remote hosts.
Note: Enable this feature to allow Secure Copy (SCP) file transfers. Since SCP uses remote command execution to transfer files, it cannot function without remote command execution.
 - Select **Allow X11 forwarding** to allow UNIX-based X Window System traffic.
Note: Enabling this feature allows up to 10 concurrent X11 sessions.
- 2 Use the **Port forwarding (tunneling)** area to control port forwarding. Port forwarding allows the TCP/IP connection of another application to be redirected through an SSH tunnel. The following options are available:
 - **Allow local port forwarding** – Select this option to allow hosts to initiate port forwarding.
 - **Allow remote port forwarding** – Select this option to allow hosts to request that the remote host initiate port forwarding.

- 3** In the **Allowed SFTP operations** area, select the SSH File Transfer Protocol (SFTP) operations you want to allow:
- **None** – Select this option to deny all SFTP operations.
 - **Any** – Select this option to allow all SFTP operations.
 - **Selected from list** – Select this option to manually select which SFTP operations to allow.
- 4** In the **Allowed non-SFTP subsystems** area, select the non-SFTP subsystems you want to allow:
- **None** – Select this option to deny all non-SFTP subsystems.
 - **Any** – Select this option to allow all non-SFTP subsystems.
 - **Specified in list** – Select this option to specify non-SFTP subsystems to allow. For each subsystem you want to allow, click **New** and type the name of the subsystem in the pop-up window.

Configuring the SSH: Client Authentication tab

Use this tab to configure client authentication methods and the client greeting banner.

Figure 157 SSH: Client Authentication tab

The screenshot shows the 'SSH: Client Authentication' configuration window. It has four tabs: 'Channels', 'Client Authentication' (selected), 'Client Advanced', and 'Server Advanced'. The 'Allowed client authentication methods' section has two radio buttons: 'Any' and 'Selected'. The 'Selected' option is chosen, and a list box below it contains 'keyboard-interactive' and 'password', both with checked checkboxes. Below the list box are four buttons: 'New', 'Delete', 'Select All', and 'Deselect All'. The 'Client greeting' section has a large text area for entering a message, with up and down arrow buttons on the right side.

To configure the Client Authentication tab:

- 1** In the **Allowed client authentication methods** area, select the authentication methods you want to allow:
- **Any** – Select this option if you want to allow any authentication method that the client and server agree on.
 - **Selected** – Select this option to allow only the authentication methods that are selected in the list.
 - Select **keyboard-interactive** to allow authentication methods based on the keyboard-interactive method defined in RFC 4252.
 - Select **password** to allow password authentication.
 - Click **New** and type the method name to add a custom authentication method. Authentication methods added in this manner are the only methods that can be deleted.
- Note:** The **publickey** and **hostbased** authentication methods are not supported.
- 2** In the **Client greeting** area, type a message to be sent to the client immediately after a secure connection is established. Clear the field if you do not want to use a client greeting.

Configuring the SSH: Client Advanced tab

Use this tab to configure advanced options for client connections.

Figure 158 SSH: Client Advanced tab

The screenshot shows the 'SSH: Client Advanced' configuration window. It features four tabs: 'Channels', 'Client Authentication', 'Client Advanced' (active), and 'Server Advanced'. The 'Client Advanced' tab is organized into three main sections. The 'Encryption' section at the top contains two checked options: 'Rekey after specified bytes' (set to 1 Gigabytes) and 'Rekey after specified time' (set to 1 Hours). To the right of these are three rows of allowed algorithms: 'Cipher algorithms' (aes128-cbc, 3des-cbc, blowfish-cbc, cast128-cbc, arc4), 'MAC algorithms' (hmac-md5, hmac-sha1, hmac-ripemd160, hmac-ripem), and 'Key exchange methods' (diffie-hellman-group-exchange-sha256, diffie-hellma), each with an 'Edit' button. The 'Proxy host keys' section below has a 'Preferred type' dropdown set to 'DSA', and two dropdowns for 'DSA Key' (Default_DSA_Key) and 'RSA key' (Default_RSA_Key). The 'Known bugs handling' section at the bottom includes a 'Software version' text box set to 'OpenSSH_4.6' and a checkbox for 'Peer bugs considered to be fatal' which is checked, with 'Inability to rekey' selected in the dropdown.

To configure the Client Advanced tab:

- 1 In the **Encryption** area, configure the rekey options for the client connection. When a rekey is triggered, Sidewinder and the client renegotiate the shared key used to encrypt the session. Configure the following options:

- **Rekey after specified bytes** – Select this option and specify a data threshold. The client connection is rekeyed when the data threshold is reached.
- **Rekey after specified time** – Select this option and specify a time threshold. The client connection is rekeyed when the specified time elapses.

Note: If both options are selected, the first threshold that is reached triggers a rekey. When a rekey occurs, both counters are reset.

- 2 In the **Encryption** area, click the appropriate **Edit** button and configure the allowed algorithms and key exchange methods.

- **Cipher algorithms** – Cipher algorithms are used to encrypt the client connection. Click **Edit** to configure which algorithms are allowed and the order in which they are presented.
- **MAC algorithms** – Message Authentication Code (MAC) algorithms are used to verify the integrity of the client connection. Click **Edit** to configure which algorithms are allowed and the order in which they are presented.
- **Key exchange methods** – Key exchange methods are used to exchange private keys between the SSH proxy and the client. Click **Edit** to configure which methods are allowed and the order in which they are presented.

- 3 In the **Proxy host keys** area, use the following drop-down lists to configure the SSH host keys that the SSH proxy presents to clients:

- **Preferred type** – Select the type of key that the proxy presents to clients by default.
- **DSA Key** – Select the DSA key that the proxy presents to clients.
- **RSA key** – Select the RSA key that the proxy presents to clients.

Note: To manage SSH host keys, select **Maintenance > Certificate/Key Management** and then click the **SSH Keys** tab. See [Managing SSH keys on page 645](#) for more information.

- 4 In the **Known bugs handling** area, configure how the SSH proxy handles bugs in the client connection:

- **Software version** – Type the server name that the SSH proxy uses to represent itself to clients. Clients use this information to work around known bugs in SSH servers. The default is *OpenSSH_4.6*.

- **Inability to rekey** – Select this option to reject connections from clients that do not have the ability to rekey.

Configuring the SSH: Server Advanced tab

Use this tab to configure advanced options for server connections.

Figure 159 SSH: Server Advanced tab

The screenshot shows the 'SSH: Server Advanced' configuration window. It includes tabs for 'Channels', 'Client Authentication', 'Client Advanced', and 'Server Advanced'. The 'Server Advanced' tab is selected. The 'Encryption' section contains two checked options: 'Rekey after specified bytes' (set to 1 Gigabytes) and 'Rekey after specified time' (set to 1 Hours). Below these are three 'Edit' buttons for 'Cipher algorithms', 'MAC algorithms', and 'Key exchange methods'. The 'Allowed server key types' section shows 'Primary key' as DSA and 'Secondary key' as RSA. The 'Key checking policy' section features a slider between 'Strict' and 'Relaxed', currently positioned at 'Medium'. The 'Known bugs handling' section at the bottom shows 'Software version' as 'OpenSSH_4.6' and 'Peer bugs considered to be fatal' checked, with 'Inability to rekey' selected.

To configure the Server Advanced tab:

- 1 In the **Encryption** area, configure the rekey options for the server connection. When a rekey is triggered, Sidewinder and the server renegotiate the shared key used to encrypt the session. Configure the following options:
 - **Rekey after specified bytes** – Select this option and specify a data threshold. The server connection is rekeyed when the data threshold is reached.
 - **Rekey after specified time** – Select this option and specify a time threshold. The server connection is rekeyed when the specified time elapses.

Note: If both options are selected, the first threshold that is reached triggers a rekey. When a rekey occurs, both counters are reset.

- 2 In the **Encryption** area, click the appropriate **Edit** button and configure the allowed algorithms and key exchange methods.
 - **Cipher algorithms** – Cipher algorithms are used to encrypt the server connection. Click **Edit** to configure which algorithms are allowed and the order in which they are presented.
 - **MAC algorithms** – Message Authentication Code (MAC) algorithms are used to verify the integrity of the server connection. Click **Edit** to configure which algorithms are allowed and the order in which they are presented.
 - **Key exchange methods** – Key exchange methods are used to exchange private keys between the SSH proxy and the server. Click **Edit** to configure which methods are allowed and the order in which they are presented.
- 3 In the **Allowed server key types** area, use the following drop-down lists to configure the types of host keys that the SSH proxy accepts from servers:
 - **Primary key** – Select the preferred server key type.
 - **Secondary key** – Select the type of server key to accept if the primary server key type is not available.

Note: The allowed server key types cannot be the same. If you do not want to configure a secondary key type, you can select <None>.

- 4 In the **Key checking policy** area, use the slider to change the level of inspection applied to server host keys. If a server's host key does not meet the requirements set by the slider, the connection is denied.

Note: Key checking policy is enforced based on the trust level of the SSH server keys in the SSH known hosts database. The known hosts database is managed on the SSH proxy agent properties; see [Configuring the SSH proxy agent on page 185](#) for more information.

5 In the **Known bugs handling** area, configure how the SSH proxy handles bugs in the server connection:

- **Software version** – Type the client name that the SSH proxy uses to represent itself to servers. Servers use this information to work around known bugs in SSH clients. The default is *OpenSSH_4.6*.
- **Inability to rekey** – Select this option to reject connections to servers that do not have the ability to rekey.

Creating Packet Filter Application Defenses

Use the Packet Filter Application Defense to configure advanced properties for rules that use filter agents.

To use a Packet Filter Application Defense, you need to create a service that uses a filter agent, which is then applied to a rule. Services can be created using the following filter agents:

- **TCP/UDP Packet Filter** – Used for creating services for the TCP and UDP protocols.
- **FTP Packet Filter** – Used for creating services for the FTP protocol.
- **ICMP Packet Filter** – Used for creating services for the ICMP protocol.
- **Other Protocol Packet Filter** – Used for creating services for a number of protocols, such as GRE and AH.

Security Alert: We strongly recommend that you use a filter agent only for non-TCP/UDP protocols, such as PUP, GRE, AH, etc. Using a filter agent for a TCP/UDP protocol will, in most cases, severely degrade the effectiveness of the Sidewinder and will expose your network to security hazards.

For more information about creating services with filter agents, see [Create and modify services on page 152](#).

To configure Packet Filter Application Defenses, select **Policy > Application Defenses > Defenses > Packet Filter**.

Configuring the Packet Filter: General tab

Use this tab to specify the request rate and the audit parameters.

Figure 160 Packet Filter: General tab

The screenshot shows the 'General' tab of the Packet Filter configuration window. At the top, there are two tabs: 'General' and 'Advanced'. Below the tabs, there is a checkbox labeled 'Limit request rate to' followed by a text input field containing '1' and the unit 'requests/second'. Below this, there is a section titled 'Auditing'. Inside the 'Auditing' section, there is a line 'Audit the first' followed by a text input field containing '1', then 'denied requests every' followed by a text input field containing '1', and finally 'seconds'. Below this, there is another checkbox labeled 'Provide informational audits every' followed by a text input field containing '1' and the unit 'requests'. At the bottom of the 'Auditing' section, there is an information icon (a blue circle with a white 'i') and a text box that reads: 'The threshold for informational audits takes effect only when the audit level is set to verbose in the rule.'

You can perform the following actions:

- **Limit the number of requests that will be allowed per second in either direction** – Select **Limit request rate to** and enter the number of packets that you want allowed per second.
- **Specify how frequently the Sidewinder will generate audit records for deny rules** – Enter the number of denied requests and the time frame in the appropriate fields of the **Audit the first x denied requests every y seconds** area. Audit will be created for the first x occurrences in every y seconds. An additional audit event will be generated to record how many other audit events were suppressed.

For example, the audit is limited to generating an audit event for the first 1 occurrences for every 1 seconds. If the firewall stopped 100 netprobes in 1 second, one record would be generated for the first denial, and then another audit record would be generated stating that 99 occurrences were suppressed.

- **Specify the number of packets allowed by a rule before an informational audit is generated** – Select **Provide informational audits every [] requests**, and enter an appropriate number of requests. To limit auditing for this Packet Filter rule to only connection or session information, set the value to zero (0).

Configuring the Packet Filter: Advanced tab

Use this tab to select the response types you want to allow for a rule.

Figure 161 Packet Filter: Advanced tab



In the **Allowed control and error responses** area, select the response types that you want to allow for a rule. These selections control the ICMP messages generated by the rule's TCP/UDP traffic.

Note: If IPv6 is enabled on your firewall, the **IPv6 Allowed control and error responses** area also appears.

Configuring Application Defense groups

Application Defense groups are used in rules to specify advanced properties for service groups.

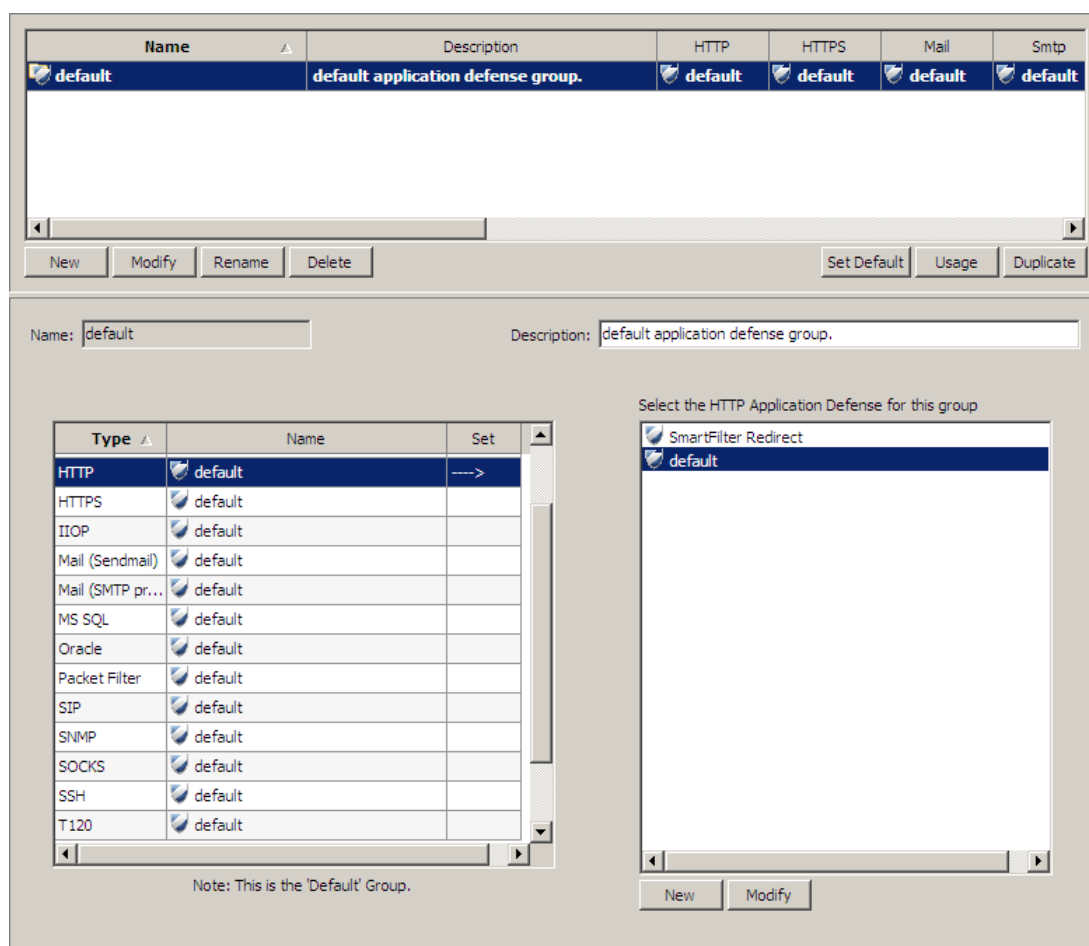
- When you create an Application Defense group, you select a single Application Defense from each category (for example, HTTP, HTTPS, FTP, etc.) to populate that Application Defense group.
- You set one Application Defense group as the default. The default group is used in all new rules using an Application Defense, unless you select a different Application Defense group in the Rules window.
- Only the Application Defenses that apply to that rule's services will be implemented in the rule.

Note: For more information on how Application Defense groups are used in a rule, see [Chapter 9, Rules](#).

To create an Application Defense group, select **Policy > Application Defenses > Groups**. The Application Defense Groups window appears.

Use this window to create and manage Application Defense groups.

Figure 162 Application Defense Groups window



- The upper pane lists all of the Application Defense groups that are currently configured. Each column shows which Application Defense is selected for the group.
- The lower pane lists each Application Defense category in the left table. When you select a category in the table, the available Application Defenses appear in the list on the right.

You can perform the following actions:

To create a new Application Defense group:

- 1 In the upper pane, click **New**. The New Groups Application Defense window appears.
- 2 Type a name for the group, then click **OK**. The group appears in the list in the upper pane.

- 3 In the lower pane, select an Application Defense for each category:
 - a In the left pane, select the appropriate Application Defense category. A list of available Application Defenses for that category appears in the right pane.
 - b In the right pane, select the Application Defense you want to associate with the selected category. The selected Application Defense appears in the Name column of the selected category.Perform this for each Application Defense category.
- 4 Save your changes. The selections in the Name column appear in the corresponding columns in the upper pane.

To modify an Application Defense group:

In the upper pane, select the appropriate Application Defense group. In the lower pane, make the desired changes for the Application Defense categories. (To make your changes in a pop-up window, select the Application Defense group and then click **Modify**.)

To rename an Application Defense group:

In the upper pane, select the appropriate Application Defense group. Click **Rename** and type a new name in the pop-up window.

To delete an Application Defense group:

In the upper pane, select the appropriate Application Defense group, then click **Delete**.

To make a group the default Application Defense group:

In the upper pane, select the appropriate Application Defense group, then click **Set Default**.

The default group is used in any rule using an Application Defense, unless you select a different Application Defense group in the Rules window.

To see which areas are using an Application Defense group:

In the upper pane, select the appropriate Application Defense group, then click **Usage**. A pop-up window appears listing the rule names that are currently using the selected group.

To duplicate an Application Defense group:

In the upper pane, select the appropriate Application Defense group, then click **Duplicate**. Type a name for the duplicated group in the pop-up window, then make the appropriate modifications to the duplicated Application Defense group.

To create or modify an Application Defense:

In the lower pane, select the appropriate Application Defense category. In the lower pane, click **New** or **Modify** and configure the Application Defense in the pop-up window.

9 Rules

Contents

[About rules](#)

[Using NAT and redirection in rules](#)

[Viewing and ordering rules and rule groups](#)

[Creating, modifying, and duplicating rules](#)

[Creating and modifying rule groups](#)

[Viewing and modifying rule elements](#)

About rules

Forcepoint Sidewinder policy is applied primarily by rules, which are made up of many elements. The table below shows the progression of a rule's creation using these elements and their corresponding chapters in this guide.

You are here in the Policy section	Use this chapter to...
Chapter 3, Policy Configuration Overview	understand the policy creation process.
Chapter 4, Network Objects and Time Periods	create or modify any network objects or time periods that will be used by rules.
Chapter 5, Authentication	create or modify authenticators that will be used by rules.
Chapter 6, Content Inspection	configure content inspection methods that will be used by rules.
Chapter 7, Services	create or modify services or service groups that will be used by rules.
Chapter 8, Application Defenses	create or modify Application Defenses that will be used by rules.
Chapter 9, Rules	create rules using the elements you created in the previous chapters in the policy section.

The basic elements that make up a rule include the service (which provides the protocol and port requirements), the source and destination, a time period, and authentication requirements; these are known as *condition elements*. If a packet matches all these parameters exactly, the firewall then refers to the rule's *action elements* for instructions on how to handle the packet. Action elements include the allow/deny/drop action, the audit level, the application defense settings, and the intrusion protection parameters.

The following sections describe the different rule elements and how to use them in a rule.

To see how these rule elements are used in a rule, see [Example of a simple rule](#).

Condition rule elements

This section describes the elements that a rule examines to see if a packet matches that rule. If the packet does not match all of these elements' values, the packet passes to the next rule. If the packet does match all of these elements' values, the rule handles the packet according to the action elements' values.

Services

Services determine a rule's protocol, port, and timeout values. There are three distinct service types:

- **Proxies** – Proxy services inspect traffic at the application layer. Proxy rules determine whether traffic will be allowed or denied using basic criteria such as protocol, port, source and destination address, but can also inspect the traffic to make sure it complies to its protocol's standards. Many proxy services also allow for advanced filtering and scanning services.

For more information, see [Configuring proxy agents and services](#).

- **Packet filters** – Packet filter services inspect traffic at the network and transport layers. Packet filter rules determine whether traffic will be allowed or denied using basic criteria such as protocol, port, source and destination address. Very little protocol and content inspection is available when using packet filter services. Because they are inherently less secure than proxies, packet filter services should be used only when necessary.

For more information, see [Configuring packet filter agents and services](#).

- **Servers** – Server services allow you to control access to firewall-hosted servers. Servers are typically used in management traffic rules where an administrator or another system needs to communicate directly with the firewall. Many of the server rules are created and enabled automatically. A few servers, such as the Sendmail server, allow for extensive configuration of its server properties, but most servers do not require changes to their default settings.

Sources and destinations

A rule's source and destination determine what can initiate traffic and what can respond to traffic that passes through, or into, the firewall. The source and destination consist of these properties:

- **Burb** – The area of the network containing the endpoint. This value can be a single burb, multiple burbs, a burb group, or multiple burb groups.
- **Endpoint** – The network object that can initiate or respond to connections or sessions. Network objects can be a domain, a Geo-Location (a way to identify the country of origin of an IP address), host, an IP address, a range of IP addresses, a subnet, a netmap (a way to map multiple IP addresses and subnets to alternate addresses without creating numerous rules), or a group that contains any combinations of those objects.
- **Network address translation (source)** – The address that replaces the original source address.

See [Using NAT and redirection in rules](#) for more information.

- **Address redirection (destination)** – The address that replaces the original destination address. Redirection can also change the original destination port to a different port.

See [Using NAT and redirection in rules](#) for more information.

Time periods

Time period rule elements determine the segment of time a rule is in effect. Time periods can be recurring, meaning the rule is active for the same time on the same day every week, or continuous, which means the rule is only active for a single period of time.

When creating a rule, you also have the option to set start and end times for rules. Delayed start times and scheduled end times are useful for making policy changes with minimal disruption to your production network.

TrustedSource

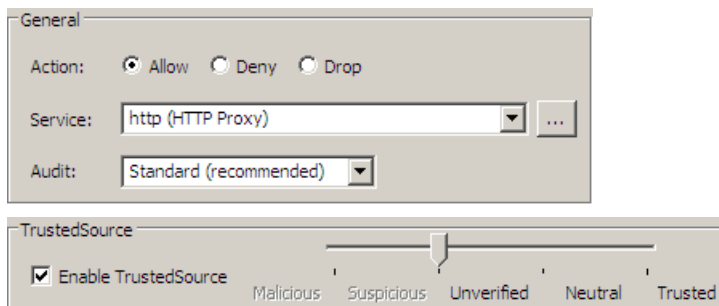
TrustedSource inspects network traffic and assigns it a reputation score. When a connection is examined by a rule with TrustedSource enabled, the firewall queries a TrustedSource server to get the reputation score of all IP addresses involved in the connection.

Traffic is not explicitly allowed or denied based on a TrustedSource score. The score is one of the elements in the rule that is examined for a match.

- In an *allow* rule, the **Unverified** to **Trusted** side of the TrustedSource slider is active by default. IP addresses with a good reputation will match this rule.


- If the reputation score is within the **Neutral** to **Trusted** range and all other elements in the rule match, the connection is allowed. No other rules are queried.
- If the reputation score is left of the **Unverified** to **Trusted** range, it is not a match. The connection is passed to the next rule.

Figure 163 TrustedSource on an allow rule



- In a **deny** rule rule, the **Suspicious** to **Malicious** side of the TrustedSource slider is active by default. IP addresses with a bad reputation will match this rule.
- If the reputation score is within the **Suspicious** to **Malicious** range and all other elements in the rule match, the connection is denied or dropped. No other rules are queried.
- If the reputation score is right of the **Suspicious** to **Malicious** range, it is not a match. The connection is passed to the next rule.

Figure 164 TrustedSource on a deny rule



A reputation is expressed in five classes:

- **Trusted** – The IP address is a source of substantial amounts of legitimate traffic.
- **Neutral** – The IP address is a source of legitimate traffic, but may send small amounts of unusual traffic or traffic requiring further inspection.
- **Unverified** – The IP address may be a legitimate sender, but data gathered to date has been either inconclusive or insufficient to make a firm reputation decision.
- **Suspicious** – The IP address has exhibited substantial suspicious behavior in the past, and connections should be treated with caution appropriate to the application protocol in question.
- **Malicious** – The IP address has a history of malicious behavior.

TrustedSource can be used for inbound and outbound rules with proxy or server services. It cannot be used for rules with filter services.

See [About TrustedSource](#) for more information.

Authentication

Authentication validates a user's identity before he or she is allowed to access a network service or server. Authentication works together with user groups to control who can access what services. Authentication can be used on rules controlling access to the firewall and through the firewall. Available authentication methods are: Password, Passport (single sign on), SafeWord®, Radius, iPlanet, Active Directory, OpenLDAP, Custom LDAP, and Windows Domain.

Action rule elements

After a rule determines that a packet matches its condition elements, the rule handles the packet according to the action elements' values.

Action

A rule's action determines what the firewall will do once it matches traffic to that rule. Options are:

- **Allow** – Permits the traffic to continue to its destination.
- **Deny** – Prevents the traffic from going through the firewall and sends the source a message that its request was rejected.
- **Drop** – Drops the packet in packet filter and UDP proxy rules. Closes the connection in TCP proxy rules.

Audit

Audit levels determine how much audit data a rule will generate on a per-rule basis. By default, all rules generate connection data that includes the packet's source, destination, and service. The amount of audit data generated can be increased to aid in troubleshooting or decreased to view errors only.

Application Defenses

Application Defenses determine advanced application-specific properties. They can be used with packet filter services, most proxy services, and the sendmail server service.

- Application Defenses for proxy services can be used to enforce RFC (Request for Comments) standards and allowed parameters. Configurable parameters include headers, commands, versions, and file sizes. Key inspection services, such as anti-virus/anti-spyware, SSL decryption, and web services management, are enabled in their respective proxy's Application Defense.
- Application Defenses for filter services can be used to control request and response rates, error and control messages, and the audit rate for denied filter rules.

Intrusion Prevention Systems (IPS)

The IPS area consists of both a signature group and a response mapping. The signature group identifies which signatures of known network-based intrusion attacks to compare to the packet. The response mappings indicate what to do if an attack payload in the packet matches an attack signature. Available options are to allow, deny, drop, or blackhole the offending packet.

Example of a simple rule

This section provides an example of a simple rule to help you better understand how the firewall uses a rule to determine whether to allow or deny a connection request, and how to handle allowed connections.

The following table lists the condition elements for a rule that permits any client in an internal burb to connect to any web server located in the external burb. Conditional elements are the elements that a rule examines to see if a packet matches that rule. [Figure 165 on page 262](#) shows where these settings are in the Rule window. The fields corresponding to the criteria described in the table are indicated in the figure.

There are also a number of action elements you can configure for each rule. After a rule determines that a packet matches its condition elements, the rule handles the packet according to the action elements' values. The action elements are whether or not to allow the connection or session, what amount of audit data to generate, if the address should be translated, what Application Defense settings to enforce, and if the traffic will be compared to a set of IPS signatures.

Table 25 Rule elements that determine if a packet will match a rule

Condition rule elements	Setting	Comments
Enable	Checked	Disabled rules do not process traffic.
Service	HTTP (HTTP Proxy)	This rule uses the default HTTP proxy service, which is for TCP traffic on port 80 with default timeout and expected connection values, and passes traffic transparently (browsers do not need to point to the firewall).
Source Burb	internal	Traffic will originate in the internal burb.
Source Endpoint	<Any>	Traffic can originate from any IP address in the internal burb.
Destination Burb	external	Traffic will be delivered to the external burb.
Destination Endpoint	<Any>	Traffic can be delivered to any IP address reachable via the external burb.
Authentication	Passport	Users must authenticate the first time they use this rule to connect to an external web server. Subsequent connection will be authenticated from a cache.

Figure 165 Screen shot of a basic rule with condition elements identified

Rules: New Proxy Rule

Name: ☒ **Enable**

Description:

General

Action: ☒ Allow ☐ Deny ☐ Drop

Service: ...

Audit:

Effective Times

Time period: ...

☐ Start on: ...

☐ Expire on: ...

Source

Burb: ...

Endpoint: ...

NAT: ...

☐ Preserve source port

Destination

Burb: ...

Endpoint: ...

Redirect: ...

Redirect port: ...

TrustedSource

☐ Enable TrustedSource

Malicious Suspicious Unverified Neutral Trusted

Neutral and trusted traffic will match this rule. (Range 14 to -255)

Inspection

Application Defense: ...

Full ☒ -- All configured settings of the application defense are enforced.

None ☐ -

IPS Signature group: ...

Response mapping: ...

Authentication

Authenticator: ...

Allow users in the following groups: ...

OK Cancel Help

Using NAT and redirection in rules

You can configure rules to perform NAT (network address translation) and redirection. NAT and redirection are essentially the same thing: replacing an original address with another specified address. NAT indicates that the firewall will rewrite the source address. Redirect indicates that the firewall will rewrite the destination address. The following sections give examples of when to use address translation and how to configure NAT and redirection in rules.

Understanding and configuring NAT


NAT refers to rewriting a packet's *source* address. When the firewall receives the packet, it removes the original source address and replaces it with the address or host name specified in the matching rule. The destination host is only aware of the translated address.

A common reason to use NAT is that your internal network uses private addressing that needs to be replaced by a publicly routable address. By default, all outbound rules are translated to use *localhost*. Localhost is a network object that automatically maps to the IP address of the specified burb, which is often the destination burb. Aliases are also frequently used in NAT.

Note: The localhost object cannot be used as the source or destination endpoint of a packet filter rule.

In the example shown in [Figure 166](#), a host on internal network 172.17.0.0 requires Telnet access to the external network 192.101.0.0. The IP address of a host on the privately addressed internal network should not be passed through the firewall; traffic sent from the internal network to the external network should appear as if it originated at the firewall's publicly routable IP address.

Figure 166 Example of network address translation



The associated outbound rule must translate the internal host address to the firewall's external address. Configure the rule's NAT information as follows:

Table 26 Outbound NAT rule

Source burb: internal	Destination burb: external
Source endpoint: 172.17.0.0 (internal subnet)	Destination endpoint: 192.101.0.2 (destination address)
NAT address: localhost	Redirect: <None>


Understanding and configuring redirection

Redirection refers to rewriting the *destination* address of the packet. The originating host sends the packet to one address, and then the firewall sends the packet to the specified redirection address. The original destination address is often the IP address of firewall's external burb or an alias assigned to that burb.

A common reason to direct traffic to one address and then redirect it to another address is when the internal object has a non-Internet routable address. Other uses include redirecting several different aliases to the same backend server for the purpose of data collection, and allowing authenticated users access to a protected server while redirecting all other uses to another server.

In the example shown in [Figure 167](#), an external network at 192.101.0.0 requires Telnet access to the internal host at 172.17.120.123. However, 192.101.0.0 is not allowed to directly route to the internal host. External hosts must initiate a Telnet connection to the firewall's external side.

Figure 167 Example of redirection



The associated inbound rule must rewrite the destination address to that of the internal host and forward the traffic onward. Configure the rule's redirection information as follows:

Table 27 Inbound redirect rule

Source burb: external	Destination burb: external
Source endpoint: 192.101.0.0 (source subnet)	Destination endpoint: 66.169.10.76 (destination address)
NAT address: <None>	Redirect: 172.17.120.123 (internal host)

Viewing and ordering rules and rule groups

Rules are the basis of your security policy. They determine what traffic will be allowed to pass through your firewall and what will be denied. To view or manage your rules, select **Policy > Rules**. The main Rules window appears.

Figure 168 The main Rules window

Name	Enabled	Action	Service	Source Burb	Source	Destination Burb	Destination	Application De
(1-14) Sidewinder Pol	<input checked="" type="checkbox"/>							
(1) Internet Services	<input type="checkbox"/>	Allow	Internet Services	internal	<Any>	external	<Any>	default
(2) VoIP SIP	<input type="checkbox"/>	Allow	sip	internal	<Any>	external	<Any>	default
(3) VoIP H.323	<input type="checkbox"/>	Allow	h323	internal	<Any>	external	<Any>	default
(4) NetMeeting	<input type="checkbox"/>	Allow	NetMeeting	internal	<Any>	external	<Any>	default
(5-5) DNS	<input checked="" type="checkbox"/>							
(6-8) Administrat	<input checked="" type="checkbox"/>							
(9-10) SmartFilter	<input type="checkbox"/>							
(11) Passport	<input checked="" type="checkbox"/>	Allow	ssod	<Any>	<Any>	<Any>	<Any>	Passport
(12) Deny All	<input checked="" type="checkbox"/>	Deny	<Any>	<Any>	<Any>	<Any>	<Any>	

This window provides an overview of your security policy. It is where you view rules, adjust rule order, and enable or disable rules. It is also the starting point for creating and modifying rules and rule groups.

Use the following sections to view and order your rules and rule groups:

- [Ordering rules within your policy](#)
- [About the default firewall policy](#)
- [Creating an alternate policy](#)
- [Using the main Rules window](#)
- [Customizing the main Rules window view](#)
- [Viewing and exporting your active policy](#)

Ordering rules within your policy

The order in which rules and nested groups appear in your security policy is significant. When the firewall receives a packet, it searches the enabled rules in sequential order (beginning with the first rule or nested group within the group, then the second, and so on). If the traffic does not match the first rule, it is forwarded on to the next rule. The first rule that matches all the characteristics of the connection request (service, source, destination, and so on) manages the connection. Once a rule match is found, the traffic is processed according to that rule and the search stops.

The following are guidelines for organizing and maintaining your security policy:

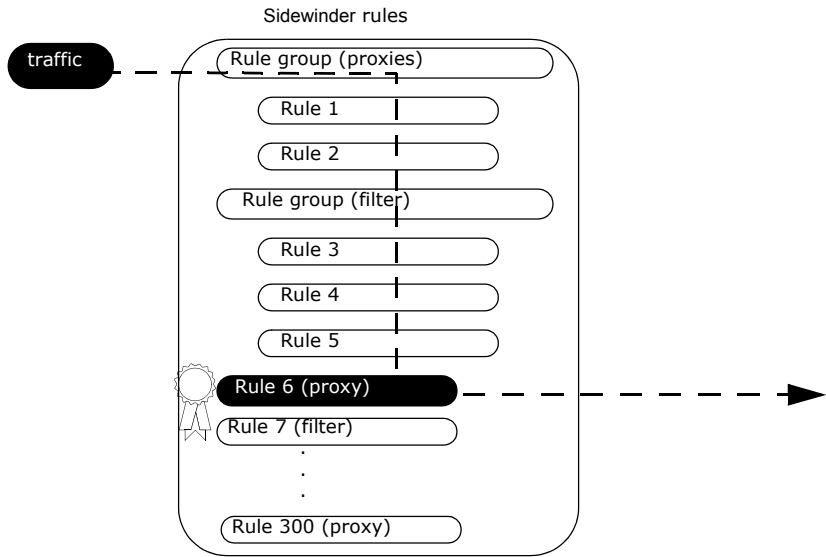
- Organize rules based on how frequently they are used. If you expect a rule to be widely used, such as a rule granting company-wide outbound HTTP access, put that rule near the beginning of your policy.
- Place specific rules before general rules. If you want to deny access to one group, such as contractors, while still allowing access for employees, put the rule denying contractors' access before the rule allowing employees' access.
- Audit your rules periodically. Look for rules that are no longer in use and rules that can be combined by using groups, such as service groups, netgroups, or application defense groups.

Caution: Do not disable or delete the login rules located in the Administration rule group, or place them below the Deny All rule. If these rules have been modified and you can no longer log in, see [Troubleshooting logging in](#) for assistance.

The default policy contains a Deny All rule at the end of the policy. This rule denies any traffic that reaches it. The rule itself is a reminder that any traffic that does not match a rule is automatically denied; even if the rule is deleted, the firewall denies any traffic that does not find an exact match in your security policy.

The following figure depicts first-match processing.

Figure 169 How traffic finds its matching rule



Note that this figure depicts rule processing at a high-level. Additional processing levels exist, but are not user-configurable.

Note: In general, proxy and filter rules can be listed in any order and will be processed sequentially. However, for proxy rules where the source or destination endpoint includes a domain object to be processed correctly, those rules must be placed after the last filter rule.

For example, suppose you want to allow access to FTP services on the Internet for all systems except those included in a netgroup called *interns*. The scenarios below illustrate both the incorrect and correct rule placement.

Incorrect placement of rules

The following shows a rule group order that is *incorrect* for this scenario.

Table 28 Incorrect rule placement

Rule 1:	Allow FTP service for all internal systems to all external systems.
Rule 2:	Deny FTP service for the netgroup <i>interns</i> to all external systems.

The first rule in the rule group allows all systems (via a wildcard) to use FTP and the second rule denies one particular netgroup.

Problem: When a system specified in the “interns” netgroup requests an FTP connection to somewhere in the Internet, the firewall will check rule 1. Because that rule allows all systems FTP service to the Internet, the firewall detects a match, stops searching the rule group, and grants the connection.

Correct placement of rules

To deny a particular netgroup, the deny rule should be placed *before* the allow rule. The correct way to order the rules in this rule group is as follows.

Table 29 Correct rule placement

Rule 1:	Deny FTP service for the netgroup <i>interns</i> to all external systems.
Rule 2:	Allow FTP service for all internal systems to all external systems.

Tip: As a basic guideline when configuring a rule group, place specific rules before any general (wildcard) rules.

The following scenario walks you through the basic process used by the firewall to process an outbound H.323 proxy connection request. This scenario assumes that the active rules consist of the following items:

- An enabled rule named **Internet Services**, which includes a service group that allows access to the most commonly used Internet services.
- An enabled rule group named **Administration**, which allows administrators to access the firewall.
- A disabled rule named **VoIP H.323** that allows voice over IP access via the H.323 proxy service.
- An enabled rule named **NetMeeting** that allows users to use audio and video conferencing components for NetMeeting®. This rule includes a service group that allows access to the H.323 and the T.120 proxy services.
- An enabled **Deny All** rule that will deny any requests that did not match any other rules.

The following steps outline the basic processing that takes place when an outbound H.323 connection request arrives at a firewall with the above rules in place:

- 1 An outbound H.323 request arrives at the firewall.
- 2 The request is processed by the first rule, which is the **Internet Services** rule. The request does not match the rule criteria.
- 3 The request is forwarded to the next rule, a rule group called **Administration**, and is inspected in sequential order by each rule contained within that group. No match is found in this rule group.
- 4 The request bypasses the **VoIP H.323** rule because the rule is disabled.
- 5 The request is forwarded to the next rule, the **NetMeeting** rule. A match is found (because the **H.323** proxy service is included in the service group used in this rule).
- 6 The request is processed according to the specifications in the **NetMeeting** rule. The request bypasses all other rules and groups contained in the active rules, and the request is granted.

About the default firewall policy

The firewall's default configuration creates a few commonly used rules for you. This policy includes outbound rules and management rules, but no inbound rules. The default rules that are deemed essential for basic management or standard functionality are enabled. During the Quick Start Wizard, you can choose to also enable a rule that allows access to a pre-configured group of commonly used Internet services. The other default rules are rules that you are likely to use at some point, but do not need to enable until required by your site's policy.

The initial enabled rules are listed in the following table:

Table 30 Initial active policy

Proxy rule name	Summary
dnsp (names vary)	Allow DNS traffic to proxy between indicated burbs. Which rules are created depends on the location of the DNS resolver IP addresses (internal burb, external burb, or assumed to be reachable by the default route) provided in the Network Information window in the Quick Start Wizard.
Admin Console	Allows administrators on the internal burb to connect to the firewall's internal interface using the Admin Console.
Login Console	Allows administrators to log in directly at the firewall using an attached keyboard and monitor.
Internet Services	<p>Allows users access to a pre-configured group of commonly used Internet services.</p> <p>Note: This rule is only enabled if you select Allow administrative and basic outbound Internet services during the Quick Start Wizard.</p> <p>The Internet Services rule regulates access to these proxies:</p> <ul style="list-style-type: none"> • FTP • HTTP • HTTPS • Ping • RealMedia • RTSP • Telnet
Passport	Allows authentication to the Passport server and facilitates the use of single sign-on authentication.
Deny All	Denies all connections from any source burb to any destination burb.

Creating an alternate policy

Many organizations need an alternate policy that is usually not in use but can be implemented quickly, such as a policy that limits inbound access if an attack is discovered. A good way to implement an alternative policy is:

- 1 Create a rule group for the alternate policy.
- 2 In that group, place all the rules needed to implement that policy. Groups can nest within groups. Be sure to create a Deny All rule as the bottom-most rule of the alternate policy.
- 3 Once the policy is finished, disable the policy by selecting the main rule group and clicking **Disable**.
- 4 When you need to use the policy, move the group to the top of the rule tree and enable it. The firewall begins enforcing your alternate policy.

Preparing policies for different disaster recovery scenarios can save valuable time in a crisis.

Using the main Rules window

This section provides information on using the main Rules window. You can perform several tasks from here, such as repositioning rules and rule groups, deleting rules and rule groups, and creating or editing rules. You can also re-arrange the columns layout, view a flat, non-nested list of all enabled rules and export that list, and use the Find feature to help you locate rules quickly.

Use the toolbar or right-click items in the table to perform the tasks in [Table 31](#).

Figure 170 The Rules window toolbar




Figure 171 Right-click menus in the Rules window

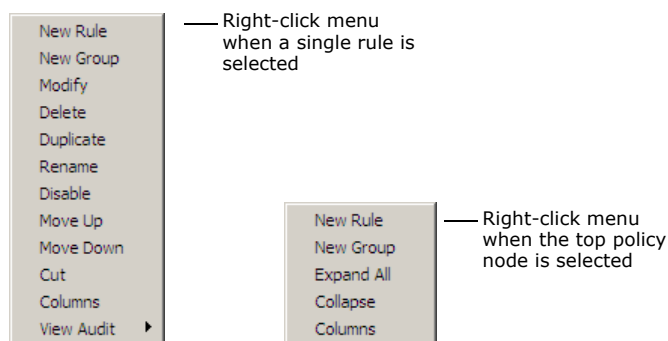



Table 31 Rules window tasks

Icon/ Menu item	Task
New Rule	Create a new rule by clicking New Rule . The New Rule window appears. See Creating, modifying, and duplicating rules for more information.
New Rule Group	Create a new rule group by clicking New Group . A window appears asking for a name and description for this group. You can add rules to a group two ways: <ul style="list-style-type: none"> Select the rules to group together and then create a new group. Create a new group and then move rules into it.
Modify	Modify a rule or rule group by double-clicking it or by selecting the rule and then clicking Modify . (Read-only administrators can click View to view a rule.) <ul style="list-style-type: none"> For rules, this opens the Modify Rule window. See Creating, modifying, and duplicating rules for more information. For rule groups, this opens the Modify Group popup, where you modify the group's description.
Delete	Delete a rule or rule group by selecting the item(s) to delete and clicking Delete . Deleting a rule group also deletes the rules in the group. If you do not want to delete a rule group member, move the rule out of the group before clicking Delete .
Cut/Paste	Cut and paste rules and groups to move items from one area of the rule tree to another. You can also move items by dragging and dropping them.
Duplicate	Duplicate a rule by selecting a rule and clicking Duplicate . The Duplicate Rule window appears, with "Copy of rule name" in the Name field. See Creating, modifying, and duplicating rules for more information. This task is useful for creating a rule that shares many properties with another rule. For example, you may need one FTP rule allowing access to one user group and one denying access to a different user group. Duplicating the first rule, then changing the action and user group, would be a quick way to accomplish this task.
Rename	Rename a rule or group by clicking Rename .

Rules

Viewing and ordering rules and rule groups

Table 31 Rules window tasks <Comment>(continued)

Icon/ Menu item	Task
View Audit	View all available audit data for a rule. You can also view audit data by right-clicking a rule and selecting the time frame: real time, last minute, last 15 minutes, last hour, or all available.
Enable/Disable	Enable or disable rules and rule groups by selecting one or more items and then clicking the appropriate icon.
Move Up/Move Down	Move rules and groups up or down one position by selecting the item and then clicking the appropriate arrow. To move a rule into a group, expand the group and then move the rule to the appropriate position. You can also move items by dragging and dropping them.
Find/Clear	Find items by entering a search term in the Find field and then clicking the magnifying glass  . The search is not case sensitive. Click the magnifying glass again to select the next instance of the search term. All columns are included in the search. Return to the full rule list by clicking Clear .
New Rule Template	Click New Rule Template to configure the rule template with custom default values. This template is used to populate the window that appears when you click New Rule .
Expand All/Collapse	Expand all rule groups so that all rules are visible or collapse the rules so that only the policy node is visible.
Active Rules	View all enabled rules in a flat list format. Use this window to sort and filter rules. See Viewing and exporting your active policy for more information.
Columns	Change the column view by clicking Columns . A window appears that allows you to choose which columns to display and in what order to display them. The Name column cannot be hidden or moved.


When you add a new rule or rule group, the placement is determined by the part of the rule tree that is selected when you click **New**. Possibilities are:

- If you select the policy node or do not have any items selected, the new rule or rule group is added to the bottom of the tree.
- If you select a group, the new rule or rule group is added to the bottom of that group.
- If you select a rule, the new rule or rule group is added directly below that rule.
- If you select multiple items, the position of the new rule or rule group depends on the last item selected.



Customizing the main Rules window view

The Column Selection window allows you to change what columns are displayed and in what order they appear. To access this window, select **Policy > Rules**, and then click **Columns**. The following window appears:

Figure 172 The Columns Selection window



Use this window to change how columns are arranged and which columns are displayed. The Name column will always be first, on the far left.


- **Hide a column** – Select one or more columns in the **Show these column in this order** list and then use the  arrow to move your selections to the **Available Columns** list. To move multiple consecutive entries, press the **Shift** key as you select the entries. To move multiple non-consecutive entries, press the **Ctrl** key as you select the entries.
- **Display a column** – To display a hidden column, select one or more columns in the **Available Columns** list and then use the  arrow to move your selections to the **Show these columns in this order** list. To move multiple consecutive entries, press the **Shift** key as you select the entries. To move multiple non-consecutive entries, press the **Ctrl** key as you select the entries.
- **Re-order the columns** – Select a single column and then use the **Up** and **Down** buttons to move it to a new location. You cannot move more than one column at a time.
- **Return to the default view** – Click **Default** to automatically display all columns in their original order.

When you finish changing the column view, click **OK** to return to the main Rules window.

Viewing and exporting your active policy

The Active Rules window lists your policy's enabled rules. This list can be exported in a comma separated value (CSV) format. To access this window, select **Policy > Rules** and click **Active Rules**. The following window appears:

Figure 173 The Active Rules window



Use the Active Rules window to view only the enabled rules. Position inconsistencies (for example, listing position 4 and then position 6) represent disabled rules, or enabled rules in a disabled group.

- **Sort the rules** – Click that column's header to sort the active rules based on the contents of a single column.
- **Filter** – Right-click a column's header to filter the active rules based on the contents of a single column.
- **Refresh** – Click the **Refresh** button in the upper-right corner to refresh the view to include rules created since this window opened.
- **View a rule in a full-window display** – Select a rule and click **View**.
- **Export the list** – Click **Export (csv)** to save this list as a .csv file. To change which columns are displayed in this file, adjust the displayed columns on the main Rules window.

When you have finished viewing the active rules, click **Close** to return to the main Rules window.

Creating, modifying, and duplicating rules

This section provides information on creating, modifying, and duplicating rules. It describes how to fill out this window, how fields can interact with each other, and valid values for fields.

To begin working with rules, select **Policy > Rules**. Several different actions provide access to a rule's parameters:

- Click **New** to start a new rule.
- Double-click a rule, or select it and then click **Modify**, to change an existing rule. (Read-only administrators can click **View** to view a rule.)
- Click **Duplicate** to create a duplicate of an existing rule. This is useful for creating a rule that shares many properties with another rule.

The following window appears:

Figure 174 New/modify/duplicate rule window

The screenshot shows the 'Rules: New Rule' dialog box. It has a title bar with 'Rules: New Rule' and a close button. The main area is divided into several sections:

- General:** Includes 'Name:' and 'Description:' text boxes, an 'Enable' checkbox, and a 'General' section with 'Action:' (radio buttons for Allow, Deny, Drop), 'Service:' (dropdown), and 'Audit:' (dropdown).
- Effective Times:** Includes 'Time period:' (dropdown), 'Start on:' (date and time), and 'Expire on:' (date and time).
- Source:** Includes 'Burb:' (dropdown), 'Endpoint:' (dropdown), 'NAT:' (dropdown), and a 'Preserve source port' checkbox.
- Destination:** Includes 'Burb:' (dropdown), 'Endpoint:' (dropdown), 'Redirect:' (dropdown), and 'Redirect port:' (text box).
- TrustedSource:** Includes an 'Enable TrustedSource' checkbox and a slider for 'TrustedSource' (Malicious, Suspicious, Unverified, Neutral, Trusted).
- Inspection:** Includes 'Application Defense:' (dropdown), 'IPS Signature group:' (dropdown), 'Response mapping:' (dropdown), and a 'Full' checkbox.
- Authentication:** Includes 'Authenticator:' (dropdown) and 'Allow users in the following groups:' (dropdown).

At the bottom right are 'OK', 'Cancel', and 'Help' buttons.

Use this window to enter all the information the firewall uses to identify and manage traffic. Each field's drop-down list contains all existing options for that field. If you know which options you want to use, type the appropriate entry or select it from the drop-down list. Burbs, burb groups, and authentication user groups support selecting multiple options. If you want to modify or search all available options or create a new entry, click **...**.

When updates are made to a rule, the window displays the user who last modified the rule along with the date and time (bottom left corner).

Fill in, modify, or view the following:

Rules

Creating, modifying, and duplicating rules

- 1 In the **Name** field, enter a name that helps identify the purpose of this rule. For example, the pre-configured rule that allows typical Internet services is called "Internet Services."

Valid values include alphanumeric characters, periods (.), dashes (-), underscores (_), and spaces (). However, the first and last character of the name must be alphanumeric. The name cannot exceed 256 characters. You can rename the rule later.

- 2 [Optional] In the **Description** field, enter any useful information for this rule (for example, a description of what makes this rule different from a similar rule).

- 3 Select the **Enable** option to enable this rule. All new rules are enabled by default. You can also change this setting on the main Rules window.

- 4 In the **Action** field, select what will happen to traffic when it matches this rule:

- **Allow** – (Default) Permits the traffic to pass. Since all traffic is denied by default, most rules you create will be allow rules.
- **Deny** – Denies the traffic and generates an audit message. It also notifies the initiator that the traffic was denied.
- **Drop** – Denies the traffic, but does not send a response to the initiator.

Note: Do not use a rule where the action is **Drop** and the service, source, and destination are set to **<Any>**. Such a rule would block traffic for servers on the firewall (such as DNS, NTP, or Admin Console). If you use Drop with a qualifier of **<Any>** for service, source, or destination, then be specific (do not use **<Any>**) for at least one of the remaining service, source, or destination fields.

- 5 In the **Service** field, select the service or service group this rule will allow or deny.

What you select here determines what values are considered valid for the rest of this window. For example, if you select a service that can use application defenses, the Application Defense field is populated with that service's application defense options.

Note: If you change your service selection, check your other selections as well, as the new service may use different options.

- 6 In the **Audit** field, set the audit level. Options are:

- **Standard (Recommended)** – (Default) This is the most common setting. It outputs major errors and informational messages.
- **Verbose (Most)** – Use this level when troubleshooting. This audit output is useful for detecting configuration issues.
- **Errors (Least)** – Use this level only if an issue with your system requires you to increase performance and reduce the size of your audit logs. Only errors are audited at this level.


See [Chapter 11, Auditing](#) for more information on audit.

- 7 In the **Effective Times** area, specify when this rule will be enforced by doing the following:

- a Select the time period during which this rule will be active. By default, all rules are always active.
- b If you want to start enforcing this rule at a specific date and time, select **Start on** and then set the date and time.
- c If you want to stop enforcing this rule at a specific date and time, select **Expire on** and then set the date and time.

8 In the **Source** area, specify where this rule's traffic can initiate:

- **Burb** – Select the burb or burbs where the source endpoint is located. You can select one or more burbs, or one or more burb groups.

You can select multiple burbs and/or burb groups by typing the names in a comma-separated list (for example, *internal, DMZ*) or by clicking  and selecting multiple options.

- **Endpoint** – Select the network object (for example, IP address, domain, netmap, etc.) that is allowed to initiate traffic.

Source and destination endpoints must have the same type of address—an IPv4 source can connect only to an IPv4 destination, and an IPv6 source can connect only to an IPv6 destination.

If you want this rule to match all endpoints in the selected source burb(s), select one of the following network objects:

- **<Any>** – This network object matches both IPv4 and IPv6 addresses.
- **<Any V4>** – This network object matches IPv4 addresses only. If IPv6 is not enabled on your firewall, selecting this endpoint ensures that this rule will not allow any traffic from IPv6 addresses if you choose to enable IPv6 in the future.
- **<Any V6>** – [Available only if IPv6 is enabled] This network object matches IPv6 addresses only.
- **NAT** – Select the network object that will replace the original source address as the traffic leaves the Sidewinder. By default, NAT is on and uses the IP address of the firewall's interface that matches the destination burb (*localhost*).

If using NAT, note the following:

- If this rule's Destination **Burb** field includes a virtual burb, do not set this field to *localhost*.
- If you selected a netmap in the Source **Endpoint** field, the appropriate NAT properties are automatically supplied based on the mapping configured for each IP address or subnet in that netmap. For more information on netmaps, see [About the Network Objects: Netmap window](#).
- [TCP/UDP packet filter service allowing source ports above 1024 only] If stateful inspection is enabled on this rule and you need to preserve the source port, you must specify an alias IP address or a subnet that contains at least one alias IP address.

View the service's **Service Properties** area to verify the service's source ports.

- [Conditional] **Preserve source port** – Check this field to translate the rule as follows: the source address is translated to the associated NAT address, but the source port will not be translated.

When using this option, the translated address is obtained one of two ways:

- If the port range included ports above 1023, this address must be an alias; it cannot be a native IP address. If the port range is below 1024, the address can be a native or *localhost*.
- From a pool of IP addresses. This requires that there be one or more alias addresses defined for the destination burb's interface and that the NAT field be set to include those addresses. The NAT field can be set to a single IP address or a subnet that includes the alias addresses. The total number of connections is therefore dependent on the number of alias addresses defined for that interface.


Caution: To use this feature with ports above 1023, you must have at least one alias configured for the destination burb's interface or traffic will not pass.

This field appears only when the selected service's agent is a filter and is most commonly used in rules handling IKE traffic when the related Security Association does not use NAT-T.

9 In the Destination area, select where this rule's traffic can go by configuring the following:

Note: When using redirection, match the destination burb to the destination endpoint, even if the redirect endpoint is in another burb.

- **Burb** – Select the burb or burbs where the destination endpoint is located. You can select a single burb, multiple burbs, a burb group, or multiple burb groups.

You can select multiple burbs and/or burb groups by typing the names in a comma-separated list (for example, *internal, DMZ*) or by clicking  and selecting multiple options.

- **Endpoint** – Select the network object (for example, IP address, subnet, netmap, etc.) to which this traffic is sent. Source and destination endpoints must have the same type of address—an IPv4 source can connect only to an IPv4 destination, and an IPv6 source can connect only to an IPv6 destination.

If you want this rule to match all endpoints in the selected destination burb(s), select one of the following network objects:

- **<Any>** – This network object matches both IPv4 and IPv6 addresses.
- **<Any V4>** – This network object matches IPv4 addresses only. If IPv6 is not enabled on your firewall, selecting this endpoint ensures that this rule will not allow any traffic to IPv6 addresses if you choose to enable IPv6 in the future.
- **<Any V6>** – [Available only if IPv6 is enabled] This network object matches IPv6 addresses only.
- **Redirect** – If the traffic needs be redirected to a different endpoint, the original destination redirects to the network object you select here.

If you selected a netmap in the Destination **Endpoint** field, the appropriate redirection properties are automatically supplied based on the mapping configured for each IP address or subnet in that netmap. For more information on netmaps, see [About the Network Objects: Netmap window](#).

- [Conditional] **Redirect Port** – This is the port to which the connection redirects. Note the following:
 - The default is blank. This means the port remains unchanged. Entering a 0 in this field also leaves the port unchanged.
 - Valid values are 1 – 65535 (inclusive).
 - This field is not available for all services.

10 In the TrustedSource area, do the following:

- Select **Enable TrustedSource** to use TrustedSource for this rule. The firewall queries a TrustedSource server to obtain a reputation score for all IP addresses involved in the connection.
 - You can whitelist objects to exempt them from TrustedSource queries. See [Using TrustedSource in rules](#) for more information.
 - Private IP addresses are not evaluated by TrustedSource or examined in rules (for example, 10.x.x.x, 172.16.x.x, 192.168.x.x).
 - TrustedSource cannot be enabled in rules with filter services.
- Move the slider to define what traffic will match the rule. The categories of traffic are Trusted, Neutral, Unverified, Suspicious, and Malicious.

Note: Default scores for reputation boundaries can be changed at **Policy > Application Defenses > TrustedSource**. See [Configuring TrustedSource](#) for more information.

Traffic is not explicitly allowed or denied based on a TrustedSource score. The score is one of the elements in the rule that is examined for a match.

- In an *allow* rule, the **Neutral** to **Trusted** side of the TrustedSource slider is active by default. IP addresses with good reputations match.
- In a *deny* or *drop* rule, the **Unverified** to **Malicious** side of the TrustedSource slider is active by default. IP addresses with bad reputations match.

11 In the Inspection area's **Application Defense** field, do the following:

- a** Select the application defense or group this rule will use to inspect this rule's advanced application-level content. The default is the Application Defense group currently set to the default.

Advanced content includes headers, commands, and filters. This is also where premium features, such as virus scanning and web filtering, are added to rules.

Some proxy services and some servers do not have configurable application defenses; this field will be grayed out when those services are selected. All filter services require the use of an application defense.

Note: Rules that use HTTPS Application Defenses with the **Decrypt Web Traffic** option enabled must have redirection configured.

- b** Move the slider to change the degree to which traffic is inspected:

- **Full** – (Default) All configured application defense settings are enforced.
- **Partial** – This prevents filtering and scanning, such as header filtering and virus scanning. Some protocol inspection is used as necessary to allow traffic to pass.
- **None** – This essentially disables defense inspection and greatly limits how deeply the traffic is inspected. Only disable defense inspection for troubleshooting purposes, or in very detailed rules created to allow non-standards compliant traffic into your site.


If the slider associated with the rule's Application Defense is set to None, services will act like a packet filter and some services may stop passing traffic typical of their protocol.

- Non-transparent functionality will be lost, which affects HTTP, FTP, and Telnet proxies.
- In-band data inspection to authorize secondary connections will be lost, which affects FTP, T.120, H.323, NetMeeting, and SOCKS proxies.
- The SIP proxy authorizes SIP calls, not point-to-point transport layer sessions, so the SIP proxy will drop all traffic when its inspection level is set to None.
- FTP traffic will be **allowed** when the HTTP Application Defense is set to None, even if the GET and PUT options are deselected in the FTP URL control tab of the HTTP Application Defense.

- c** In the **IPS Signature Group** field, select the IPS signature group to search when inspecting this rule's traffic.

- d** In the **Response Mapping** field, select the response mapping this rule will use when it finds a suspected IPS attack.

12 [Optional] In the Authentication area, select how authentication will be used for this rule:

- a** From the **Authenticator** drop-down list, select the authenticator that will be used to authenticate this rule.
- b** From the **Allow users in the following groups** drop-down list, select a group of users who will be allowed to authenticate.
 - If the rule is an *allow* rule, those users will be allowed to use the service.
 - If the rule is a *deny* or *drop* rule, the users will authenticate and then be denied access to the service. You can use authentication in a deny rule to deny a service to one group while allowing others access. For example, you can use a deny rule to deny corporate insiders access to stock trading web sites during blackout windows to prove due diligence.
 - You can select multiple user groups by typing the user group names in a comma-separated list (for example, *contractors, interns*) or by clicking  and selecting multiple options.
 - Almost all proxies can be authenticated using the Passport authenticator.
 - Services that support authentication even if not using Passport include:
 - Proxies: FTP, HTTP, HTTPS, SOCKS, and Telnet
 - Servers: login, Admin Console, Telnetd, sshd, and ssod
 - You are not allowed to create a rule using a service group if one of the services does not support that authenticator.
 - Not all filter services and related service groups support authentication.

13 Click **OK**.

14 Save your changes.

This rule is now a part of your security policy. For additional information on how to configure each option, see [Viewing and modifying rule elements](#).

Creating and modifying rule groups

This section provides information on creating and modifying rules groups. You can create an empty rule group, or you can select existing rules and add them directly to a new group. You can also nest groups within another group. To begin working with rule groups, select **Policy > Rules**.

To create an empty rule group

- 1 Determine where in the rule list your new rule group will go and then select the rule or rule group that will be directly above it. If nothing is selected, the rule group will be added to the bottom of the list.
- 2 Click **New Group**.
- 3 Enter a name and a description for the new group.
- 4 Click **OK**.

You can now add rules and other rule groups to this new rule group. Be sure to save your changes.

To place existing rules into a new rule group

- 1 Select the rules and rule groups to add to the new group.
- 2 Click **New Group**.
- 3 Enter a name and a description for the new group.
- 4 Check **Move selected items into new group**.
- 5 Click **OK**.
- 6 Verify that the rules are in the desired order.


You can now add rules and other rule groups to this new rule group. Be sure to save your changes.

To modify an existing rule group

- 1 Expand the rule group.
- 2 Select the rules to move into or out of the group. Hold down the **Shift** key to select multiple adjacent rules or the **Ctrl** key to select multiple non-adjacent rules.
- 3 Move the rules using any of these methods: dragging and dropping the rules, using cut and paste, or using the Up and Down arrows.
- 4 Verify that the rules are in the desired order.
- 5 If needed, modify the description.

You can now add rules and other rule groups to this new rule group. Be sure to save your changes.

Viewing and modifying rule elements


This section provides additional information on each part of a rule. It also describes the windows that appear when you click  next to a field.

- [Services](#)
- [Time periods](#)
- [Source burbs, endpoints, and NAT](#)
- [Destination burbs, endpoints, and redirection](#)
- [Application Defenses](#)
- [IPS response mapping and signature groups](#)
- [Authentication](#)

Services

Clicking ... button next to the Service field brings you to the Rule: Service window, where you can view the full list of existing services and service groups. You can also create new services and service groups or modify an existing service's properties while using this window. Services are methods for getting traffic through the firewall (proxies and filters) or into the firewall (servers).

Figure 175 The Rules: Service window



This window displays a list of all configured services. They are grouped by service (proxies, filters, servers) and then alphabetized.

Rules

Viewing and modifying rule elements

Use this window to do the following:

- Find a service or service group by entering a character string related to the object you are searching for in the **Find** field. The search function searches all columns, and filters as you type. The search is not case sensitive. For example, if you are searching a service based on the HTTP proxy, typing “http” reduces the list to only the services containing that character string.

Clear the **Find** field to show all options again.

- Add another service or service group by clicking **New** in the appropriate area. Once the new item is created, it is added to the list and can be used in this rule.
- Modify an existing service by selecting it and clicking **Modify**.

If the service is referenced by another area, the Usage window appears. Click **Yes** to modify the service.

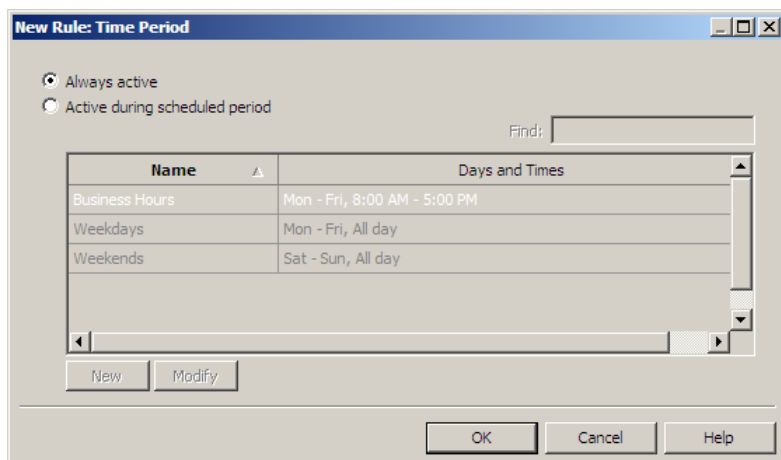
After you have determined which service or service group to use in this rule, select that item and then click **OK**. This service appears on the dependent rule, and the application defense options change accordingly.

To learn more about services and how to modify them, see [Chapter 7, Services](#).

Time periods

Clicking ... button next to the Time Period field brings you to the Time Period window, where you can view the full list of existing time periods. You can also create new time periods or modify an existing time period's properties while using this window. Time periods determine the specific times when a rule will be active.

Figure 176 The Rules: Time Period window



On a new rule, this window defaults to *always active*.

Use this window to do the following:

- Add another time period by clicking **New**. Once the new time period is created, it is added to the list and can be used in this rule.
- Modify an existing time period by selecting it and clicking **Modify**.

If the time period is referenced by another area, the Usage window appears. Click **Yes** to modify the time period.


After you have determined which time period to use in this rule, select **Active during scheduled period**. Then the select the time period to use and click **OK**.

To learn more about time periods and how to modify them, see [Creating time periods](#).

Source burbs, endpoints, and NAT

Clicking the ... button in the Source area brings you to the Source Options window, where you can view the full list of existing burbs and burb groups, network objects that can be used as endpoints, and NAT values. You can also create new burbs, burb groups, and network objects, or modify an existing item's properties. A rule's source is what can initiate a connection through, or into, the firewall.

Figure 177 The Rules: Source Options window



Note: On a new rule, the Source area defaults are an endpoint of <Any>, the NAT address of localhost (Host), and Preserve source port disabled. These defaults can be changed by modifying the new rule template.

Use the Source Options window to select the values for this rule's source. After you have selected the appropriate item or items, click **OK**. These values will appear on the dependent rule.

You can make the following selections and entries:

Burb

- **Any burb** – Select this to use any burb as the source endpoint for this rule.
- **Selected burbs** – Select the check box next to the burb or burbs to use in this rule.
 - Use **Select All/Deselect All** to select or clear the check boxes next to all available burbs.
 - Add another burb by clicking **New**. Once the burb is created, it is added to the list and can be used in this rule.
 - Modify an existing burb by selecting the burb name and clicking **Modify**.
- **Selected burb groups** – Select the check box next to the burb group or groups to use in this rule.
 - Use **Select All/Deselect All** to select or clear the check boxes next to all available burb groups.
 - Add another burb group by clicking **New**. Once the burb group is created, it is added to the list and can be used in this rule.
 - Modify an existing burb group by selecting the burb group name and clicking **Modify**.

Endpoint

- **Any endpoint** – Select this to allow any network object to initiate traffic. From the drop-down list, select the type of addresses to allow.
 - All Addresses
 - All IPv4 Addresses
 - All IPv6 Addresses (If IPv6 is enabled)
- **Selected endpoint** – Select a specific network object to initiate traffic.
 - Search for a particular network object by entering a character string in the **Find** field. You can search for both the name and properties. For example, if you are searching for a network object in the 192.168 subnet, typing **192.168** reduces the list to only network objects containing that character string. The search is not case sensitive.

Clear the **Find** field to show all options again.
 - Add another network object by clicking **New**. Once the network object is created, it is added to the list and can be used in this rule.
 - Modify an existing network object by selecting it and clicking **Modify**. If the network object is referenced by another area, the Usage window appears. Click **Yes** to modify the network object.

NAT

- **None** – Select this to turn NAT off for this rule.
- **Selected NAT** – Select the network object that will replace the original source address as the traffic leaves the firewall.
 - Search for a particular network object by entering a character string in the **Find** field. You can search for both the name and properties. For example, if you are searching for a network object in the 192.168 subnet, typing **192.168** reduces the list to only network objects containing that character string. The search is not case sensitive.

Clear the **Find** field to show all options again.
 - Add another network object by clicking **New**. Once the network object is created, it is added to the list and can be used in this rule.
 - Modify an existing network object by selecting it and clicking **Modify**. If the network object is referenced by another area, the Usage window appears. Click **Yes** to modify the network object.
 - [Conditional] The **Preserve source port** option is configurable only when the selected service's agent is a filter. This option is most commonly used in rules handling IKE traffic when the related VPN definition does not use NAT-T.

When **Preserve source port** is selected, the source address is translated to the associated NAT address, but the source port is not translated.

The translated address is obtained in one of these ways:

- From the address in the **NAT** field. This address must be an alias; it cannot be a native IP address.
- From a pool of IP addresses. There must be one or more alias addresses defined for the destination burb's interface and the NAT field must be set to **localhost**. The total number of connections is therefore dependent on the number of alias addresses defined for that interface.

Caution: To use this feature, you must have at least one alias configured for the destination burb's interface or traffic will not pass.

To learn more about the elements that make up a source and how to modify them, see the following sections:

- [Creating network objects](#)
- [Configuring burbs](#)

Destination burbs, endpoints, and redirection

Clicking ... button in the Destination area brings you to the Destination Options window, where you can view the full list of existing burbs and burb groups, network objects that can be used as endpoints, and redirect values. You can also create new burbs, burb groups, and network objects, or modify an existing item's properties. A rule's destination is what can receive or respond to traffic initiated by the rule's source.

Figure 178 Rules: Destination Options window

New Rule: Destination Options

Burb

☐ Any burb

☒ Selected burbs

external
internal

New Modify Select All Deselect All

Selected burb groups

New Modify Select All Deselect All

Endpoint

☒ Any endpoint All Addresses

☐ Selected endpoint Find:

Name	Properties
localhost	localhost
external primary DNS re...	10.65.240.246
Firewall	127.0.0.1
DNS resolvers	external primary DNS resolver
<SmartFilter server >	127.0.0.1
<CommandCenter serve...	

New Modify

Redirect

☒ None

☐ Selected redirect Find:

Name	Properties
localhost	localhost
external primary DNS re...	10.65.240.246
Firewall	127.0.0.1
<SmartFilter server >	127.0.0.1

New Modify

Redirect Port:

OK Cancel Help

Note: On a new rule, the Destination area defaults are an endpoint of **<Any>**, a Redirect address of **<None>**, and Redirect port set to blank or 0 (do not translate). These defaults can be changed by modifying the new rule template.

Use the Source Options window to select the values for this rule's source. After you have selected the appropriate item or items, click **OK**. These values will appear on the dependent rule.

You can make the following selections and entries:

Burb

- **Any burb** – Select this to use any burb as the destination endpoint for this rule.
- **Selected burbs** – Select the check box next to the burb or burbs to use in this rule.
 - Use **Select All/Deselect All** to select or clear the check boxes next to all available burbs.
 - Add another burb by clicking **New**. Once the burb is created, it is added to the list and can be used in this rule.
 - Modify an existing burb by selecting the burb name and clicking **Modify**.
- **Selected burb groups** – Select the check box next to the burb group or groups to use in this rule.
 - Use **Select All/Deselect All** to select or clear the check boxes next to all available burb groups.
 - Add another burb group by clicking **New**. Once the burb group is created, it is added to the list and can be used in this rule.
 - Modify an existing burb group by selecting the burb group name and clicking **Modify**.

Endpoint

- **Any endpoint** – Select this to allow traffic to be sent to any network object. From the drop-down list, select the type of addresses to allow.
 - All Addresses
 - All IPv4 Addresses
 - All IPv6 Addresses (If IPv6 is enabled)
- **Selected endpoint** – Select a specific network object for traffic to be sent to.
 - Search for a particular network object by entering a character string in the **Find** field. You can search for both the name and properties. For example, if you are searching for a network object in the 192.168 subnet, typing **192.168** reduces the list to only network objects containing that character string. The search is not case sensitive.

Clear the **Find** field to show all options again.
 - Add another network object by clicking **New**. Once the network object is created, it is added to the list and can be used in this rule.
 - Modify an existing network object by selecting it and clicking **Modify**. If the network object is referenced by another area, the Usage window appears. Click **Yes** to modify the network object.

Redirect

- **None** – Select this to turn redirection off for this rule.
- **Selected redirect** – Select the network object that traffic will be redirected to.
 - Search for a particular network object by entering a character string in the **Find** field. You can search for both the name and properties. For example, if you are searching for a network object in the 192.168 subnet, typing **192.168** reduces the list to only network objects containing that character string. The search is not case sensitive.

Clear the **Find** field to show all options again.
 - Add another network object by clicking **New**. Once the network object is created, it is added to the list and can be used in this rule.
 - Modify an existing network object by selecting it and clicking **Modify**. If the network object is referenced by another area, the Usage window appears. Click **Yes** to modify the network object.
- **Redirect Port** – This is the port the connection is redirected to. Note the following:
 - The default is blank. This means the port remains unchanged. Entering a 0 in this field also leaves the port unchanged.
 - Valid values are 1–65535 (inclusive).

- This field is not available for all services.

To learn more about the elements that make up a destination and how to modify them, see the following sections:

- [Creating network objects](#)
- [Configuring burbs](#)

Application Defenses


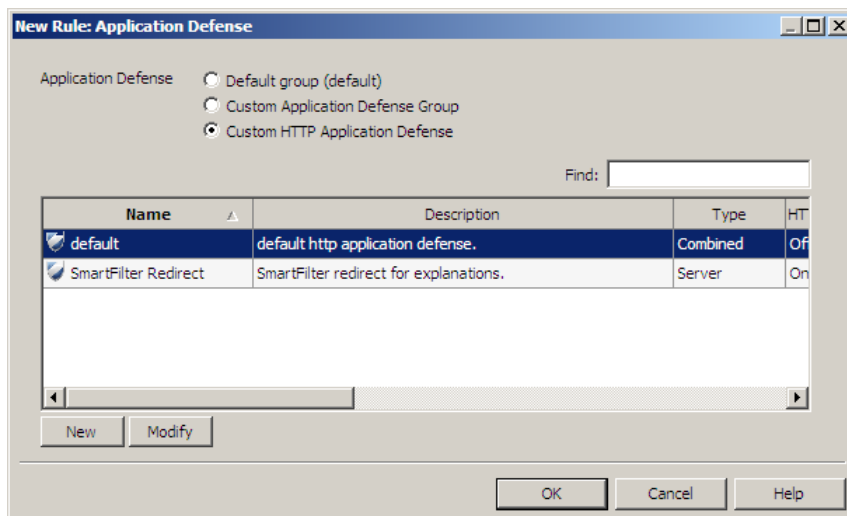
Clicking  next to the Application Defense field brings you to the Rule: Application Defense window, where you can view the full list of application defenses and application defense groups that are appropriate for the selected service. You can also create new application defenses, or modify an existing application defense's properties, while using this window. Application defenses contain the settings for inspecting advanced application-level content, such as headers, commands, and filters. They also enable additional features such as virus scanning and web filtering.

Figure 179 The Rules: Application Defense window



On a new rule, this window defaults to the default group associated with the selected service or service group.

Note: Some servers do not use configurable application defenses.

Use this window to do the following:

- Add another application defense by clicking **New**. Once the new application defense is created, it is added to the list and can be used in this rule.
- Modify an existing application defense by selecting it and clicking **Modify**.

If the application defense is referenced by another area, the Usage window appears. Click **Yes** to modify the application defense.

- To select a different application defense group to use in this rule, first select **Custom Application Defense Group**. Then select the application defense group and click **OK**. This item appears on the dependent rule.
- To select a different application defense to use in this rule, first select **Custom Application Defense**. Then select the application defense and click **OK**. This item appears on the dependent rule.

To learn more about application defenses and how to modify them, see [Chapter 8, Application Defenses](#).

IPS response mapping and signature groups


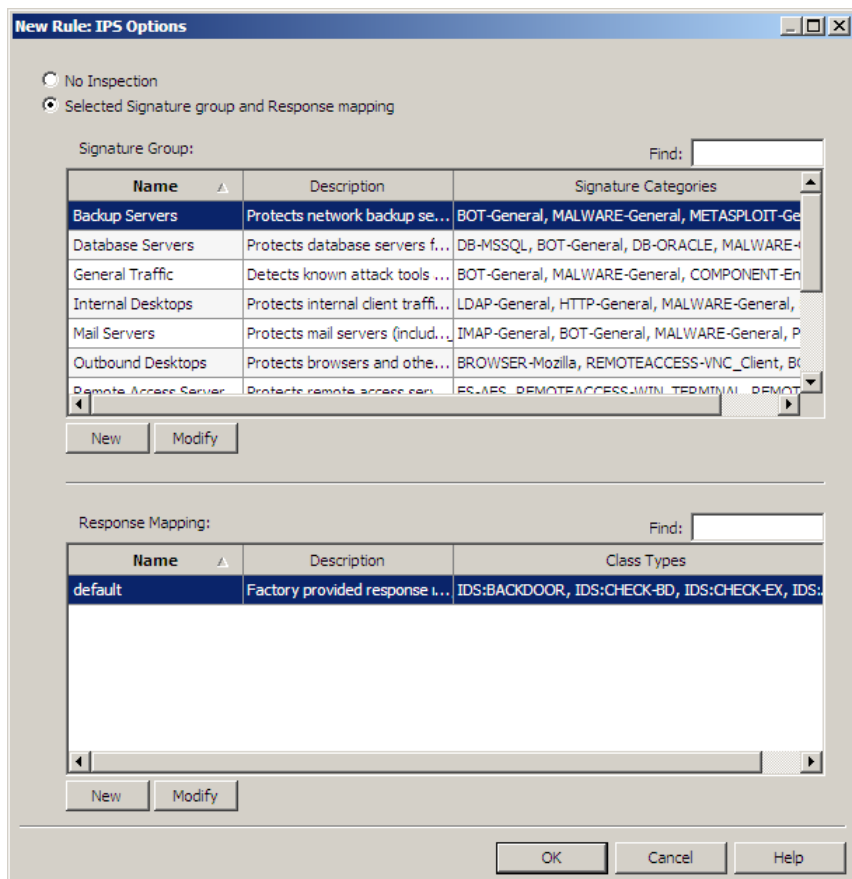
IPS inspection consists of two rule elements: Signature Groups and Response Mappings. Clicking  next to the Response Mapping field opens the Rule: IPS Options window.

Figure 180 Rules: IPS Options window



This window contains two elements:

- Signature groups – Each group contains one or more signature categories which identify the type of intrusion for which this rule is searching.
- IPS mappings – Mappings contain the settings for how the firewall responds when it identifies a known network-based attack.

You can perform the following actions on this window:

- View the full list of signature groups and response mappings
- Create new groups and mappings
- Modify the contents of an existing signature group or mapping (signature groups each contain one or more signature categories)

On a new rule, this window defaults to No Inspection. Some services do not support the use of IPS inspection.

Rules

Viewing and modifying rule elements

Use this window to do the following:

- Add another signature group or response mapping by clicking **New**. The new mapping is then added to the list and can be used in this rule.
- Modify an existing item by selecting it and clicking **Modify**. If the item is referenced by another area, the Usage window appears. Click **Yes** to modify the item.

After you have determined which signature group and response mapping to use in this rule, select the items and click **OK**. These items appear on the window.

To learn more about intrusion protection services, see [Chapter 6, Content Inspection](#).

Authentication


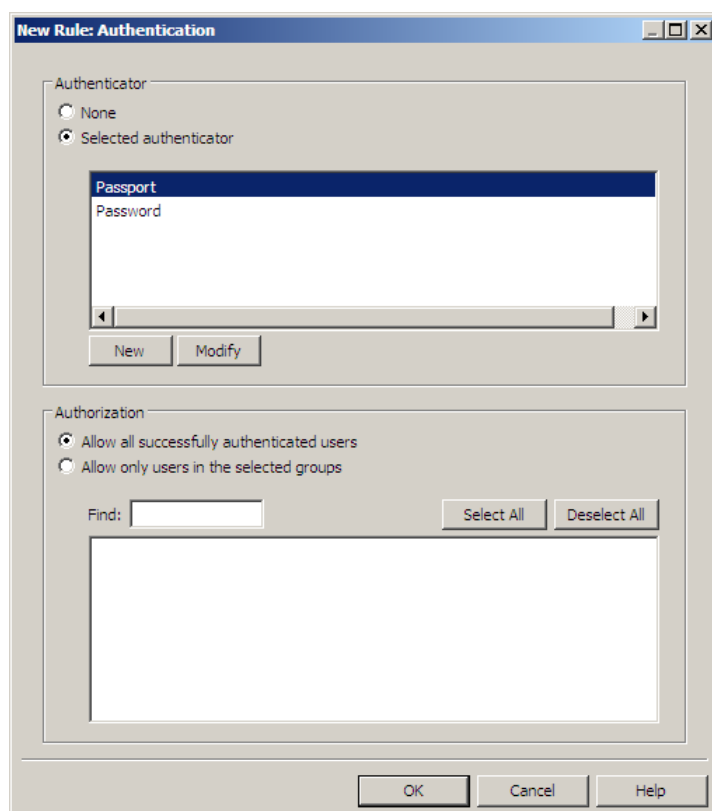
Clicking  in the Authentication area brings you to the Authentication window, where you can view the list of possible authenticators for the selected service, as well as corresponding authorization properties. You can also create new authenticators or modify an existing authenticator's properties. Authenticators are applications that validate a person's identity before he or she is allowed to log into a network service. Authorization determines which users will use that authentication method.

Figure 181 The Rules: Authentication window



On a new rule, this window defaults to no authentication (None).

Use this window to do the following:

- Find a user group by entering a character string related to the group you are searching for in the **Find** field. The search is not case sensitive. For example, if you are searching for an Engineering department user group, typing “Eng” reduces the list to only network objects containing that character string.

Clear the Find field to show all options again.

- Add another authenticator by clicking **New** in the appropriate area. Once the new item is created, it is added to the list and can be used in this rule.
- Modify an existing authenticator by selecting it and clicking **Modify**.

If the authenticator is referenced by another area, the Usage window appears. Click **Yes** to modify it.

- Check the item or items to use in this rule. Use **Select all/Deselect all** to select or clear the check boxes next to all the user groups.

After deciding which authenticator to use in this rule, do the following:

1 Select **Selected authenticator**.

2 Select an authenticator.

3 [Conditional] If authentication will be limited based on user group or external group, select **Allow only users in the select groups** and then select one or more groups.

4 Click **OK**.

To learn more about authentication and authorization, see [Chapter 5, Authentication](#).

SECTION 3

Monitoring

Chapter 10, The Dashboard

Chapter 11, Auditing

Chapter 12, Service Status

Chapter 13, IPS Attack and System Event Responses

Chapter 14, Network Defenses

Chapter 15, The SNMP Agent

10 The Dashboard

Contents

[Monitoring Sidewinder status using the dashboard](#)

[Viewing device information](#)

[Viewing network traffic information](#)

[Viewing IPS attack and system event summaries](#)

Monitoring Sidewinder status using the dashboard

The Admin Console allows you to monitor status information on your Forcepoint Sidewinder using its dashboard. The `monitord` server records data about the system and traffic status. Auditbots detect packets and traffic patterns that may be of interest to administrators. The dashboard gathers this data from those and other firewall components and provides a centralized view of important system and audit data. This window displays summary data and specific audit events.

The dashboard allows you to monitor the following areas:

- Device information (version, uptime, configuration state, etc.)
- Network traffic (active VPN and proxy sessions, interface status, etc.)
- Recently detected attack activity
- System events (hardware and software failures, log overflows, etc.)

You can set this information to refresh automatically or on demand.

While this window is a useful tool to observe your firewall, you may also want to take advantage of other audit and monitoring tools:

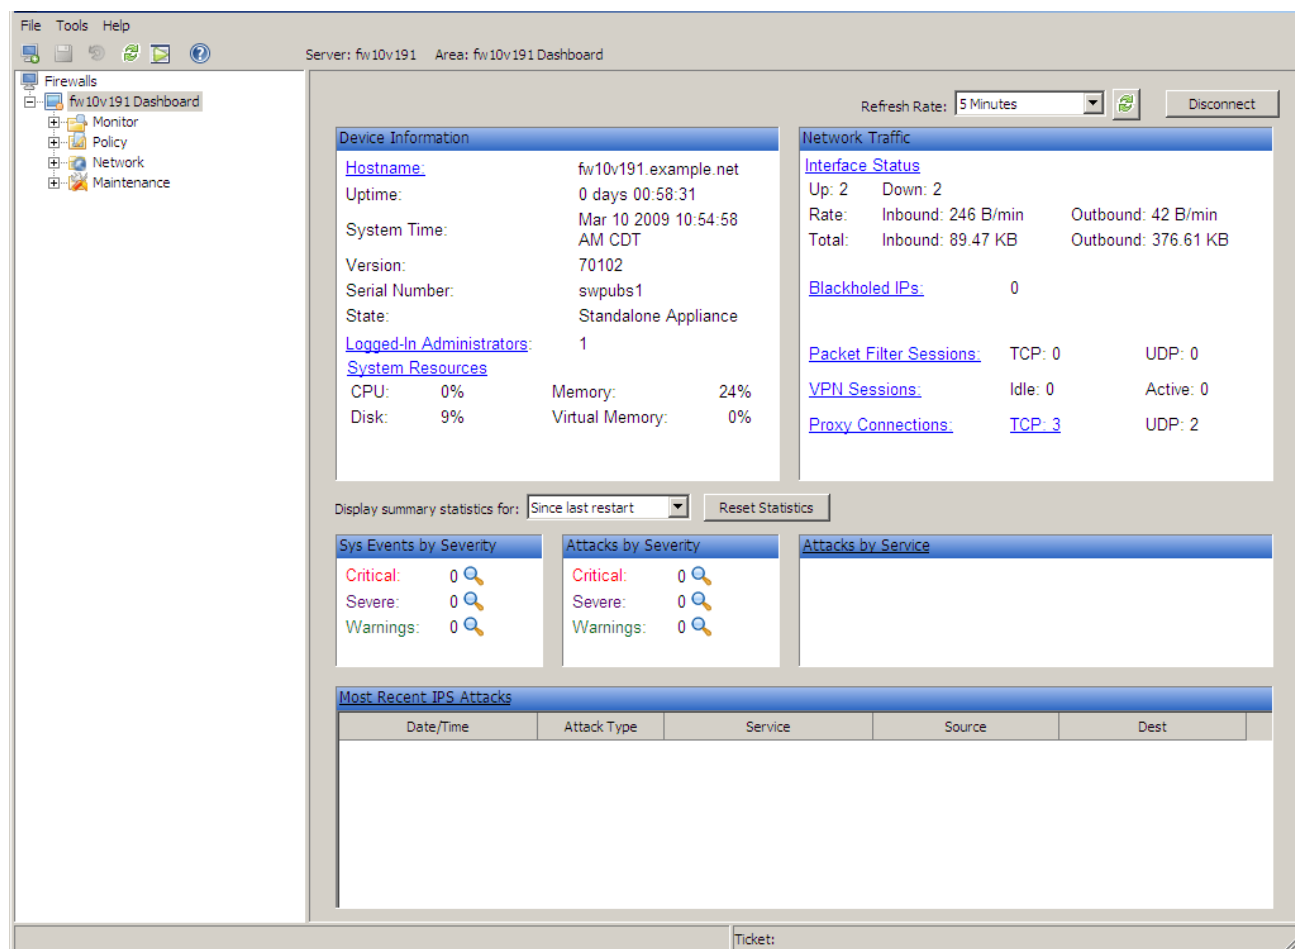
- For additional audit information, see [Chapter 11, Auditing](#).
- For information on commands that monitor the firewall, see [Troubleshooting system status](#).

The Dashboard

Monitoring Sidewinder status using the dashboard

When you log into the Admin Console, the dashboard displays. To view the dashboard at any other time, click the root node of the tree labeled **firewall_name Dashboard** (where *firewall_name* is the name of your firewall in the tree). A window similar to the following appears:

Figure 182 The dashboard



The dashboard allows you to monitor various firewall areas. It displays statistics recorded since the last reboot. From the dashboard, you can:

- **Monitor the firewall's status** – Monitor general system information, what traffic is passing through the firewall, and system and attack events. For more information on each area, see the following sections:
 - [Viewing device information](#)
 - [Viewing network traffic information](#)
 - [Viewing IPS attack and system event summaries](#)
- **View additional information** – Learn more about any given area by clicking the appropriate link or magnifying glass icon.
- **Change the refresh rate** – Indicate how often the dashboard will refresh by using the **Refresh Rate** field. Valid values range from 30 seconds to 30 minutes. There is also a Manual Refresh option. The default is 5 minutes. When you modify the refresh rate, the change will not take effect until the next scheduled refresh time. To make the change take effect immediately, change the refresh value and click the **Refresh** icon.
- **Manage blackholed IP addresses** – View a list of the IP addresses the firewall is currently blackholing. You can also delete addresses that do not need to be blackholed and manually add new addresses to the list. To manage blackholed IP addresses, click **Blackholed IPs**.
- **Disconnect** – Disconnect the current Admin Console session by clicking the **Disconnect** button. If you hover the mouse pointer over the Disconnect button, a tool tip appears that includes the connected firewall's IP address.

Viewing device information

The dashboard's Device Information area, shown in [Figure 183](#), displays basic system information. The device information that this area monitors includes: the firewall host name, the amount of time since the last reboot, the date and time, the current version, the serial number, data about logged-in administrators, and basic system resource data for the whole system, with the option to view process-specific data as well.

Figure 183 Dashboard: Device Information area

Device Information			
<u>Hostname:</u>		fw10v191.example.net	
Uptime:		0 days 00:52:08	
System Time:		Mar 10 2009 10:48:35 AM CDT	
Version:		70102	
Serial Number:		swpubs1	
State:		Standalone Appliance	
<u>Logged-In Administrators:</u>		1	
<u>System Resources</u>			
CPU:	1%	Memory:	24%
Disk:	9%	Virtual Memory:	0%

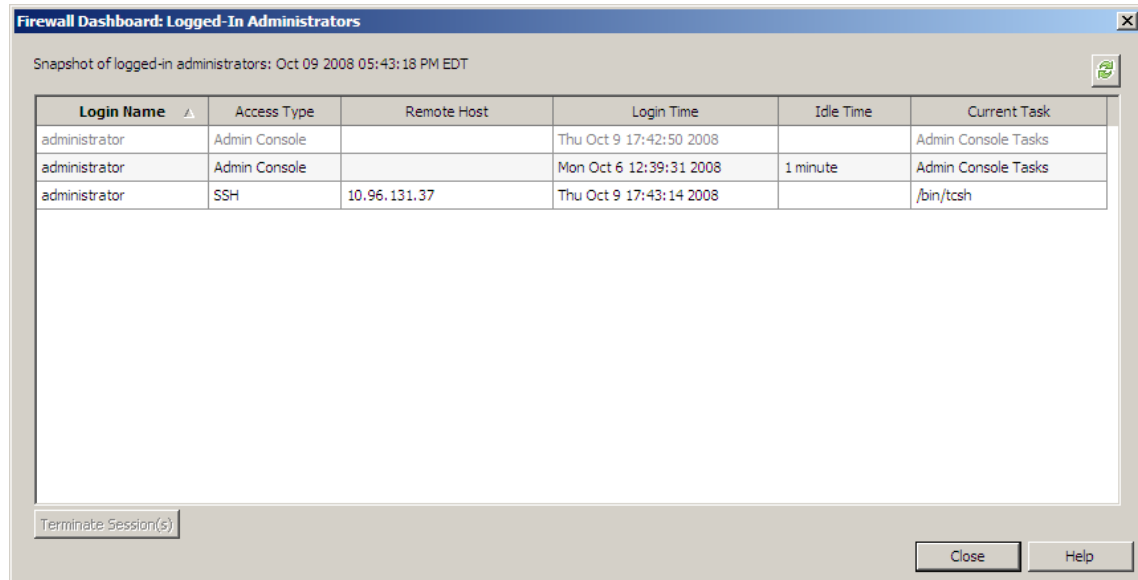
In this area, you can do the following:

- Change the firewall's host name by clicking **Hostname**.
 - Changing the host name affects your DNS configuration, sendmail configuration, and all entries in your */etc/resolv.conf** files. You must manually change any necessary entries to ensure proper functioning.
 - You will be prompted to restart the firewall. The firewall must be restarted for the change to take effect.
- View information about administrators who are logged into this firewall by clicking **Logged-In Administrators**.
- View process use and disk use information by clicking **System Resources**. The relevant information appears on separate tabs in the pop-up window.
- Receive feedback that a system resource may be experiencing trouble. If the value turns red, the memory or disk may be getting too full and requires attention. Click **System Resources** to view more information.

About the Logged-In Administrators window

Use this window to view information about administrators who are currently logged into this firewall.

Figure 184 The Logged-In Administrator window



The Logged-In Administrators window displays the following information:

- **Login Name** – Logged-in administrators' user names
- **Access Type** – Management program/protocol (Admin Console, SSH, Telnet, System Console)
- **Remote Host** – If not using the Admin Console, the IP address or host name of the host that initiated the management session.
- **Login Time** – Time stamp of the most recent successful login
- **Idle Time** – Time since the administrator's last action
- **Current Task** – What each administrator is doing when the window is refreshed (if known)

On this tab, you can do the following:

- Select one or more administrator's rows and then click **Terminate Session(s)** to close an open session.
- Click **Refresh** to view current information. This window does not automatically refresh.
- Click **Close** to close this window and return to the dashboard.

About the Process Use tab

This tab displays the status of each process that is currently running on this firewall.

Figure 185 System Resources: Process Use tab

Firewall Dashboard: System Information

Snapshot of system resources: Oct 09 2008 05:44:39 PM EDT

Process Use | Disk Use | CPU Use

Process	CPU %	Process Size	Resident Memory
nmbresd	0.00	1.53 MB	880.00 KB
auditbotd	0.00	2.42 MB	1.77 MB
auditd	0.00	2.13 MB	1.19 MB
cron	0.00	1.59 MB	1.17 MB
httpd	0.00	16.05 MB	3.83 MB
pupd	0.00	2.80 MB	1.61 MB
AdminConsole	0.00	22.75 MB	21.24 MB
kms	0.00	2.58 MB	1.90 MB
scobrap	0.00	3.57 MB	2.95 MB
kms	0.00	2.58 MB	1.89 MB
nss	0.00	3.92 MB	3.02 MB
nss	0.00	3.91 MB	3.02 MB
dnsp	0.00	2.89 MB	1.88 MB

Close Help

It provides the following details for each process:

- **Process** – This column displays the name of each running process.
- **CPU** – This column displays the percentage of CPU currently being used.
- **Process Size** – This column displays the amount of memory a process is using.
- **Resident Memory** – This column displays the amount of physical memory a process is using.

On this tab, you can do the following:


- Click **Refresh** to update this tab's data.
- Click **Close** to close this window and return to the Dashboard.

About the Disk Use tab

This tab displays how much of the appliance's hard disk space is currently being used.

Note: The /dev file system is a virtual file system and does not actually occupy space on your hard drive. The Percent Used column should display 100% used, and the Used value should be 1.00 KB.

Figure 186 System information: Disk Use tab



Firewall Dashboard: System Information

Snapshot of system resources: Oct 09 2008 05:44:39 PM EDT

Process Use | **Disk Use** | CPU Use

Mounted On	Percent Used ▾	Used	Available	Description
/dev	100	1.00 KB	0 B	Devices
/usr	34	309.10 MB	601.00 MB	System tools and utilities such as man pages, Perl, and .
/vcdrom	24	667.01 MB	2.02 GB	Package data
/tmp	21	48.89 MB	179.03 MB	Temporary files
/	17	24.82 MB	117.54 MB	Standard UNIX libraries and system configuration settin_
/secureos	3	78.68 MB	2.46 GB	Configuration, policy, AV files
/var/spool	1	80.81 MB	5.23 GB	Queued mail traffic, Admin Console package data
/var	1	20.30 MB	1.90 GB	Configuration backups, process run directories, packag.
swap	0	0 B	5.00 GB	Swap space
/var/log	0	6.05 MB	6.86 GB	Auditing and logging information
/home	0	112.00 KB	1.64 GB	Home directories for administrator accounts

Close Help

It provides the following details for each disk partition:

- **Mounted On** – This column displays the name of each disk partition.
- **Percent Used** – The column displays the percent of that partition being used.
- **Used** – This column displays the amount of a given partition being used.
- **Available** – This column displays the amount of disk space available for use in the given partition.
- **Description** – This column displays a description of the disk partition.

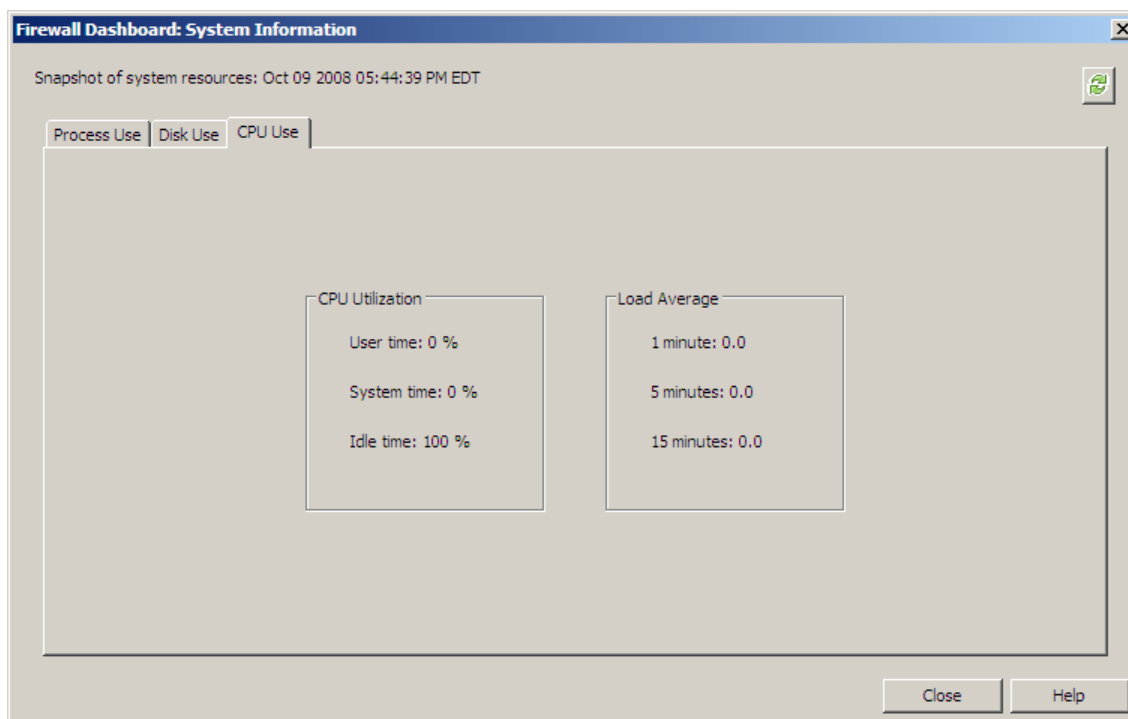
On this tab, you can do the following:

- Click **Refresh** to update this tab's data.
- Click **Close** to close this window.

About the CPU Use tab

This tab displays appliance utilization and load average information. You can click **Refresh** to instantly update the numbers.

Figure 187 System information: CPU Use tab



CPU utilization

The CPU Utilization area displays the percentage of a one-second interval that the appliance spent in the following areas:

- **User time** – Handling non-kernel processes
- **System time** – Handling kernel processes
- **Idle time** – Doing nothing

Load average

Load average, a UNIX term for the average system load over a specified time period, measures how hard the appliance is working. Load average helps you understand how long processes have been waiting in the queue over the previous 1-minute, 5-minute, and 15-minute periods.

Viewing network traffic information

The dashboard's Network Traffic area, shown in [Figure 188](#), displays information on network traffic passing through the firewall. View information such as number of interfaces up and receiving traffic, number of active filter rules, number of active VPN sessions, and number of active proxy and server service connections.

Figure 188 Dashboard: Network Traffic area

Network Traffic		
Interface Status		
Up: 2	Down: 2	
Rate:	Inbound: 300 B/min	Outbound: 0 B/min
Total:	Inbound: 87.67 KB	Outbound: 364.25 KB
Blackholed IPs: 0		
Packet Filter Sessions:	TCP: 0	UDP: 0
VPN Sessions:	Idle: 0	Active: 0
Proxy Connections:	TCP: 3	UDP: 2

Use this area of the dashboard to monitor the following:

- **Interface Status** – Displays the status of all physical and VLAN interfaces in the firewall. The displayed rate is data on the transfer speed in the last minute, reported in bytes per minute. The displayed total is the number of inbound/outbound bytes processed since the last reboot.

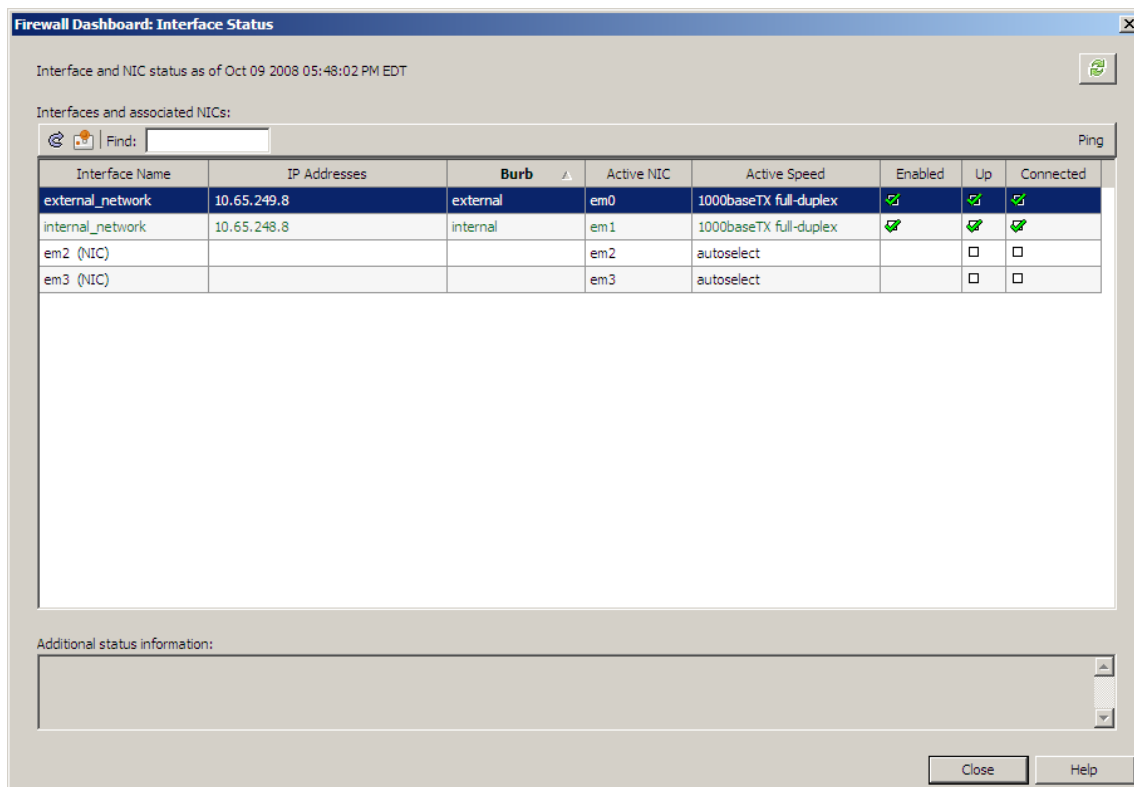
Click **Interface Status** to view additional information about each interface. See [About the Interface Status window](#) on the following page for more information.
- **Blackholed IPs** – Click **Blackholed IPs** to view and manage the currently blackholed IP addresses. See [About the Blackholed IPs window](#) for more information.
- **Packet Filter Sessions** – Displays the number of packet filter sessions that are currently open. The link turns red when 90 percent of the sessions are being used.

A rule's filter service must have Stateful Packet Inspection enabled to create a session.
- **VPN Sessions** – Click **VPN Sessions** to view additional information about configured VPNs. See [About the Active VPNs window](#) for more information.
- **Proxy Connections** – Displays the current number of TCP and UDP sessions. Click **Proxy Connections** to view a list of each proxy and server service that is currently passing traffic and the number of instances each service. Click **TCP:** to display how many connections are in each state. See [About the Proxy Connections window](#) and [About the TCP State Information window](#) for more information.

About the Interface Status window

Use this window to review the status of each interface, NIC group, and unused NIC on the firewall.

Figure 189 Network Traffic: Interface Status window



The table shows the following information:

- The timestamp at the top of the window shows the time that the information was gathered. Click **Refresh** to update the information.
- The name, IP address, and burb for the interface are listed, and whether the interface is enabled and connected.
- The Active NIC column shows the NIC that is passing traffic for that interface. If an interface is using a NIC group, the active NIC in the NIC group is listed rather than the NIC group itself.
- The Active Speed column shows the currently running speed of the active NIC. This is useful when the selected media type for the NIC is **autoselect**.
- The Up column shows interface availability. If this column is checked, the interface is ready for an active network connection. If the column is cleared, the interface will not accept an active network connection.
- The Connected column shows network cable connection. If this column is checked, the network cable is plugged into the active NIC. If the column is cleared, the network cable is not plugged into the NIC.
- The Additional status information box gives a brief explanation of the state indicated in the table.

The following alerts indicate interface problems:

- A warning icon appears next to the interface name if traffic is passing but failure is likely. This can happen during NIC failover or when a standby NIC in a NIC group is down.
- Red text and a slash icon next to the interface name indicate that a problem is preventing traffic from passing through the interface or NIC. This might be caused by a NIC being down or disconnected.

You can perform these actions:


- To restart a NIC that is down, select the NIC in the list and click **Restart NIC**.
- To view how an interface, NIC, or NIC group is being referenced, select it in the list and click **Usage**.
- To search for a specific element(s) in the list, type your search criteria in the **Find** field, and interfaces with matching elements will appear in the list. Clear this field to see the full list again.
- To test interface connectivity, click **Ping** and use the pop-up window to send a ping to a specified address. See [About the Ping Test window](#) for more information.

Click **Close** to close this window.

About the Blackholed IPs window

Use this window to view and manage the currently blackholed IP addresses.

Figure 190 Network Traffic: Blackholed IPs window



Each entry in the table displays the IP address, burb, and the date and time at which the IP address will no longer be blackholed. You can perform the following actions in this window:


- **Add an IP address to blackhole** – To add an IP address to this list, click **Add**. In the Add Blackhole IP pop-up window, enter the IP address you want to blackhole and how long, in seconds, before the firewall will accept and respond to traffic from that IP address. This address is then automatically blackholed on all configured burbs.
- **Delete one or more entries** – To remove one or more entries from the list, select the row you want to delete and click **Delete**. To select multiple rows, press and hold the **Ctrl** key as you select the items.
- **Delete all IP entries** – To remove all of the entries that are listed in the table, click **Delete All**.
- **Update the window** – To retrieve an updated list of blackholed IP addresses, click the **Refresh Now** icon. The date and time when displayed data was captured is listed in the upper portion of the window.

Click **Close** to exit the window. Changes are saved automatically.

About the Active VPNs window

Use this window to monitor the status of all configured VPNs.

Figure 191 Network Traffic: Active VPNs window



The statuses include:


- **Idle** – No active session.
- **Active** – One or more VPNs have active sessions established for this VPN.

Click **Refresh** to update the information. Click **Close** to return to the main window.

About the Proxy Connections window

Use this window to monitor the type and number of active proxy sessions going through the firewall.

Figure 192 Network Traffic: Proxy Connections window



Name	Count
scobrap	4
sshd	2
dnsp	2
AdminConsole	2

Information provided includes:

- **Name** – Name of the proxy passing traffic
- **Count** – Number of current instances


On this window, you can:

- Click **Refresh** to update the information.
- Click **Close** to return to the main window.

About the TCP State Information window

Use this window to monitor the various states of the TCP proxy connections going through the firewall.

Figure 193 Network Traffic: TCP State Information window



TCP State	Count	Description
Established	7	An open connection ready to transmit and receive data
Time Wait	0	A connection waiting to ensure the remote TCP session received the...
FIN Wait 1	0	A connection waiting for a termination request or an acknowledgem...
FIN Wait 2	0	A connection waiting for a termination request
Close Wait	0	Waiting for a termination request from the local user

Information provided includes:

- **TCP State** – Lists the different possible states of a TCP connection.
- **Count** – Number of TCP sessions in that state
- **Description** – Describes that row's TCP state.

On this window, you can:

- Click **Refresh** to update the information.
- Click **Close** to return to the main window.

Viewing IPS attack and system event summaries

The statistics summary area of the dashboard displays a summary of the audit events that the firewall detects. By default, the firewall audits packet and traffic patterns it assumes to be an attack. It also audits system events administrators tend to consider important. Each predefined audit event is related to a severity. The dashboard summarizes the audit events for a given time frame, providing administrators a quick overview of audit activity. View additional details by clicking the magnifying glasses, links, and audit rows.

Understanding audit event severities

IPS attack audit events are based on anomaly detection. They are not necessarily detecting a specific attack attempt, but are detecting unexpected or suspicious deviations from allowed packets and patterns. The severities represent the assumed risk to the firewall and its protected systems if the attack had not been blocked. For example, an attack event generated by a commonly occurring packet that is used to gather information is considered a *warning*. An attack event made up of packets that appear to be crafted and, if not blocked, could crash a vulnerable system are considered *severe* or *critical*. Administrators should immediately investigate all critical attacks.

System audit events are generated by expected and unexpected system behavior. The severities are generally based on the type of action, if any, an administrator should take in response to the event. Whereas a critical event generally requires immediate investigation, a warning generally requires no action from the administrator.

[Table 32](#) defines each severity in more detail.

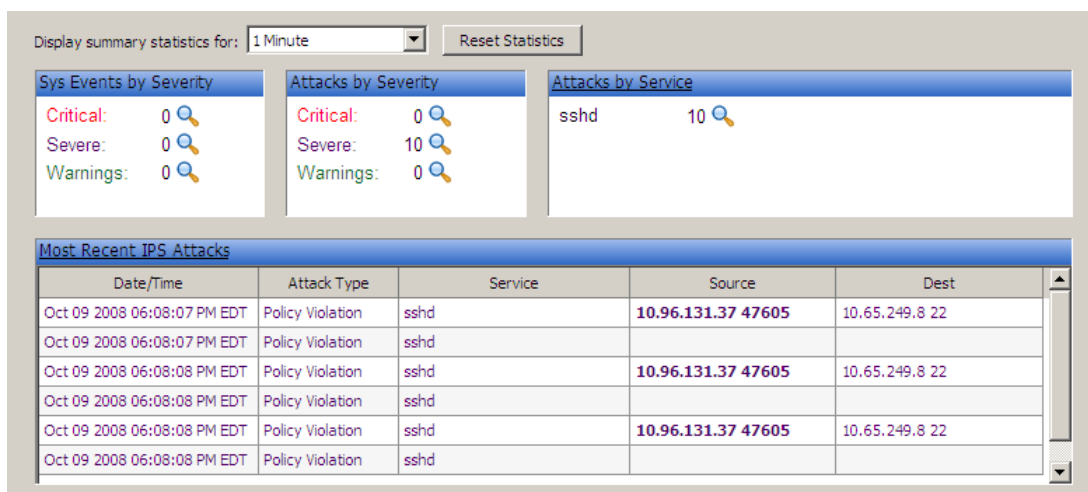
Table 32 Definitions of IPS attack and system event severities

Severity	Definition
Critical	<ul style="list-style-type: none">Indicates activity that is definitely an attack and that could have significantly affected a protected system had it not been prevented.Indicates that a system component or subsystem stopped working, that the system is going down (expectedly or unexpectedly), or that the system is not expected to work again without intervention. <p>At the command line, these audit events are classified as <i>emergency</i>, <i>alert</i>, <i>critical</i>, and <i>fatal</i> priorities.</p>
Severe	<ul style="list-style-type: none">Indicates activity that represents a likely significant attack or policy violation.Indicates something is occurring in the system that an administrator should know. <p>At the command line, these audit events are classified as a <i>major</i> priority.</p>
Warning	<ul style="list-style-type: none">Indicates activity that may be an attack or information gathering, or that represents a minor attempted violation of the site security policy (for example, attempting to use a restricted FTP command).Indicates something is occurring in the system that an administrator might want to know or might consider trivial. <p>At the command line, these audit events are classified as <i>minor</i> or <i>trivial</i> priorities.</p>


Viewing the summary statistics

The summary statistics area is located in the lower portion of the dashboard, as shown in [Figure 194](#).

Figure 194 Summary statistics area



In this area, you can:

- Change the displayed statistics based on a time period by selecting different options in the **Display summary statistics for** drop-down list. The range of options vary depending on the firewall's uptime.
- Reset the displayed statistics to 0 by clicking **Reset Statistics**.
- View audit data for any system event or attack category by clicking the magnifying glass .
- View a snapshot of all attacks listed by service by clicking **Attacks by Service**. See [About the Attacks by Service window](#) for more information.
- View and save attack audit data by clicking **Most Recent IPS Attacks**.
- View an individual audit record by double-clicking that audit event's row. See [About the Audit Record window](#) for more information.

Use this area of the dashboard to monitor the following:

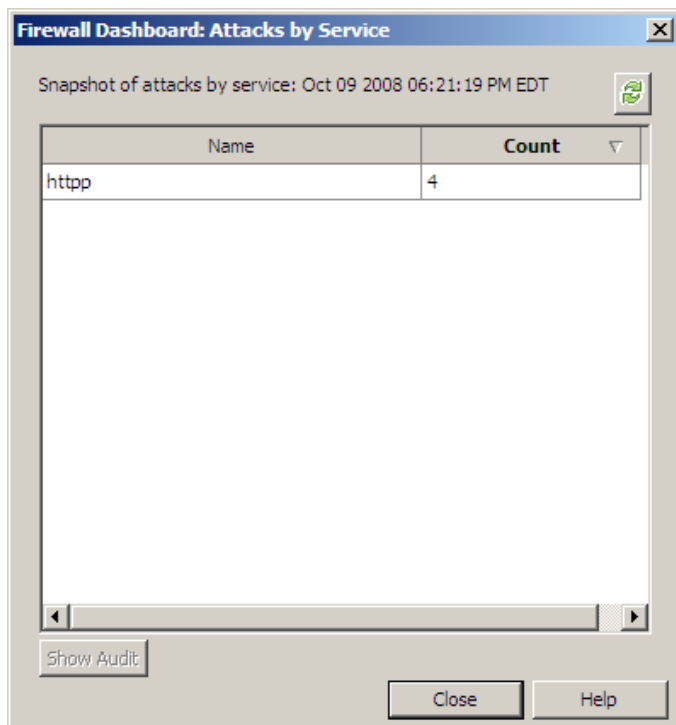
- **System events by severity** – Lists system audit events according to severity.
- **Attacks by severity** – Lists audit attack events according to severity.
- **Attacks by service** – Lists audit attack events according to service.
- **Most recent IPS attacks** – Displays the audit events for recent attacks.

Note: Use the Admin Console's IPS Attack Responses and System Event Responses to determine how the firewall reacts to different audit events. For more information, see the "IPS Attack and System Event Responses" chapter.

About the Attacks by Service window

Use this window to view audit of suspect traffic.

Figure 195 Attacks by Service window



Information provided includes:

- **Name** – Name of the service being attacked
- **Count** – Number of attack instances

On this window, you can:

- Click **Refresh** to update the information.
- Select a service and click **Show Audit** to see the audit output. You can also view the audit by clicking the magnifying glass on the main window.
- Click **Close** to return to the main window.

About the Audit Record window

When you double-click an audit event in the table, the detailed audit information for that attack appears in a pop-up window.

Figure 196 Audit Record window



The displayed fields vary, depending on the audit type. In general, the data in an audit message is a tag name followed by a colon and the tag's value. The following table provides examples and descriptions of fields that may appear in an audit record. Most administrators begin troubleshooting by noting the reason an event was audited and then examining the source and destination information.

More information on audit fields is available using `acat -c |more` at a command line interface and in the Sidewinder Export Format application note at <https://support.forcepoint.com>.

Table 33 Audit filter fields

Field	Description
facility	Specify an event facility code (such as <code>f_login</code> , <code>f_proxy</code> , etc.).
type	Specify an event type code (for example, type <code>t_nettraffic</code>).
category	Specify an event category code (for example, <code>c_policy_violation</code>).
eventid	Specify an event identifier code (for example, <code>r_licexceeded</code>).
hostname	Specify a host name.
username	Specify a user name.
src_ip	Specify the source IP address. Separate optional mask bits with a slash (/).
dst_ip	Specify the destination IP address. Separate optional mask bits with a slash (/).
src_port	Specify the TCP or UDP source port.
dst_port	Specify the TCP or UDP destination port.
src_burb	Specify the source burb name or index number.
dst_burb	Specify the destination burb name or index number.
service	Specify the service name. (To filter on an agent, use the facility field.)
vpn_l_gw	Specify a VPN local gateway using the standard dotted decimal IP version 4 notation with optional mask bits separated by a slash (/).
vpn_r_gw	Specify a VPN remote gateway using the dotted decimal IP version 4 notation with optional mask bits separated by a slash (/).

11 Auditing

Contents

[Understanding the Sidewinder audit process](#)

[Viewing audit information](#)

[Managing log files](#)

[Exporting audit data to Firewall Reporter and syslog servers](#)

Understanding the Sidewinder audit process

Monitoring, auditing, reporting, and attack and system event responses are closely related pieces of the audit process. They function together to provide information to you about the activity on your Forcepoint Sidewinder. You can monitor the status of various processes in real time, view stored audit information, generate detailed reports, and have the firewall respond to audit events by alerting administrators and ignoring hosts sending malicious packets.

Auditing is one of the firewall's most important features. It provides information on what is happening with your system and fulfills compliance regulations. The firewall generates audit information each time it or any of its services are stopped or started. It generates audit data for what configuration changes are made and who made them. Other relevant audit information includes:

- Identification and authentication attempts (successful and failed)
- Network communication (including the presumed addresses of the source and destination subject)
- Administrative connections (using srole)
- Modifications to your security policy or system configuration (including all administrator activity, such as changing the system time)

Because audit records are important, storing them is a high priority. The audit facilities monitor the state of log files to minimize the risk of lost data. Log files are compressed, labeled, and stored on a daily basis, and a new "current" log file is created. Using this mechanism, no audit data is lost during the storage transition.

The amount of available audit storage space is monitored very closely via the rollaudit and logcheck utilities. Those utilities monitor the log file size and rotate log files as needed.

Learn more about the Sidewinder audit process in the following sections:

- [Audit components](#)
- [Audit file names](#)
- [Understanding audit messages](#)
- [Tools for viewing and customizing audit events](#)
- [Supported log file formats](#)
- [Exporting audit data to Firewall Reporter and syslog servers](#)

For information on using rollaudit, see [Monitoring disk space using cron jobs](#). For information on using the logcheck utility, refer to the logcheck man page.

Audit components

There are three main components to the Sidewinder audit process:

- **auditd** – This is the audit logging daemon. This daemon listens to the Sidewinder audit device and writes the information to log files. The log files provide a complete record of audit events that can be viewed by an administrator. By default, **auditd** sends all audit data to a binary file called */var/log/audit.raw*.
- **auditbotd** – This is the daemon that listens to the audit device and gathers the security-relevant information it finds. It tracks these events and uses its configuration to determine when the data might be indicating a problem and require a response, such as an attempted break-in. If it does detect an audit event that has a configured response, the firewall responds accordingly. For more information on configuring IPS attack and system event responses, refer to [Chapter 13, IPS Attack and System Event Responses](#).
- **auditdbd** – This daemon maintains the audit database. auditdbd monitors the audit stream and sends reporting information to an audit database. The auditdbd server is disabled by default.

To use the on-box reporting service (**cf reports**), you must first enable the following components by entering the following commands:

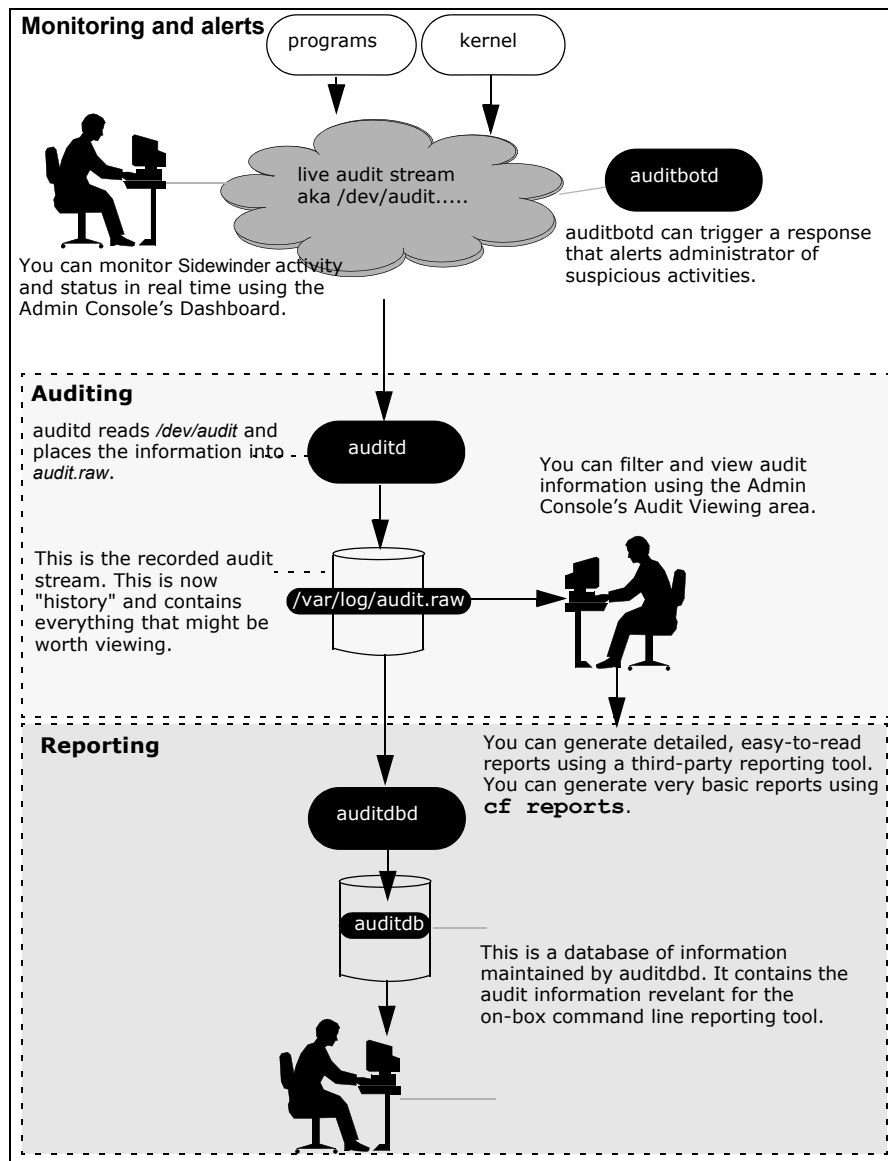
```
cf daemond enable agent=auditsql
cf daemond enable agent=auditdbd
```

Note: The auditsql agent must be enabled before the auditdbd agent.

If you are not using the Sidewinder on-box reporting tool, leave these agents disabled.

The following diagram demonstrates how these pieces are related in the audit flow.

Figure 197 Audit flow



Audit file names

The audit information is saved in a binary format in the `/var/log/audit.raw` files. When the file is rolled, a timestamp is appended to the file name. The easiest method for viewing the contents of the `audit.raw` files is to use the Admin Console's Audit Viewing window. Refer to [Viewing audit information](#).

Tip: If you prefer to view the file contents via command line, refer to the **acat** and **showaudit** man pages.

Audit log files use one of two file suffixes:

- ***.gz** – This suffix is for files in compressed format. These files may be decompressed using **acat** or **showaudit**. The default file name format is `audit.raw.YYYYMMDDhhmmssZZZ.YYYYMMDDhhmmssZZZ.gz`, where the variables represent date and time (including time zone) of the beginning and end of that audit file's contents. For example, `20051231020000CST.20060101020000CST.gz` is a file that contains audit data from December 31, 2005 at 2:00 am to January 1, 2006 at 2:00 am.
- ***.raw** – This suffix is for files in raw audit format. These are binary formatted files that can be viewed in ASCII format using the Admin Console or command line.

Understanding audit messages

When viewing audit messages in the Admin Console, the form may vary depending on the purpose and content of the message. The form of the first two lines is the same for all audit messages and provides general information about the process generating or causing the audit. The third line will vary but usually includes Type Enforcement information and possibly some additional information. The other lines of an audit message will vary depending on the type of audit message.

Note: To view audit message files, see [Viewing audit information](#).

The message below is an example of a Type Enforcement audit message (using the **te_filter** filter). The numbers have been added to link the example line with the bullets below.

```
(1)Nov 22 11:38:46 2006 EST  f_kernel a_tepm t_attack p_major
(2)pid: 11124 ruid: 100 euid: 100 pgid: 11124 logid: 100 cmd: 'cat'
(3)domain: User edomain: User hostname: python.a.net category: policy_violation
(4)event: ddt violation srcdmn: User filedom: Pass filetype: file
(5)reason: OP: OP_FS_PERM_CHECK perm wanted: 0x1<read> perm granted: 0x0
(6)information: open /etc/spwd.db
```

- **Line 1** – This line lists the date and time, the facility that audited the message (such as the Kernel, FTP, or Telnet), the location (known as the *area*) in the facility that audited the message (such as general area or type enforcer), the type of audit message (such as attack, Type Enforcement violation, or access control list) and the priority of the message (such as major or minor).
- **Line 2** – This line lists the process ID, the real user ID, the effective user ID, the process group ID, the log ID, and the command associated with the process ID.
- **Line 3** – This line lists the real domain the process is running in and the effective domain (the domain of the process for which permission is given). It also lists the firewall's host name and the audit event's category.
- **Lines 4, 5, and 6** – The fourth line contains the integer representation of the permissions requested by the process and granted to the process, the domain of the requesting process, and the type of file that the process is requesting access to. The last two lines often contain the reason the audit event was generated and any additional information.

Tools for viewing and customizing audit events

Many tools allow you to interact with the audit data. Use the following tools to generate, view, and respond to audit events:

- Audit output can be customized using Network Defenses (**Policy > Network Defenses**).
- Audit output can be viewed using these tools:
 - The Admin Console's dashboard
 - The Admin Console's audit viewing area (**Monitor > Audit Viewing**)

- The Admin Console's configuration backup area (**Maintenance > Configuration Backup**)
 - Audit output can be viewed in these formats:
 - Sidewinder Export Format (SEF)
 - WebTrends Log Format
 - HTTP
 - ASCII
 - Verbose ASCII
 - XML
- See [Supported log file formats](#) for more information.
- Audit output can be configured to trigger alerts using these tools:
 - IPS Attack Responses (**Monitor > IPS Attack Responses**)
 - System Responses (**Monitor > System Responses**)

Supported log file formats

Table 34 lists the log formats the firewall supports, as well as some uses for each format and other important information.

Table 34 Supported log formats and their uses

Format	Use	Comments
Sidewinder Export Format (SEF)	Firewall Reporter, various third-party tools	SEF is the format used when exporting logs to Firewall Reporter. Note: Firewall Reporter is a legacy feature and is not actively supported. If using SmartFilter and SEF, set the audit level on the appropriate HTTP proxy rules to Verbose (Policy > Rules) . Note: SmartFilter is not actively supported in Sidewinder 7.x.
WebTrends Extended Logging Format (WELF)	WebTrends® reporting tools	
W3C Extend Log Format (HTTP)	various third-party reporting tools	If using this format, set the audit level on the appropriate HTTP proxy rules to Verbose (Policy > Rules) .
Extensible Markup Language (XML)	various third-party reporting tools	
Binary or RAW (bin)	various third-party reporting tools	Using the acat command is optional as this output is an exact copy of the audit raw file.
American Standard Code of Information Interchange (ascii)	various third-party reporting tools	ASCII is the standard format and therefore does not require any arguments with acat.
Verbose American Standard Code of Information Interchange (vascii)	various third-party reporting tools	

Viewing audit information

Use the Audit Viewing window to monitor the activity on your Sidewinder. You can filter the records to focus on the information you want to see, and you can transfer audit records to different locations.

To view audit records, select **Monitor > Audit Viewing**. The Audit Viewing window appears.

Figure 198 Audit Viewing window

Audsits		Description	Date	Time	Syslog	Type
Common						
All audit		Detects all attack and system events, re	Oct 10 2008 EDT	11:11:41	Debugging (7)	geninfo
Attack events		This detects all severities of Application D	Oct 10 2008 EDT	11:11:41	Debugging (7)	lcm
Configuration changes		Detects when the Sidewinder's configura	Oct 10 2008 EDT	11:11:41	Debugging (7)	interface
System events		Detects all system events of all severitie	Oct 10 2008 EDT	11:11:41	Debugging (7)	interface
TrustedSource		Detects attacks identified as spam by Tru	Oct 10 2008 EDT	11:11:41	Debugging (7)	interface
VPN		Detects VPN audit events.	Oct 10 2008 EDT	11:11:41	Debugging (7)	interface
Custom			Oct 10 2008 EDT	11:10:41	Debugging (7)	geninfo
Advanced			Oct 10 2008 EDT	11:10:41	Debugging (7)	lcm
			Oct 10 2008 EDT	11:10:41	Debugging (7)	interface
			Oct 10 2008 EDT	11:10:41	Debugging (7)	interface
			Oct 10 2008 EDT	11:10:41	Debugging (7)	interface
			Oct 10 2008 EDT	11:10:41	Debugging (7)	interface
			Oct 10 2008 EDT	11:10:00	Informational (6)	info
			Oct 10 2008 EDT	11:09:48	Notifications (5)	important
			Oct 10 2008 EDT	11:09:47	Notifications (5)	auth_attempt
			Oct 10 2008 EDT	11:09:47	Informational (6)	cfg_change
			Oct 10 2008 EDT	11:09:41	Debugging (7)	geninfo
			Oct 10 2008 EDT	11:09:41	Debugging (7)	lcm
			Oct 10 2008 EDT	11:09:41	Debugging (7)	interface
			Oct 10 2008 EDT	11:09:41	Debugging (7)	interface
			Oct 10 2008 EDT	11:09:41	Debugging (7)	interface
			Oct 10 2008 EDT	11:09:41	Debugging (7)	interface

- The left pane lists audit filters. Select a filter and the audit records returned by that filter appear in the right pane.
 - Use predefined common and advanced filters.
 - Create custom filters.
 - View audit records in real time or restrict them to a specified time span.


See [Filtering audit data](#) for more information on using filters.

- The right pane displays the audit records of the selected filter.
 - View details of each audit record.
 - Export or copy audit records to off-box locations.
 - Modify how the records appear in the list.

See [Viewing and transferring audit records](#) for more information on working with audit records.

- The toolbar has controls to modify filters and records.

Figure 199 Audit Viewing toolbar



Use the toolbar to perform the following actions:

Table 35 Audit Viewing toolbar

Icon	Action
Pause audit/Continue audit	To temporarily stop records from loading, click Pause audit . To resume loading records, click Continue audit .
Export audit	Export audit data to a file that can be viewed and printed by clicking Export audit . A pop-up window appears where you set the appropriate properties.
Set time span	View audit information in real time or for a specific time span by clicking Set time span .
Filter	<p>The Filter field shows the expressions that make up the selected filter.</p> <ul style="list-style-type: none"> Add commonly used expressions to the field by clicking Expression and selecting from the drop-down list. Select Advanced from the drop-down list for a wider selection. Edit more complex filter expressions in a pop-up window by clicking Filter. See the Filter field history by clicking the down arrow—the last ten filters appear in reverse order. Clear the field by clicking Clear. Display or refresh the audit records for the selected filter by clicking Apply.
Open as new application	<p>Open the Audit Viewing window as a separate application by clicking Open as new application. Enter the same administrator user name and password you use to log in to the Sidewinder.</p> <p>Use this feature to view audit while performing other tasks on the firewall.</p>
Color settings	Modify the audit record's on-screen appearance by clicking Color settings . Use the pop-up window to select predefined text and background color schemes or create custom color schemes.
Columns	Select which columns appear in the audit record pane by clicking Columns . Use the pop-up window to add or remove columns.

Filtering audit data

Audit filters control the audit information that you want to see by displaying or excluding certain types of audit records. Using filters can greatly reduce your audit output and simplify troubleshooting.

The left pane of the Audit Viewing window lists predefined and custom audit filters. Select a filter in the list that corresponds to the audit output you want to see, and any audit records matching the filter's conditions appear in the right pane.

You can perform these tasks to define and modify audit filters:

- [Define a time span for audit filters](#)
- [Modify a predefined filter](#)
- [Create a custom filter](#)

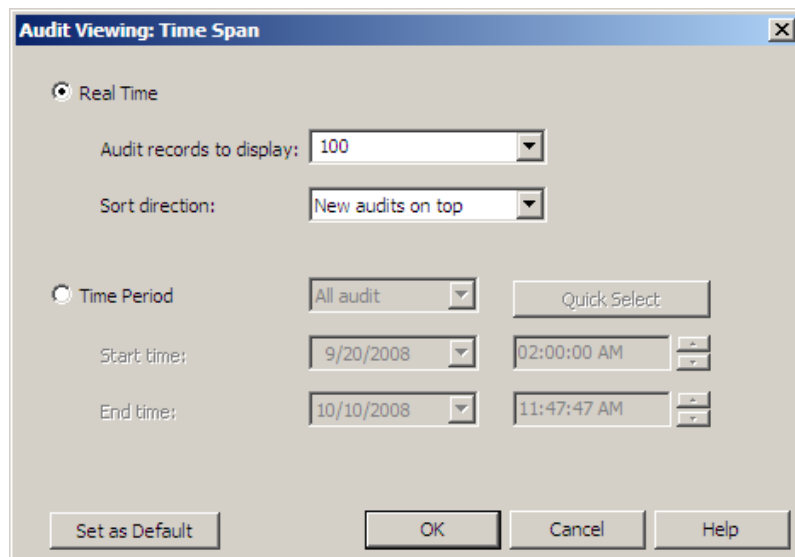
Define a time span for audit filters

Audit records are filtered for time parameters you specify. You can view audit records in real time or filter records within a designated time period.

To define a time span:

- 1 In the left pane, select an audit filter.
- 2 On the toolbar, click **Current time span**. The Time Span window appears.

Figure 200 Audit Viewing: Time Span window



You can make the following selections:

- **Real Time** – Select this option to view streaming audit data in real time. Use the drop-down lists to select how many records to display and in what order.
Note: A selection of **Unlimited** can impact firewall performance.
- **Time Period** – Select a preset time span from the drop-down list, or select **Custom** and then select a start and end time. You can click **Quick Select** to modify custom dates and times from the pop-up menus.
- **Set as Default** – Click this to save the selected time span as the default time the next time you open the Audit Viewing window. If you select **Real Time**, the **Audit Records to display** and **Sort direction** settings are set as defaults. This option is not available for a custom time span.

Modify a predefined filter

Common and Advanced filters are predefined on the Sidewinder to give you convenient access to frequently used filters. You can modify predefined filters to further refine the records that are displayed. You can also save a predefined filter as a new custom filter.

- You cannot delete predefined audit filters.
- For an explanation of the event types that the predefined filters audit, see [About the predefined audit filters](#).

- For an explanation of filter expressions, see [About filter syntax](#).

To modify a predefined filter:

- 1 In the left pane, select an audit filter.
- 2 On the toolbar, click **Filter**. The Current Filter window appears.

Figure 201 Audit Viewing: Current Filter window

Audit Viewing: Current Filter

Source burb:

Source IP address:

Destination burb:

Destination IP address:

Service:

Ticket ID:

☒ Custom:

Expression | Validate

(type AUDIT_T_ATTACK)

☐ Save as new filter

OK Cancel Help

Note: You can also use the toolbar to modify filters. See [Table 35](#) for information.

The current filter appears in the lower field. You can perform these actions:

- Modify the filter in the lower field.
 - Undo or redo a change to the filter by clicking **Undo** or **Redo**.
 - Add commonly used expressions to the filter by clicking **Expression** and selecting from the drop-down list. Select **Advanced** from the drop-down list for a pop-up window with a wider selection of expressions.
 - Confirm the corrected filter by clicking **Validate**. If the syntax is incorrect, a red underline appears under the invalid part of the expression. If the syntax is correct, no red underline appears.
- Make simpler changes to the filter: Clear the **Custom** check box and make selections and entries in the upper part of the window.
 - **Source burb** – To view audit records generated by a source burb, select a burb from the drop-down list.
 - **Source IP address** – To view audit records generated by a specific source IP address, enter an IP address. You can also enter the number of significant bits needed to create the subnet you want to filter.
 - **Destination burb** – To view audit records generated by a destination burb, select a burb from the drop-down list.
 - **Destination IP address** – To view audit records generated by a specific destination IP address, enter an IP address. You can also enter the number of significant bits needed to create the subnet you want to filter.
 - **Service** – To view audit records generated by a service, select a service from the drop-down list. The list includes all configured services.
 - **Ticket ID** – To view audit records generated during a change ticket, select this check box and enter the ticket name.
- Create a new filter from this filter: After you have made the desired modifications, select **Save as new filter** and then click **OK**. The New Filter window appears. See [Create a custom filter](#) for information on completing this window.

Create a custom filter

Audit filters that you create appear under the **Custom** group in the Audit Viewing window. You can also modify and delete custom filters.

- You can modify a predefined filter and save it as a custom filter. See [Modify a predefined filter](#) for more information.
- For an explanation of filter expressions, see [About filter syntax](#).

To create a custom audit filter, right-click an existing filter and select **New filter** from the pop-up menu. The New Filter window appears.

Figure 202 Audit Viewing: New/Modify Filter

The screenshot shows a window titled "Audit Viewing: New Filter". It contains the following elements:

- Name:** A text input field.
- Description:** A text input field.
- Filter Type:** Two radio buttons, "Attack filter" (which is selected) and "System filter".
- SNMP trap:** A text input field containing the value "0".
- Expression/Validate Section:** A tabbed interface with "Expression" and "Validate" tabs. The "Expression" tab is active, showing a text area with the text `(type AUDIT_T_ATTACK)`.
- Buttons:** "OK", "Cancel", and "Help" buttons at the bottom right.

- You can right-click a predefined filter or a custom filter to create a new custom filter.
- The expressions of the selected filter appear in the New Filter window. If you want to use certain expressions as building blocks for a custom filter, right-click a similar filter.

You can make these entries and selections:

- **Name** – Enter an easily identifiable name for the filter. (You cannot modify the filter name.)
- **Description** – You can optionally enter a description to further distinguish the filter.
- **Filter type:**
 - **Attack filter** – Select this option for the filter to appear as an event in the IPS Attack Responses window. The description appears in this window.
 - **System filter** – Select this option for the filter to appear as an event in the System Responses window. The description appears in this window.
- **SNMP trap** – If you want to send an alert message in case of an audit event, enter the number that corresponds to the trap on your SNMP station. If the entry is **0**, no trap is sent.
- Enter or modify the filter in the lower field.
 - Undo or redo a change to the filter by clicking **Undo** or **Redo**.
 - Add commonly used expressions to the filter by clicking **Expression** and selecting from the drop-down list. Select **Advanced** from the drop-down list for a pop-up window with a wider selection of filters.
 - Confirm the corrected filter by clicking **Validate**. If the syntax is incorrect, a red underline appears under the invalid part of the expression. If the syntax is correct, no red underline appears.

About the predefined audit filters

Use the tables below to see lists of predefined filters and descriptions of the event types that each filter audits.

- [Table 36](#)
- [Table 37](#)

Table 36 Common predefined audit filters

Audit types	Description
All Audit	Detects all attack and system events, regardless of type.
Attack All	Detects attack events of all severities. This option also detects all severities of Application Defense violation attacks, buffer overflow attacks, DOS attacks, general attacks, policy violation attacks, protocol violation attacks, virus attacks, and spam attacks.
Config Change	Detects when the Sidewinder's configuration changes.
System All	Detects all system events of all severities, including power failures, hardware and software failures, failover events, license expiration, host license exceeded, log overflows, and IPsec errors.
TrustedSource	Detects attacks identified as spam by TrustedSource.
VPN	Detects VPN audit events.

The following list displays the default Advanced audit filters and describes the event types that each filter audits.

Table 37 Advanced predefined audit filters

Audit types	Description
Access Control List	Detects all ACL audit events.
ACL Allow	Detects when a connection is allowed by a rule in the active policy.
ACL Deny	Detects when a connection is denied by a rule in the active policy.
Application Defense Violation All	Detects attacks of all severities that violate active policy defined by Application Defenses. This attack category includes mime and keyword filter failure attacks.
Application Defense Violation Severe	Detects when severe attacks violate active policy defined by Application Defenses, including mime and keyword filter reject audits. Severe attacks indicate something is occurring that an administrator should know.
Attack Severe	Detects severe attacks, including Application Defense violation attacks, buffer overflow attacks, general attacks, DOS attacks, policy violation attacks, protocol violation attacks, and virus attacks
Buffer Overflow Attack	Detects attempted buffer overflow attacks targeted at systems protected by the Sidewinder.
Denied Authentication	Detects when a user attempts to authenticate and enters invalid data. For example, if a user is required to enter a password and entered it incorrectly, the denied auth event would log the event.
DOS All	Detects Denial of Service attacks of all severities. This attack category also detects all severities of TCP SYN attacks and proxy flood attacks.
DOS Severe	Detects severe Denial of Service attacks. This attack category also detects TCP SYN attacks and proxy flood attacks. Severe attacks indicate something is occurring that an administrator should know.
Error	Detects all system events identified as AUDIT_T_ERROR in the audit stream.
General Attack All	Detects general attacks of all severities that do not fall into the predefined categories.
General Attack Severe	Detects severe general attacks that do not fall into the predefined categories. Severe attacks indicate something is occurring that an administrator should know.
HA Failover	Detects when a failover IP address changes because a High Availability cluster failed over to its secondary/standby.
Hardware Software Failure	Detects some hardware failures, such as RAID, hard drive, and AMIR monitor failures.
Host License Exceeded	Detects when the number of hosts protected by the Sidewinder exceeds the number of licensed hosts.
IPFilter Deny	Detects when a connection is denied by the active IP Filter policy.
IPsec Error	Detects when traffic generates IPsec errors.
Keyword Filter Failure	Detects when an SMTP mail message is rejected due to a configured keyword filter.
License Expiration	Detects when a licensed feature is about to expire.
Log Overflow	Detects when the log partition is close to filling up.
Network Probe	<p>Detects network probe attacks, which occur any time a user attempts to connect or send a message to a TCP or UDP port when the security policy does not include a service that is expecting to receive traffic on that port.</p> <p>Note: The firewall does not blackhole netprobe attacks, as they are likely to be Denial of Service attacks from spoofed source addresses.</p>
Network Traffic	Detects all connections that successfully pass through the Sidewinder.
Not Config Change	Detects all attack and system events that are not configuration changes.
Policy Violation All	Detects attacks of all severities that violate the active policy. This attack category also detects all severities of failed authentication attacks, ACL and IP Filter deny attacks, and Type Enforcement error attacks.
Policy Violation Severe	Detects severe attacks that violate the active policy. This attack category also detects failed authentication attacks, ACL and IP Filter deny attacks, and Type Enforcement error attacks. Severe attacks indicate something is occurring that an administrator should know.
Power Failure	Detects when an Uninterruptible Power Supply (UPS) device detects a power failure and the system is running on UPS battery power.
Protocol Violation All	Detects attacks of all severities that violate protocol compliance.
Protocol Violation Severe	Detects severe attacks that violate proxy protocols (HTTP, Telnet, FTP, etc.). Severe attacks indicate something is occurring that an administrator should know.
Proxy Flood	Detects potential connection attack attempts. A connection attack is defined as one or more addresses launching numerous proxy connection attempts to try and flood the system. When NSS receives more connection attempts than it can handle for a proxy, new connections to that proxy are briefly delayed (to allow the proxy to "catch up"), and the attack is audited.

Table 37 Advanced predefined audit filters <Comment>(continued)

Audit types	Description
Signature IPS Intrusion All	Detects all attacks identified by the signature-based IPS. This category detects attacks that were denied, dropped, or rejected, as well as suspected attacks that were allowed but were audited by IPS.
Signature IPS Intrusion Blackholed	Detects attacks identified by the signature-based IPS where the attacker was blackholed.
Signature IPS Intrusion Deny	Detects attacks identified by the signature-based IPS where the offending network session was dropped or rejected, or the attacker was blackholed.
Spam	Detects attacks of all severities that are spam.
Spam Severe	Detects severe attacks that are spam.
Syslog	Detects all audit attacks and system events created via syslog.
System Critical	Detects all critical system events, including power failures, hardware failures, critical software failures, and failover events. Critical system events indicate a component or subsystem stopped working, that the system is going down (expectedly or unexpectedly), or that the system is not expected to work again without intervention.
System Critical And Severe	Detects critical and severe system events including power failures, hardware failures, critical and severe software failures, failover events, license expiration, log overflows, and IPsec errors. Critical system events indicate a component or subsystem stopped working, that the system is going down (expectedly or unexpectedly), or that the system is not expected to work again without intervention. Severe attacks indicate something is occurring that an administrator should know.
TCP SYN attack	Detects a possible attempt to overrun the firewall with connection attempts.
Type Enforcement	Detects when there is a TE violation due to an unauthorized user or process attempting to perform an illegal operation.
UPS System Shutdown	Detects when a UPS is running out of battery power or has been on battery power for the estimated battery time.
Virus	Detects attacks of all severities that are viruses.
Virus Severe	Detects severe attacks that are viruses.

About filter syntax

Use the following syntax when building expressions:

- Identify a filter using either single quotes (') or double quotes ("). All examples shown below use single quotes.
- Express "and" using either **and** or **&&**.
- Express "or" using either **or** or **||**.
- Express "not" using either **not** or **!**.

A filter should include:

- The *type* or *facility* you want to search for, using one of these formats:
 - The Name format (AUDIT_T_*TYPE* as in AUDIT_T_ATTACK, AUDIT_F_*FACILITY* as in AUDIT_F_LOGIN)
 - The Short Message format (attack, login)
 - The Short Message format prepended with classification indicator (t_attack, f_login)

Note: This last format appears in audit records and is useful when copying or pasting directly from audit output.

- Additional fields to further specify the audit results; fields can be separated by Boolean operators (and, or, not) and grouped by parentheses

Example

This filter expression:

```
dest_burb external and (src_ip 10.69.101.34 or src_ip 10.69.101.36)
```

returns this audit record:

```
Aug 22 02:02:20 2008 CDT f_ping_proxy a_proxy t_nettraffic p_major
pid: 3728 ruid: 0 euid: 0 pgid: 3728 logid: 0 cmd: 'pingp'
domain: Ping edomain: Ping hostname: mixer.ext.b.test
event: proxy traffic end service_name: ping netsessid: 48ad640e000e0151
srcip: 10.69.101.34 srcburb: internal protocol: 1 dstip: 10.66.6.22
dstburb: external bytes_written_to_client: 83079240
bytes_written_to_server: 83087396 acl_id: Internet Services cache_hit: 1
request_status: 0 start_time: Thu Aug 21 07:48:14 2008
```

A source IP address of 10.69.101.34 and an external destination burb matches the filter expression

Viewing and transferring audit records

View audit records to monitor the activity on your Sidewinder.

The right pane of the Audit Viewing window displays the audit records filtered by the selection in the left pane. Each audit record appears as a single row in the table. You can view audit records on-screen or export or copy the data to another location.

Note: Some audit types will not contain information for each table column. If a column is blank, that type of information does not apply to that particular audit record.

You can perform these tasks to view and transfer audit records:

- [View and copy audit record details](#)
- [Export audit records](#)
- [Add or remove columns in the audit records table](#)
- [Configure on-screen color schemes for audit records](#)

A high volume of audit records can affect firewall performance. You can take these steps to lessen the impact of a large audit record list:

- Click **Pause audit** on the toolbar to temporarily stop the records from loading. Click **Continue audit** to resume loading the records.
- Filter the audit records to reduce the number of records that are displayed: In the right pane, right-click a column head and select a value to filter the list.

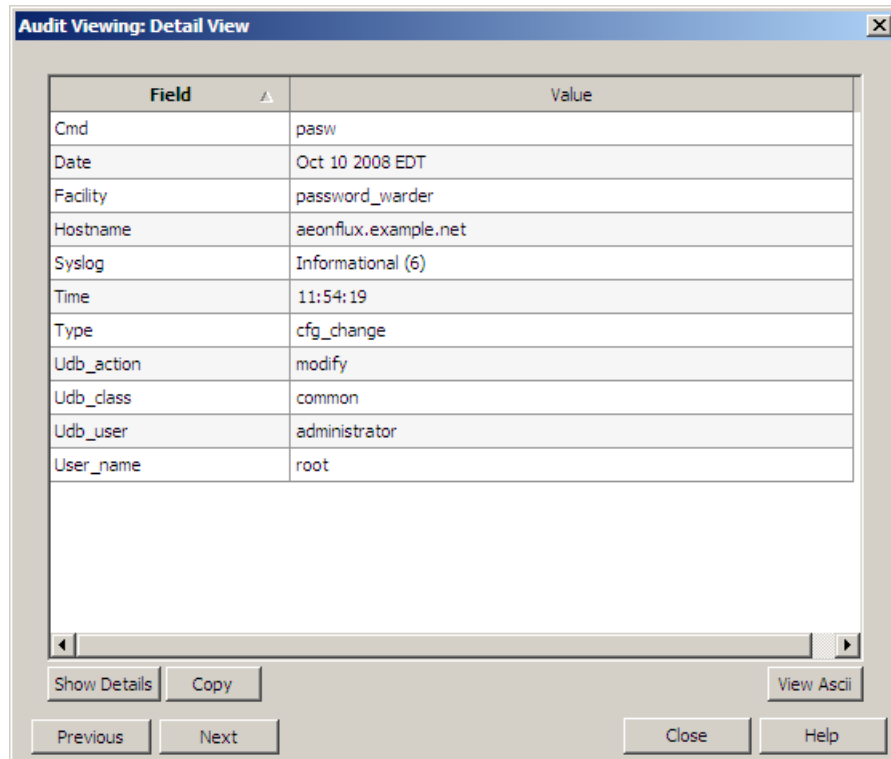
View and copy audit record details

You can view details about each record that appears in the audit records table. You can copy the record details to a document or spreadsheet or view them in ASCII format.

To see details about an audit record:

In the records table, double-click an audit record. (Each audit record appears as a single row in the table.) The Detail View window appears.

Figure 203 Audit Viewing: Detail View window



You can perform these actions:

- See more fields by clicking **Show Details**. See fewer fields by clicking **Hide Details**.
- To view the man page for the **Cmd** field value, right-click the row and select **Investigate binary** from the pop-up menu.
- Copy the data to a document or spreadsheet:
 - Click **Copy** and select **As text** from the pop-up menu, and then paste the data to a document.
 - Click **Copy** and select **As table** from the pop-up menu, and then paste the data into a spreadsheet.
- To see the record data in ASCII format, click **View Ascii**.
- To move through the audit records table and view details for other records without closing this window, click **Previous** or **Next**. The selected audit record details appear in this window.

Export audit records

You can export audit records to another location, where they can be printed, viewed directly, or opened in a reporting or editing tool.

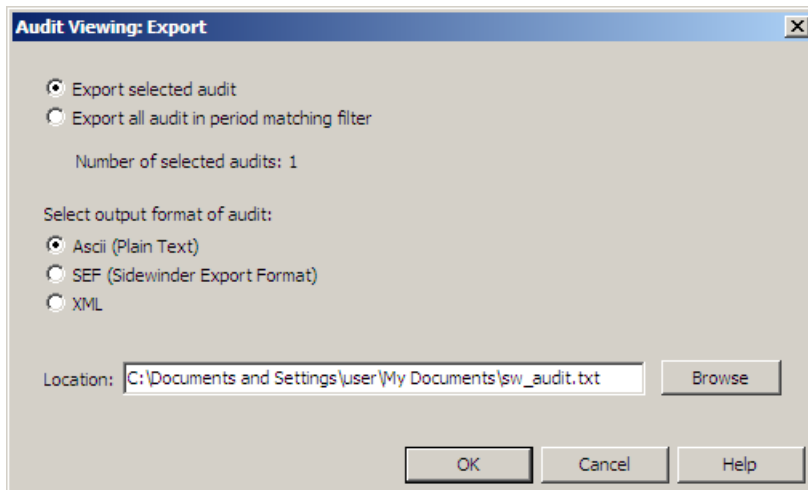
To export audit records:

- 1 In the right pane, select one or more audit records.

To select multiple records, press the **Ctrl** key as you select each record. To select multiple consecutive records, press the **Shift** key as you select the first and last record.

- 2 On the toolbar, click **Export audit**. The Export window appears.

Figure 204 Audit Viewing: Export window



- 1 Select the records you want to export.

- **Export selected audit** – Select this option to export audit records you have selected in the right pane of the Audit Viewing window.
- **Export all audit in period matching filter** – Select this option to export audit records created during the period selected on the Time Span window.

Note: Exporting all records matching a time period can take a significant amount of time and disk space.

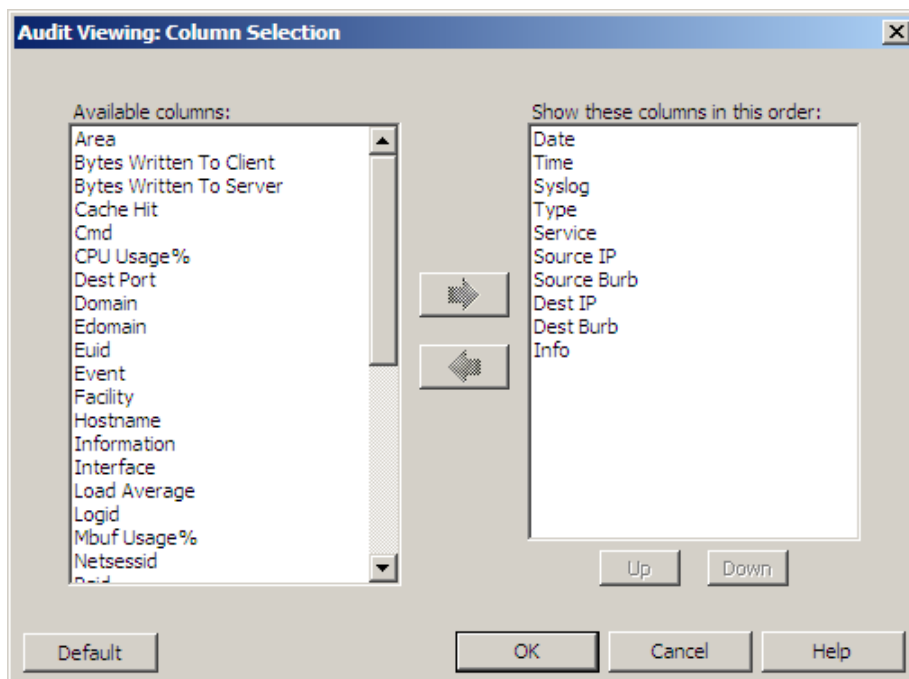
- 2 Select the output format for the records.
- 3 Click **Browse** and navigate to the location you want the audit record saved to.
- 4 Click **OK**. The Export window closes and the audit record is saved in the location you specified.

Add or remove columns in the audit records table

Further refine the information that appears in the audit records table by selecting the table's columns.

In the toolbar, click **Columns**. The Column Selection window appears.

Figure 205 Audit Viewing: Column Selection window



You can perform these actions:


- To add columns to the records table, select a column in the **Available columns** list and click the right arrow to move it to the **Show these columns in this order** list.
- To remove a column, select it in the **Show these columns in this order** list and click the left arrow.
- Select multiple columns by pressing and holding the CTRL key while selecting the appropriate columns.
- Select a range of columns by selecting the first column in the range, pressing and holding the SHIFT key, and then selecting the last column in the range.
- To order the columns in the table, select a column in the **Show these columns in this order** list and click **Up** and **Down** to move it to the desired location. The top-to-bottom columns in the list appear from left to right in the table.
- To return the displayed columns to the firewall's default format, click **Default**.

Configure on-screen color schemes for audit records

You can configure the on-screen color scheme of the audit records to easily identify types of records or for other organizational purposes.

On the toolbar, click the **Color settings** button. The Color Settings window appears.

Figure 206 Audit Viewing: Color Settings window



- **System colors**, **Minimal color**, and **Color** are preset color schemes. When you select one of these options, the table shows the background color, text color, and an example that corresponds to each Severity type. (The Severity type appears in the Syslog column of the Audit Viewing window.)
- To create a custom color for a Severity type:
 - Select **Minimal color**, **Color**, or **Custom**.**
 - Click the table cell for the Background Color or Text Color of the desired type. The system color window appears.**
 - Select a basic color or move the crosshair and slider bar to create a custom color.**
 - Click **OK**. The **Custom** option is selected and your color selection appears in the table.**

Click **OK** to close the Color Settings window. You must click **Apply** for the new color scheme to take effect.

Managing log files

Use the Audit Management window to manage your audit log files, including:

- Exporting log files in a variety of formats to a specified host
- Scheduling exports
- Adding a signature to the log files
- Rolling the log files
- Identifying changes using change tickets

Generally, you will set up this service during system startup, then test to make sure you are getting the results you intended. Once setup is complete, the log files transfer and roll automatically, giving you the audit data you need and keeping the Sidewinder running freely.

Select **Monitor > Audit Management**. The Audit Management tab appears.

Note: Use the Firewall Reporter/Syslog tab to configure the export of audit data to designated syslog servers. For more information, see [Exporting audit data to Firewall Reporter and syslog servers](#).

Figure 207 Audit Management tab

Audit Management | Firewall Reporter/Syslog

Audit Options

- ☒ Show system statistics in audit log
- ☐ Require change ticket
- ☐ Create backups before each change ticket

Number of automatic backups to keep: 50

Logfile Options

+ [Icons] Find: [Text Box]

Export Entry Name	Type	Summary
-------------------	------	---------

Export logfiles: disabled (35 minutes after every hour.) [Change] [Export All Now]

Roll logfiles: daily at 2:00am. [Change] [Roll Now]

☐ Delete logs after export

Signature Options

- ☐ Sign exported files
- ☐ Append signature to exported file ☐ Put signature in separate file
- Sign with: <None>

Use this tab to capture network and system utilization statistics, to automatically create change tickets, and to configure and schedule the export of audit log files.

The Audit Management window contains two panes:

- [Audit Options pane](#)
- [Logfile Options pane](#)

Audit Options pane

The Audit Options pane contains settings to capture network and system utilization statistics and to automatically create change tickets to identify changes made to the firewall.

- **Show system statistics in audit log** – Select this option to capture network and system utilization statistics, which appear in the dashboard.

Note: This option is enabled by default and should rarely, if ever, need to be disabled.

- **Require change ticket** – Select this option to automatically open the Change Ticket window when you save changes made to the firewall. A change ticket identifies specific changes to the firewall.

To view change ticket audit records, use the Current Filter window. See [Modify a predefined filter](#) for more information.

- **Create backups before each change ticket** – Select this option to automatically create a configuration backup when you start a change ticket.

- A Lite configuration backup is created of the firewall state before the ticket was started.
- A Lite backup does not contain the home directories or support bundle, making it smaller than a full configuration backup.

Enter the number of automatic backups that you want to keep.

Note: You can view audit records for this backup on the Configuration Backup window. See [Manage configuration backup files](#) for more information.

Logfile Options pane

From the Logfile Options pane, you can create export entries that allow the Sidewinder to transfer its log files in a variety of formats to a specified host. From this pane, you can also schedule the exports, include a signature, and roll log files.

Use the toolbar to perform the actions described in [Table 38](#).

Figure 208 Logfile Options toolbar




Table 38 Logfile Options toolbar tasks

Button	Action
New	Click New to create an export entry. Complete the fields as described in Creating or modifying an export entry . Click OK ; the entry appears in the Logfile Options pane.
Modify	Double-click the entry you want to change (alternately, click once to highlight the entry, then click Modify). Make your changes, click OK , and save your changes.
Delete	Select the entry you want to delete, then click Delete . Click Yes to delete the entry or No to cancel the action.
Export selection now	Click Export selection now to immediately export a selected entry.
Search	To find an export entry, enter all or part of the name. When the system finds a match, it appears highlighted in the pane. If the system does not find a match, the pane appears blank. Use the Backspace key to find partial matches or delete the search term to return to the Logfile Options pane.

Creating or modifying an export entry

On the Audit Management window toolbar, click **New** or **Modify**. The Export File window appears.

Figure 209 Export File window

Export: Export File

Entry Name:

Export Type

☒ Sidewinder Export Format ☐ WebTrends
☐ HTTP ☐ XML
☐ Binary ☐ ASCII
☐ Verbose ASCII

Export with:

Host:

Directory:

Username:

Password:

OK Cancel Help

- **Entry Name** – Enter a descriptive, single-word name.
- **Export Type** – Select the export type. For more information about supported log formats, refer to [Supported log file formats](#).
- **Export with** – From the drop-down list, select **FTP** or **SCP** transfer protocol.
Note: We recommend using the SCP protocol if it is supported by the destination host.
- **Host** – Enter the host name or IP address of the host that will receive the exported file.
- **Directory** – Enter the name of the directory that will store the exported file.
- **Username** – Enter the username for the host you specified.
- **Password** – Enter the password for the host you specified.

Once you have created the export entries, test them to make sure the results are what you intended. See [Signing export files](#) and [Exporting and rolling log files](#).

Signing export files

Log files can be cryptographically signed to ensure data integrity. To add a signature:

- 1 In the Logfile Options pane, select the **Sign exported files** check box.
- 2 In the **Sign with** field, use the drop-down list to select the signature certification.
- 3 In the **Signature Options** area, select how you want to store the signature file:
 - **Append signature to file** – This option creates one .gz file that includes the signature at the end of the file.
 - **Put signature in separate file** – This option creates two files: a .gz file that contains the actual audit and a .gz.pem file that contains the signature.

Exporting and rolling log files

Once you configure and enable a schedule, the Sidewinder will automatically check to see if it should export any log files and, if so, export those files. You can also export log files on request for a single export entry or all entries.

The firewall automatically rolls log files every morning at 2:00 a.m. You can change the schedule and export or roll log files on request. By default, the Sidewinder maintains 20 rolled instances of the audit.raw file. This setting can be reconfigured in the `/etc/sidewinder/rollaudit.conf` file.


Configure a schedule for exporting or rolling log files

Use the Crontab Editor to schedule an export program.

To configure a schedule for the Sidewinder to export log files:

- 1 Click **Change**. The Crontab Editor window appears.

Figure 210 Crontab Editor window



The screenshot shows the 'Export: Crontab Editor' window. It has a title bar with a close button. Inside, there's a 'Frequency' dropdown set to 'Hourly' and an 'Enable' checkbox. Below that is a spinner for '35 minutes after the hour'. There's a 'Custom' checkbox and five input fields for 'Minute', 'Hour', 'Day of month', 'Month', and 'Day of week', each with a default value of '*' or '35'. A 'Command' field contains 'cf export latest' and a 'Description' field contains 'Run export utility 35 minutes past every hour'. At the bottom are 'OK', 'Cancel', and 'Help' buttons.

- 2 Select the **Enable** check box to activate this schedule. If you leave the check box clear, the entry will be saved, but the Sidewinder will not act on it.

3 [Conditional] To designate a standard frequency for exporting files (for example, every day at 2:00 a.m.):

- **Frequency** – From the drop-down list, select the frequency for exporting the file (hourly, daily, or weekly).
 - If you selected Hourly, enter the number of minutes after the hour.
 - If you selected Daily, enter the time for export.
 - If you selected Weekly, enter the time and day. You can select multiple days.
- **Description** – Enter a descriptive name for the task (such as Run export utility 35 minutes past every hour).

4 [Conditional] To define a custom frequency for exporting files:

- **Custom** – Select this check box and complete the fields. Refer to `man 5 crontab` for options.

Note: The Crontab Editor allows custom syntax. Make sure your syntax is correct, and verify your entry with `cf crontab query`.

- **Description** – Enter a descriptive name for the task (such as Run export utility the 1st and 15th day of every month at 2:00 a.m.).

5 Click **OK** to accept the schedule.

Export or roll log files on request

Click **Export All Now** to immediately export all log files.

Click **Roll Now** at the bottom of the Logfile Options pane to immediately roll all log files. This option is generally used for testing and troubleshooting purposes.

Delete exported log files

To delete log files from the firewall after they have been exported, select **Delete logs after export**.

Monitoring disk space using cron jobs

The roll audit cron job serves an important function in monitoring available disk space. There are two `rollaudit` jobs. The first job checks the size of various audit and log files daily at 2:00 a.m. The second job runs each hour and rotates files found to be growing too quickly. When these jobs run, they check the `/secureos/etc/rollaudit.conf` configuration file to see which files should be rotated. The following files are checked by `rollaudit`:

- `/var/log/audit.raw` (The firewall generates reports when these files are rolled.)
- `/var/log/cron`
- `/var/log/daemon.log`
- `/var/log/daemond.log`
- `/var/log/messages`
- `/var/log/maillog` (This file is rotated once a week.)
- `/var/log/SF.log`
- `/var/log/snmpd.log`

You can edit the `/secureos/etc/rollaudit.conf` file to specify how large files are allowed to get before they are rotated and the maximum amount of time that should elapse between rotations. See the `rollaudit` man page for details on editing this file.

Caution: To avoid serious system problems, do not allow the `/var/log` partition to become full. The `/sbin/logcheck` job will generate an e-mail message warning you if the `/var/log` partition becomes 85% full and then again if it becomes 100% full.

Identifying changes using change tickets

You can create change tickets to identify specific changes made to the firewall.

1 Start a change ticket:

- **Manually** – On the Admin Console toolbar, click **Start ticket**. The Change Ticket window appears.

- Automatically – On the Audit Management tab, select **Require change ticket**. *Ticket (Required)* appears at the bottom of the Admin Console.

When you save changes made to the firewall, the Change Ticket window appears.

- 2 In the **Ticket** field, enter a name to identify the ticket.

The ticket name can be 1–32 characters, using letters, numbers, symbols, underscores, and spaces. Quotes, double quotes, and back quotes are not allowed.

Note: If you enter an existing ticket name, the existing ticket is added to and, if enabled, a new automatic configuration backup is not created.

- 3 [Optional] In the **Description** box, enter information to further identify this ticket.
- 4 Click **OK**. The open ticket name appears at the bottom of the Admin Console. All changes made while the ticket is open are associated with the ticket.
- 5 Make the appropriate changes on the firewall and save your changes.
- 6 On the Admin Console toolbar, click **Stop ticket**.

Exporting audit data to Firewall Reporter and syslog servers

The Sidewinder uses the UNIX syslog facility to log messages sent by programs running on the firewall. These messages can be useful in tracking down unauthorized system users or in analyzing hardware or software problems. All syslog data is stored in the audit log files.

Note: Firewall Reporter is a legacy feature and is not actively supported.

Listed below are some basic points about syslog and how it works on the Sidewinder.

- syslog runs as a daemon process called *syslogd*.
- Each application determines whether it will use syslog and the types of messages that will be generated. Normally, applications generate messages of different severity levels, such as informational and critical.
- Hackers will often try to edit syslog files to cover any evidence of their break-ins. The firewall uses Type Enforcement to protect the syslog files from being modified by unauthorized users.
- A copy of the syslog data is sent to the firewall's audit log files.
- The log files generated by *syslogd* can get large and start using a lot of hard disk space. To solve this problem, the log files are periodically rotated. See [Troubleshooting system status](#) for more information on file rotation.

To send audit data from your Sidewinder to a Firewall Reporter or to designated syslog servers: Select **Monitor > Audit Management** and click the **Firewall Reporter/Syslog** tab. The Audit Management tab appears.

Note: Use the Audit Management tab to capture network and system utilization statistics, to automatically create change tickets, and to configure and schedule the export of audit log files. For more information, see [Managing log files](#).

The Firewall Reporter/Syslog tab consists of two panes:

- [Export audit to Firewall Reporter](#)
- [Export audit to syslog servers](#)

Auditing

Exporting audit data to Firewall Reporter and syslog servers

Figure 211 Firewall Reporter/Syslog tab

The screenshot shows the 'Firewall Reporter/Syslog' tab within the 'Audit Management' window. It is divided into two main sections: 'Export audit to McAfee Firewall Reporter' and 'Export audit to syslog servers'.

Export audit to McAfee Firewall Reporter

- ☐ Use Firewall Reporter
- IP address: [DNS Lookup icon]
- Remote facility:
- Filter:

Export audit to syslog servers

Below this section is a table with the following columns: Enabled, IP Address, Remote Facility, and Description. The table is currently empty, with a scroll bar at the bottom.

Export audit to Firewall Reporter

Use this pane to configure your firewall to export audit data to Firewall Reporter. Firewall Reporter provides more advanced reporting capabilities than what is available directly on the Sidewinder.

Note: Firewall Reporter is a legacy feature and is not actively supported.

To configure your firewall to export audit to Firewall Reporter, make the following entries and selections:

- **Use Firewall Reporter** – Select this check box to enable real-time transmission of firewall audit data to Firewall Reporter.
- **IP address** – Enter the IP address of the Firewall Reporter's syslog server. To find the IP address for a host name, type the name and click **DNS Lookup**.
- **Remote facility** – Select a syslog facility to help identify the audit export.
- **Filter** – Select a filter to include or exclude certain types of audit records from your export file. See [About the predefined audit filters](#) for more information.

Export audit to syslog servers

The firewall provides you with the option to convert audit data into various formats used by third-party reporting tools. To generate reports based on the log files, you must format the audit data and then export those files to the workstation or host that contains the software needed to generate log reports (for example, Firewall Reporter). You can then generate the Sidewinder log reports on that machine.

Use this pane to export audit data to a syslog server, to generate and display reports based on the Sidewinder log files.

Use the toolbar to perform the actions described in [Table 39](#).

Figure 212 Export audit to syslog servers toolbar




Table 39 Export audit to syslog servers toolbar tasks

Button	Action
New	Click New to create an export entry.
Advanced	Select an entry and click Advanced to further define the parameters for audit export. See Configuring advanced settings for exporting audit to syslog servers for more information.
Delete	Select the entry you want to delete, then click Delete . Click Yes to confirm the deletion.

To redirect audit output to a syslog sever:

- 1 On the toolbar, click **New**.
- 2 Click in the **IP Address** cell and enter the address of the syslog server you are sending audit data to.
- 3 From the **Remote Facility** drop-down list, select a syslog facility to help identify the audit export.
- 4 If desired, click in the **Description** cell and enter information to further identify the audit export.
- 5 Save your changes.

Configuring advanced settings for exporting audit to syslog servers

Use the Firewall Reporter/Syslog Advanced Settings window to further define the parameters for audit exports.

You can optionally make the following entries and selections:

- **Port** – The default port that the firewall exports audit data through is port 514.
- **Filter** – Select a filter to include or exclude certain types of audit records from your export file.
- **Format** – Select a format to convert the audit data into.
- **Max PDU size** – Enter the maximum size of a syslog record.
- **PDU exceed behavior** – Select how to audit export records that exceed the maximum PDU size:
 - To delete the remainder of the export record past the maximum PDU size, select **Truncate**.
 - To divide the export record into segments that each match the maximum PDU size, select **Fragment**.

12 Service Status

Contents

[Understanding processes that control server status](#)

[Viewing service status](#)

[Viewing a service's process information](#)

Understanding processes that control server status

There are two significant processes involved in controlling and monitoring service status: **daemon** and **NSS**. **daemon** controls the starting, stopping, and restarting of services, and **NSS** handles port assignments.

Learn more about these processes in the following sections:

- [daemon](#)
- [Network Services Sentry \(NSS\)](#)

daemon

If you have administered a standard UNIX system, you are probably familiar with **init**, which manages process control initialization. On Forcepoint Sidewinder, **init** has been augmented with the **daemon** process. **daemon** is a powerful component that enhances overall security. It monitors and controls all of the major software components on the firewall.

The **daemon** process also detects and audits some classes of attacks against the firewall. For example, should someone try to attack a firewall service (such as **sendmail**), causing the component to crash, the **daemon** process detects the failure, immediately restarts the failed component, and creates a critical event audit entry, which allows the administrator to be notified and respond to the attack.

daemon starts during the firewall boot process. On start up, it reads the `/secureos/etc/daemon.conf` file to determine its configuration options. By default, **daemon** runs in its normal mode. This means that **daemon** attempts to start all enabled components in the `/etc/server.conf` and `/secureos/etc/nss.common.conf` files. **daemon** is capable of restarting and stopping processes both automatically and manually. A full description of **daemon**'s usage is available on the **cf_daemon** man page. If **daemon** detects certain failure events, it switches to failure mode. Failure mode is explained in [About failure mode](#).

Restarting processes

If a component dies unexpectedly, **daemon** restarts that component and audits the event in both the audit log and the **daemon** log. The message in `/var/log/daemon.log` is similar to this:

```
Jan 7 12:39:55 2007 EST: Starting 'ftp' (6896): '/usr/libexe/pftp -f
/secureos/etc/proxy/pftp.conf -I 0'
```

If a component quits within five seconds of starting three times in a row, **daemon** does not attempt to restart the component until the next time **daemon** gets a **SIGHUP**. This event will also be audited to both the audit log and the **daemon** log. The message in `/var/log/daemon.log` will look similar to this:

```
Jan 17 13:26:38 2007 EST: ftp (7061) died after restart; not restarting
```

You can manually restart an agent using the Service Status window or **cf daemon restart agent=agentname**.

Note: When you restart an agent, you restart all the processes related to that agent. If you have multiple services using the same agent, all those services are restarted.

Responsive service processes

daemonD monitors some services for responsiveness. If a service does not respond to periodic messages within 15 seconds, daemonD gathers diagnostic information, kills the process, and then restarts it.

These services are monitored:

- NSS
- HTTP proxy
- HTTPS proxy
- TCP proxy
- DNS proxy

The diagnostic information is gathered in a tar archive and stored in `/var/diagnostics`. If three tar archives are gathered for a service, daemonD kills and restarts the process, but stops gathering diagnostics.

You can change the number of seconds that daemonD waits for a response before starting the diagnostic program from the command line using

```
cf daemonD set ping_timeout=seconds
```

For more information, see the diagnostic man page.

Stopping processes

daemonD is also responsible for *stopping* processes. If a firewall administrator chooses to disable a process (using the Admin Console or `cf` commands), the configuration files are changed and a `SIGHUP` command is sent to daemonD. The `SIGHUP` command signals daemonD to reread the configuration files. If daemonD finds an entry associated with a currently running process that is now marked as disabled, daemonD will stop that process. The process will not be started again until it is re-enabled by an administrator. Re-enabling a process will cause another `SIGHUP` command to be sent to daemonD, which will reread the configuration files and attempt to restart the process.

All component failure events are logged in the `/var/log/daemonD.log` file and the audit log. If daemonD fails during system startup, the daemonD log file will record the reason for this failure. It will also record information each time daemonD restarts a process that died unexpectedly. This is useful for tracking attacks on a particular component.

About failure mode

When a failure event occurs, daemonD will start in failure mode. This mode is also called *safe mode*. This means that daemonD only starts those components necessary to administer the system. Components that are not enabled for failure mode will not be started, which includes most proxy agents.

Failure mode is set under any of the following circumstances:

- a license check fails
- the audit partition overflows

Once the problem that sent the firewall into failure mode has been corrected, use `cf daemonD set failure_mode=off` to resume normal operation.

About High Availability and daemonD

If you configure a failover High Availability (HA) cluster, the standby firewall will run in standby mode with a limited set of services. If the primary becomes unavailable and the standby is required to take over as the primary, daemonD will start all services for that firewall.

If the primary in an HA cluster goes into failure mode and the secondary/standby is not available, the primary will remain as the primary, but the priority value for that firewall will change to one, ensuring that if a secondary/standby becomes available, it can take over as the primary. For information on HA, see [Chapter 23, High Availability](#).

Network Services Sentry (NSS)

If you have administered a standard UNIX system, you are probably familiar with `inetd`, which listens for connections and manages daemons for network services. Daemons are server processes that run continuously in the background and wait until they are needed. On the Sidewinder, `inetd` has been replaced with the Network Services Sentry (NSS). There is an NSS configuration file for each burb defined on your firewall. NSS regulates the ability to change the default port. For example, the files are updated whenever you change a service's ports.

You may use the Admin Console or the command line to edit a service's default ports. The NSS configuration files are updated for you when you make these changes. For example, you might want to alter ports when the default conflicts with the port of another service, or when you want to create a portlist with non-continuous numbers. When changing the port for a service, be sure to consider the criteria listed in [Table 40](#) below.

Table 40 Criteria for modifying a service port

Port type	Criteria
Port	<ul style="list-style-type: none">Valid port values are between 1–65535.Must be unique within ports assigned to other enabled services of the same type
Port range	<ul style="list-style-type: none">Must be two valid ports separated by a single hyphenMust be listed in ascending orderThe range must have a maximum of 1995 ports. If a service requires more than 1995 ports, use a portlist.
Portlist	<ul style="list-style-type: none">May be non-continuousValid ports and/or valid ranges separated by commas


Viewing service status

Knowing a service's status is an important part of monitoring your Sidewinder. It can help you verify that a service is configured correctly, and it can help you determine if the service is running as expected.

The Service Status window allows you to view configuration and status information on all services that are enabled on your firewall. It has shortcuts to audit and usage information so you can easily gather information about individual services. You can also restart a service from this window, which is sometimes required after certain configuration changes or as a troubleshooting step.

To view the services that are currently used in enabled rules, select **Monitor > Service Status**. The main Service Status window appears.

Figure 213 The main Service Status window



Status	Service	Burbs	Ports	Active Rules
	Admin Console	external, internal	9003/tcp	Admin Console
	dns	external, internal	53/tcp, 53/udp	dnsp all to external resolv
	http	external	80/tcp	Internet Services, inbound
	nss	Firewall, external, internal		
	sshd	external, internal	22/tcp	Secure Shell Server
	ssod	external, internal	8111/tcp	Passport

You can accomplish the tasks listed in the following table using the toolbar shown here.

Figure 214 Tasks available in the Service Status window

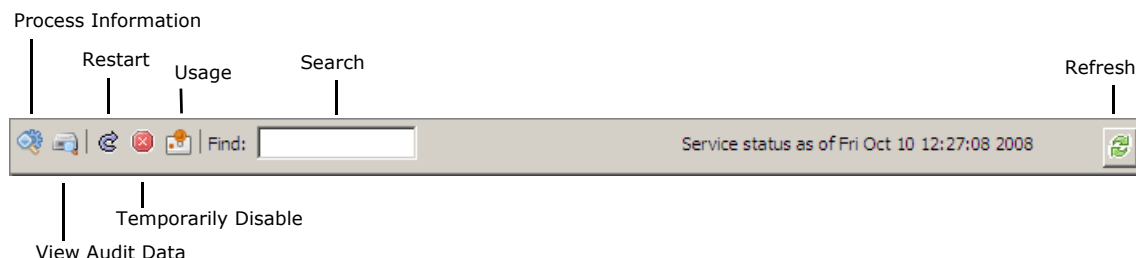












Table 41 Tasks that can be performed from the main Service Status window

Icon/ Menu item	Tasks
Process Information	View the service's status (running as expected, running with errors, not running), its ports, burbs, and where it is listening by selecting a service and clicking Process Information . Opening this window is most useful for checking that a service is listening on the expected ports or for monitoring the status of a single service. For more information, see Viewing a service's process information .
View Audit Data	View a service's audit data by selecting a service and clicking View Audit Data . This displays the past 24 hours of data. Additional audit viewing is available at Monitor > Audit .
Restart	Restart services by selecting one or more services and then clicking Restart . Clicking Restart also re-enables a disabled service. In this case, the firewall first checks the policy to verify that the service should be enabled. Caution: Before you restart a service, make sure you know which agent the service is using. A restart disables and enables the underlying agent, which means all connections using this agent will be dropped as opposed to just dropping the connections using this service.
Temporarily Disable	Temporarily disable services by selecting one or more services and then clicking Temporarily Disable . Tip: A quick way to safely re-enable all stopped agents is to change a rule or service's description and save the changes.
Usage	View the rules that currently use a given service by selecting a service and then clicking Usage .
Find	Find a service by entering a character string related to the service you are searching for in the Find field. The search function searches all columns, and filters as you type. For example, if you are searching all services running in the DMZ burb, typing "DMZ" reduces the list to only the services containing that character string. Clear the Find field to show all options again.
Refresh	View current information for all services by clicking Refresh .

This window displays the following information about each service:

Table 42 Service Status window service information

Field	Information Provided									
Status	<p>Indicates if the service is running as expected</p> <table><tr><td></td><td>Running</td><td>The service is processing traffic as expected.</td></tr><tr><td></td><td>Running with errors</td><td>The service is processing traffic but it is generating errors and needs to be investigated, or is temporarily disabled.</td></tr><tr><td></td><td>Not running</td><td>The service is not running, or no information is available about the service's status. The service needs to be investigated.</td></tr></table>		Running	The service is processing traffic as expected.		Running with errors	The service is processing traffic but it is generating errors and needs to be investigated, or is temporarily disabled.		Not running	The service is not running, or no information is available about the service's status. The service needs to be investigated.
	Running	The service is processing traffic as expected.								
	Running with errors	The service is processing traffic but it is generating errors and needs to be investigated, or is temporarily disabled.								
	Not running	The service is not running, or no information is available about the service's status. The service needs to be investigated.								
Service	<p>The service's name</p> <p>Note: kvmfilter and virus-scan appear when these options are selected on an enabled rule's application defense. They are not associated with an agent. See the related service, such as sendmail, for complete burb and port information.</p>									
Burbs	<p>The burbs where a service is enabled</p> <p>When a service is used in a rule, the service is enabled in that rule's source burb. All source burbs for rules that use this service are listed here. The  icon indicates that the service is enabled in all burbs valid for that service.</p> <p>Note: Certain services display the Firewall burb. This burb is used for firewall internal processing and cannot be modified. Sendmail only runs in two burbs, even if the source burb is set to <Any>.</p>									
Ports	The ports configured for the service									
Active Rules	The enabled rules that use this service									


Note: If a warning message states that the firewall is in failure mode, you must take action to restore the firewall to a normal operating state. See [About failure mode](#) for information.

Viewing a service's process information

This section provides information on the Service Status Process Information window. You can access this window by selecting **Monitor > Service Status** and then double-clicking a service, or selecting a service and then clicking **Process Information** on the toolbar.

The Service Status Process Information window appears.

Figure 215 The Service Status: Process Information window



Use the Process Information window to view the burbs and ports on which the service should be listening, as well as the service's current status.

From this window, you can do the following:

- Refresh the data – Click **Refresh** to display current status information.
- Check a service's status – Status is displayed near the top of the window. Possible statuses are:

Table 43 Service status options

Icon	Service Status	Status Description
	Running	The service is processing traffic as expected.
	Running with errors	The service is processing traffic but it is generating errors and needs to be investigated, or is temporarily disabled.
	Not running	The service is not running, or no information is available about the service's status. The service needs to be investigated.

Service Status

Viewing a service's process information





- View where the service is running and listening for connections – Configuration areas are:

Table 44 Service Status Process Information window configuration areas

Configuration area	Description
Configured Burb	When a service is used in a rule, the service is enabled in that rule's source burb. All source burbs for rules that use this service are listed here.
Configured Port	All ports that are configured for this service are listed here.
Listening	When a service is listening (accepting connections) on a port, a green check mark appears in this column. Note: If a port does not have a check mark next to it, there is a problem with the service that needs to be investigated. Contact Technical Support for assistance.

Note: The kvmfilter and virus-scan services appear when these options are selected on an enabled rule's application defense. They are not associated with an agent. See the related service for complete burb and port information.

Table 45 Service status Process Information window buttons

Icon	Action	Description
	Restart	Click this button to restart or re-enable the agent used by this service. Restarting a service disables and then immediately enables the service's agent. This action drops all current connections and resets any audit counts (for example, if an IPS attack response is checking the frequency of an attack before issuing an alert). Do not restart an agent unless it is part of a procedure, you have completed other troubleshooting measures, or have been instructed to by Forcepoint support. Note: Restarting a service drops all current connections for that agent, not just the selected service.
	Temporarily Disable	Click this button to halt the agent used by this service. Temporarily disabling a service stops the service's agent. The agent is restarted as soon as any policy configuration changes are saved. Do not temporarily disable an agent unless it is part of a procedure, you have completed other troubleshooting measures, or have been instructed to by Forcepoint support. Tip: A quick way to safely re-enable all stopped agents is to change a rule or service's description and save the changes.
	Refresh	Click this button to view the most current information.
		Click this button to view audit output that is filtered to show this service's activity over the past 24 hours.

Service Status

Viewing a service's process information

13 IPS Attack and System Event Responses

Contents

[Understanding attack and system event responses](#)

[Creating IPS attack responses](#)

[Creating system responses](#)

[Ignoring network probe attempts](#)

[Sidewinder SNMP traps](#)

Understanding attack and system event responses

Forcepoint Sidewinder IPS attack responses and system event responses allow you to monitor your network for abnormal and potentially threatening activities ranging from an attempted attack to an audit overflow. Using the Admin Console, you can configure how many times a particular event must occur within a specified time frame before it triggers a response.

When the Sidewinder encounters audit activity that matches the specified type and frequency criteria, the response you configured for that system event or attack type determines how the firewall will react. The firewall can be configured to respond by alerting an administrator of the event via e-mail and/or SNMP trap and by ignoring packets from particular hosts for a specified period of time (known as a *Strikeback™*).


Some default attack and system event responses are automatically created on the firewall during its initial configuration. The additional configuration options you select will depend mainly on your site's security policy and, to some extent, on your own experiences using the features. You may want to start with the default options and make adjustments as necessary to meet your site's needs.

Creating IPS attack responses

IPS (Intrusion Prevention System) attack responses allow you to configure how the firewall responds when it detects audit events that indicate a possible attack, such as Type Enforcement violations and proxy floods. When you create a new response, the Add IPS Attack Response Wizard guides you through the options. You can modify these options at any time from the IPS Attack Responses main window.

To launch the wizard, view or configure attack responses, or change who should receive attack alerts, select **Monitor > IPS Attack Responses**. The following window appears:

Figure 216 IPS Attack Responses main window



Name	Attack Type	Frequency	Response
ACL Deny	ACL Deny	5 in 30 seconds	E-mail
Denied Authentication	Denied Authentication	20 in 60 seconds	E-mail
IPFilter Deny	IPFilter Deny	5 in 30 seconds	E-mail
IPS	Attack Severe	Every time	E-mail
Keyword Filter Failure	Keyword Filter Failure	5 in 30 seconds	E-mail
Network Probe	Network Probe	50 in 30 seconds	E-mail
Proxy Flood	Proxy Flood	5 in 10 seconds	E-mail, Blackhole
Signature IPS Intrusion All	Signature IPS Intrusion All	Every time	E-mail
Signature IPS Intrusion t...	Signature IPS Intrusion Blackholed	Every time	E-mail, Blackhole
Signature IPS Intrusion L...	Signature IPS Intrusion Deny	Every time	E-mail
Spam Filter Failure	Spam Severe	5 in 30 seconds	E-mail
TCP SYN Attack	TCP SYN Attack	Every time	E-mail
TrustedSource	TrustedSource	Every time	E-mail, Blackhole
Type Enforcement	Type Enforcement	Every time	E-mail
Virus Filter Failure	Virus Severe	5 in 30 seconds	E-mail

New Modify Delete Disable Response Settings

Use this window to perform the following tasks:

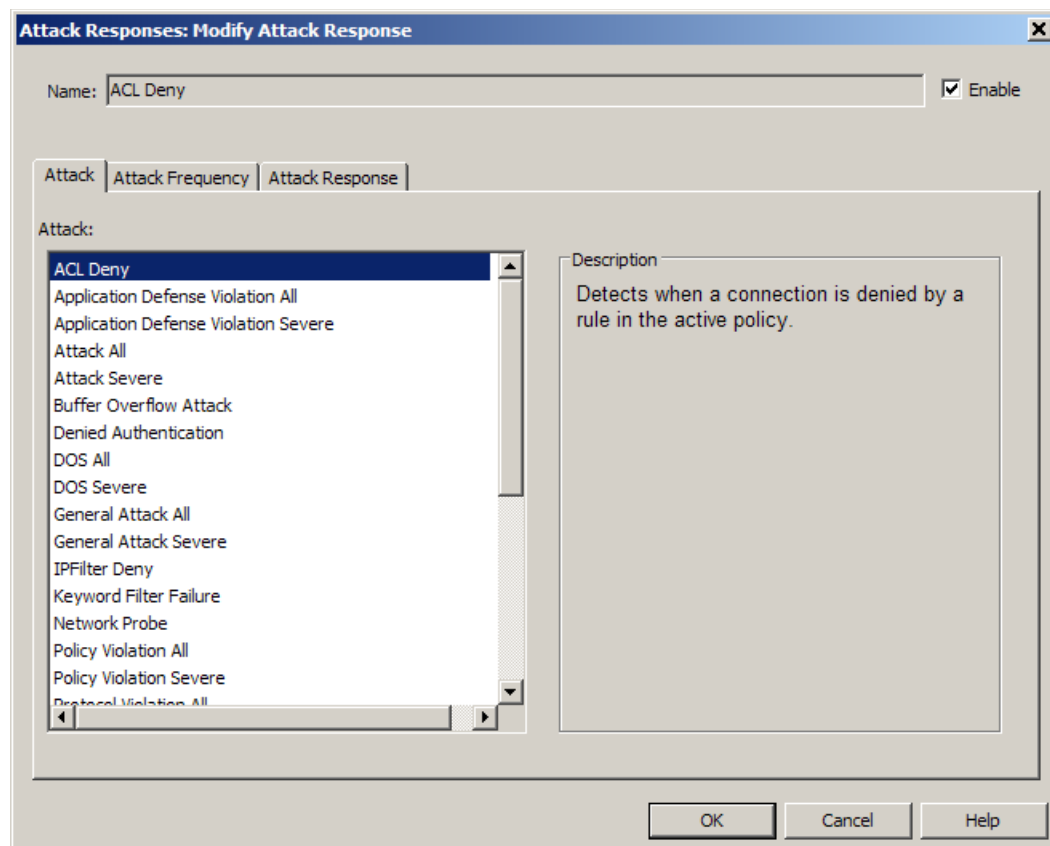
- **Configure a new IPS attack response** – To configure a new IPS attack response, click **New**. The Add Attack Response Wizard appears. Follow the on-screen instructions.
- **Modify an existing IPS attack response** – To modify an existing IPS attack response, select the appropriate item within the list and click **Modify**. (Read-only administrators can click **View** to view an IPS attack response.)
See [Modifying an IPS attack response](#) for more information.
- **Filter the list of IPS attack responses** – To modify the displayed list, right-click a column name and select from the current list of filters or create a custom filter. The list then displays only IPS attack responses of that type.
- **Delete an existing IPS attack response** – To delete an IPS attack response, select the list item you want to delete and then click **Delete**.
- **Disable/enable an IPS attack response** – The disable and enable options depend on an IPS attack response's current status. If one or more responses with the same status are selected, their status can be changed to its opposite (for example, if all selected responses are enabled, you may disable all of them). When multiple responses with mixed statuses are selected, the only available action is enabling the responses.

- **Create the e-mail list to notify in the event of an attack** – To create or modify the list of e-mail addresses to notify if any IPS attack triggers an alert, click **Response Settings**. You can also blackhole a source IP address if the attack IP cannot be confirmed. See [Configuring the e-mail response settings](#) for more information.

Modifying an IPS attack response

When you modify an IPS attack response, the following window appears:

Figure 217 IPS Attack Responses: Modify window



About the Modify Attack Responses: Attack tab

Use this tab to change this attack response's attack filter. An attack is generally defined as suspect traffic at either the network or application level. Each attack filter identifies a different attack audit event.

- 1 Select the attack for which you want the firewall to send out a response. A complete list of pre-defined attacks is provided in the following table.

To create additional attack filters, see [Create a custom filter](#).

- 2 Click **OK** or the next tab you want to modify.

Note: For descriptions of the audit severities, see [Understanding audit event severities](#).

Table 46 Descriptions of pre-defined attacks

Attack	Description
ACL Deny	Detects when a connection is denied by a rule in the active policy.
Application Defense Violation All	Detects attacks of all severities that violate active policy defined by Application Defenses. This attack category includes mime and keyword filter failure attacks.
Application Defense Violation Severe	Detects when severe attacks violate active policy defined by Application Defenses, including mime and keyword filter reject audits.
Attack All	Detects attack events of Application Defense violation attacks, buffer overflow attacks, DOS attacks, general attacks, policy violation attacks, protocol violation attacks, virus attacks, and spam attacks.
Attack Severe	Detects severe attacks. This detects Application Defense violation attacks, buffer overflow attacks, general attacks, DOS attacks, policy violation attacks, protocol violation attacks, virus attacks, and spam attacks. Severe attacks indicate something is occurring that an administrator should know.
Buffer Overflow Attack	Detects attempted buffer overflow attacks targeted at protected systems.
Denied Authentication	Detects when a user attempts to authenticate and enters invalid data. For example, if a user is required to enter a password and entered it incorrectly, the denied auth event would log the event.
DOS All	Detects Denial of Service attacks of all severities. This attack category also detects all severities of TCP SYN attacks and proxy flood attacks.
DOS Severe	Detects severe Denial of Service attacks. This attack category also detects TCP SYN attacks and proxy flood attacks. Severe attacks indicate something is occurring that an administrator should know.
General Attack All	Detects general attacks of all severities that do not fall into the pre-defined categories.
General Attack Severe	Detects severe general attacks that do not fall into the pre-defined categories. Severe attacks indicate something is occurring that an administrator should know.
IPFilter Deny	Detects when a connection is denied by the active IPFilter policy.
Keyword Filter Failure	Detects when an SMTP mail message is rejected due to a configured keyword filter.
Network Probe	Detects network probe attacks, which occur any time a user attempts to connect or send a message to a TCP or UDP port which has no service. Note: The firewall does not blackhole netprobe attacks, as they are likely to be denial of service attacks from spoofed source addresses.
Policy Violation All	Detects attacks of all severities that violate the active policy. This attack category also detects all severities of failed authentication attacks, ACL and IPFilter deny attacks, and Type Enforcement error attacks.
Policy Violation Severe	Detects severe attacks that violate the active policy. This attack category also detects failed authentication attacks, network probe attacks, ACL and IPFilter deny attacks, and Type Enforcement error attacks. Severe attacks indicate something is occurring that an administrator should know.
Protocol Violation All	Detects attacks of all severities that violate protocol compliance.
Protocol Violation Severe	Detects severe attacks that violate proxy protocols (HTTP, Telnet, FTP, etc.). Severe attacks indicate something is occurring that an administrator should know.
Proxy Flood	Detects potential connection attack attempts. A connection attack is defined as one or more addresses launching numerous proxy connection attempts to try and flood the system. When NSS receives more connection attempts than it can handle for a proxy, new connections to that proxy are briefly delayed (to allow the proxy to "catch up"), and the attack is audited.
Signature IPS Intrusion All	Detects all attacks identified by the signature-based IPS. This category detects attacks that were denied, dropped, or rejected, as well as suspected attacks that were allowed but were audited by IPS.
Signature IPS Intrusion Blackholed	Detects attacks identified by the signature-based IPS where the attacker was blackholed.

Table 46 Descriptions of pre-defined attacks <Comment>(continued)

Attack	Description
Signature IPS Intrusion Deny	Detects attacks identified by the signature-based IPS where the offending network session was dropped, or rejected, or the attacker was blackholed.
Spam	Detects attacks of all severities that are spam.
Spam Severe	Detects severe attacks that are spam.
TCP SYN Attack	Detects a possible attempt to overrun the firewall with connection attempts.
TrustedSource	Detects attacks identified as spam by TrustedSource.
Type Enforcement	Detects when there is a TE violation due to an unauthorized user or process attempting to perform an illegal operation.
Virus	Detects attacks of all severities that are viruses.
Virus Severe	Detects severe attacks that are viruses.

About the Modify Attack Response: Attack Frequency tab

Use this tab to modify the parameters to be met before the firewall generates a response. The options are:

- **Always respond** – Select this option to have the firewall respond each time the attack specified on the Attack tab occurs.
- **Limit responses** – Select this option to respond only when the attack pattern matches the parameters set here:
 - **Respond if x attacks in y seconds** where:
 - Valid values for x are between 2 and 100000. The firewall responds when the x attack occurs.
 - Valid values for y are between 1 and 100000. This represents a buffer of y seconds, so the firewall checks the current time - y.

For example, if you have configured a response to filter for netprobe attempts, and you want to trigger an attack response if 5 or more probe attempts occur within a 30-second period, you would enter “Respond if **5** attacks in **30** seconds.”
 - **Reset attack count to zero after responding** – After x attacks, the firewall zeroes out its attack counter and waits until another x attacks occur in y seconds before sending out the next e-mail alert or SNMP trap.

If this option is not selected, the same attacks may be used to generate additional alerts.

About the Modify Attack Response: Attack Response tab

Use this tab to configure how the firewall should respond when the attack type's pattern matches the criteria on the Attack Frequency tab. The options are:

- **Configure an alert** – The firewall can send an alert using an e-mail, an SNMP trap, or both.
 - **Send e-mail to:** Select this option and select a group from the drop-down list to send an e-mail to each address in the selected group.

You can create different groups to receive e-mails for different types of attacks. Create groups of e-mail addresses from the main IPS Attack Response window. Additional information is available in [Configuring the e-mail response settings](#).

- **Send SNMP trap:** Select this option to send an SNMP trap to the location(s) configured for the snmpd server. (Configure the SNMP server at **Policy > Rule Elements > Services > snmpd**. Additional information is available in [About Sidewinder SNMP traps](#).)
- [Conditional] If configuring an alert, specify how long the firewall should wait before sending the next e-mail or SNMP trap for the same attack type by using the **Time to wait between alerts (seconds)** option.

For example, suppose you configure an alert to trigger when 5 or more denied authentication attempts occur in a 30-second period, and you instruct the firewall to wait 300 seconds (five minutes) between alerts.

In this configuration, if an intruder attempts to authenticate 5 times in a 30 second period, a response is triggered. However, if the intruder tries 5 more authentication attempts during the next 30 seconds, the firewall will not send another alert. Note that if the response calls for a Strikeback (see next section), traffic will continue to be blackholed.

After five minutes, if the threshold is again reached, another alert will be triggered.

- **Configure Strikeback** – The firewall can *blackhole*, or ignore, traffic from a host that is sending suspect traffic.
Caution: The firewall blackholes based on source address, as opposed to traffic type. If you choose to blackhole a host, **all** traffic from that host will be ignored.

Blackhole – Select this option to ignore all traffic from the suspect traffic's source(s) for a set time period. The source of the attack is recorded in the audit event's `attack_ip` field. The source of the suspect traffic may be the connection's source IP address (a peer or a client) or destination IP address (if a server is attacking a client). If the firewall considers it likely that the source IP address could have been forged, it will leave the `attack_ip` field blank and not blackhole any IP address for this audit event. The apparent source and destination IP address is still recorded in the audit event.

If you select the Blackhole option, you must also specify for how long you want to blackhole traffic. Set a time limit in the **Blackhole packets for x seconds** field, where x is a value between 1 and 100000.

Tip: If you find you need to blackhole traffic for more than 100,000 seconds (a little over 24 hours), consider creating a TCP/UDP Packet Filter deny rule for that traffic.

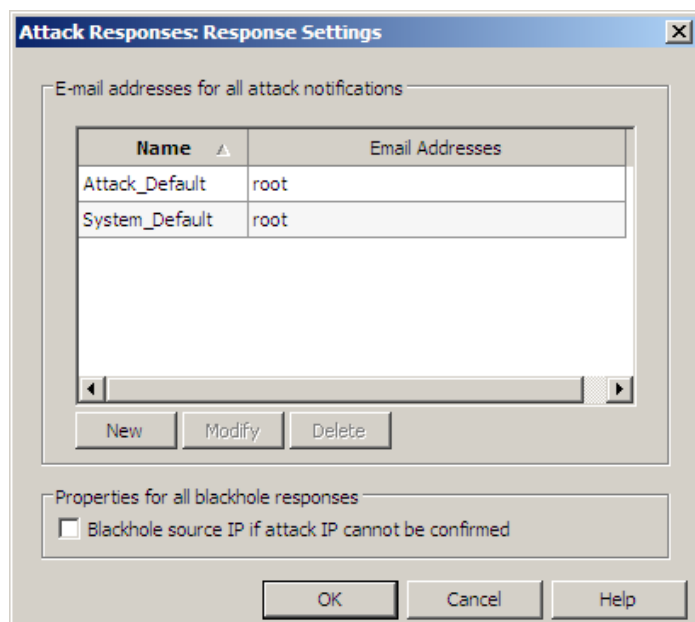
- **All attacking hosts:** Select this option to blackhole all hosts involved in triggering the alert. For example, if you want an alert after 5 occurrences in 30 seconds and host A sent 4 occurrences and host B sent 1, all traffic from hosts A and B would be ignored for the set amount of time.
- **Each host responsible for y% of the attacks:** Select this option to limit blackholing on a percentage basis. For example, if you set the percentage at 50% and host A caused 4 out of 5 attacks and host B caused 1 out of 5 attacks, only traffic from host A would be ignored.

Use the Dashboard's Blackholed IP window to view, delete, and manually add blackholed IP addresses.

Configuring the e-mail response settings

To view, add, modify, or delete the e-mail addresses that will receive alerts, click **Response Settings** in the IPS Attack Responses main window's lower-right corner. The following window appears:

Figure 218 Attack Responses: Settings window



Use this window to configure groups of e-mail addresses that will receive alerts. The groups you create here can be selected in the Attack Response tab. For every triggered attack response that is set to send an e-mail alert, the selected group of e-mail addresses will receive an alert.

You can configure entries by using the buttons described here:

Tip: If you have not already done so, create an off-box alias for the root and administrators mail accounts. This ensures that system messages are sent to an account that is checked regularly. If mail is not forwarded or checked regularly, the local mailbox could fill up too much hard disk space and cause problems. See [Setting up e-mail aliases for administrator accounts](#).

- **New** – Click this button to define a new group of e-mail addresses to receive attack alerts.
- **Modify** – Select an entry and click this button to modify an existing group of e-mail addresses.
- **Delete** – Select an entry and click this button to delete that group of e-mail addresses.
- **Blackhole source IP if attack IP cannot be confirmed** [Attack Responses only] – Select this check box to blackhole a source IP when the related audit message does not have an Attack IP field. No connections will be accepted from the IP address originating the attack.
 - This can be used to enforce thresholds on otherwise allowed behaviors (for example, limiting a connection rate for SSH traffic).
 - This feature can also be used to configure blackholing on netprobes, UDP attacks, and SYN attacks (all audit messages that do not contain an Attack IP field).


Caution: For netprobes, UDP attacks, and SYN attacks, it is possible for the attacker to forge the source IP address. A configuration which blackholes source addresses found in these audits may allow an attacker to trigger a blackhole for an unrelated third party, potentially interrupting desired traffic.

Creating system responses

System responses allow you to configure how the firewall responds when it detects audit events that indicate significant system events, such as license failures and log overflow issues.

To view or configure system responses, select **Monitor > System Responses**. The System Responses window appears.

Figure 219 System Responses main window



Name	Event Type	Frequency	Response
HA Failover	HA Failover	Every time	E-mail
Hardware Software Fail...	Hardware Software Failure	Every time	E-mail, SNMP
Host License Exceeded	Host License Exceeded	Every time	E-mail
IPsec Error	IPsec Error	10 in 60 seconds	E-mail
License Expiration	License Expiration	Every time	E-mail
Log Overflow	Log Overflow	Every time	E-mail
Network Traffic	Network Traffic	1000 in 30 seconds	E-mail
Power Failure	Power Failure	Every time	E-mail
Profiler Update Failure	Profiler Update Failure	Every time	E-mail, SNMP
UPS System Shutdown	UPS System Shutdown	Every time	E-mail

New Modify Delete Disable Response Settings

Use this window perform the following tasks:

- **Filter the list of system responses** – To modify the displayed list, right-click a column name and select from the current list of filters or create a custom filter. The list will then display only that system responses of that type.
- **Configure a new system event response** – To configure a new system response, click **New**. The Add System Response Wizard appears.
- **Modify an existing system response** – To modify an existing system response, select the appropriate item within the list and click **Modify**. (Read-only administrators can click **View** to view a system response.)

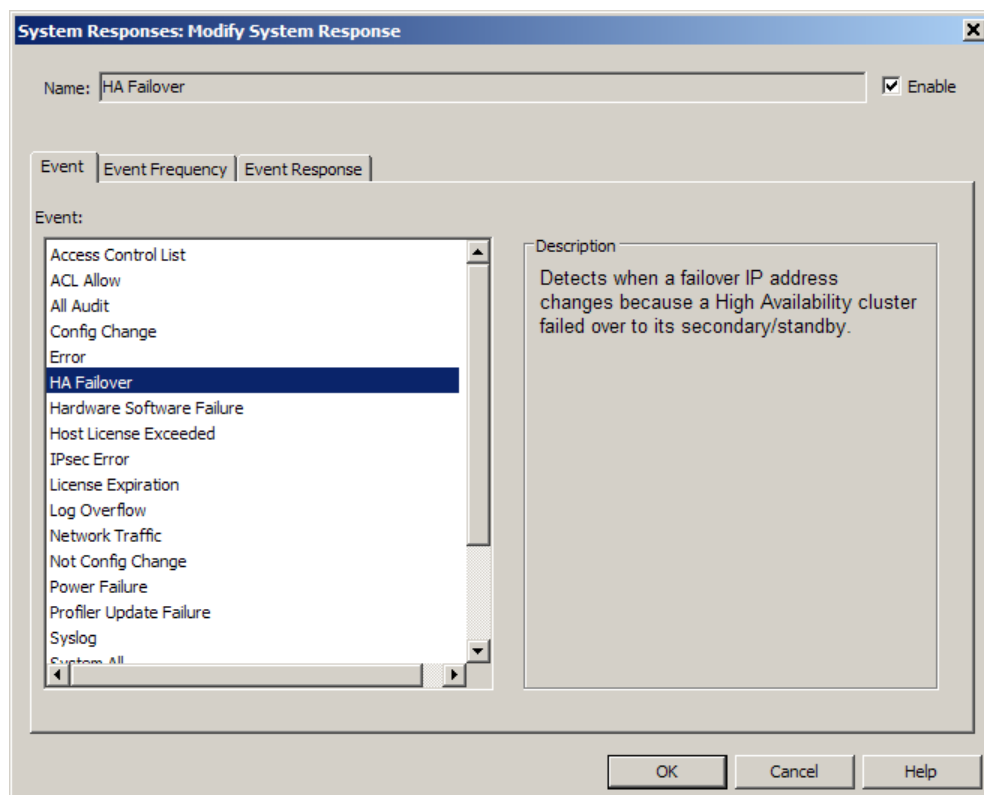
For more information, see [Modifying a system response](#).

- **Delete an existing system response** – To delete a system response, select the list item you want to delete and then click **Delete**.
- **Disable/enable a system response** – The disable and enable options depend on a system response's current status. If one or more responses with the same status are selected, their status can be changed to its opposite (for example, if all selected responses are enabled, you may disable all of them). When multiple responses with mixed statuses are selected, the only available action is enabling the responses.
- **Create the e-mail list to notify in the event of a system event** – To create or modify the list of e-mail addresses to notify if any system event triggers an alert, click **Response Settings**. See [Configuring the e-mail settings](#) for more information.

Modifying a system response

When you modify a system response, the following window appears:

Figure 220 System Responses Modify window



About the Modify System Responses: Event tab

Use this tab to change this system response's event type. An event is generally defined as an important, generally unexpected, change in your system. Each event type identifies a different set of system changes.

- 1 Select the event for which you want the firewall to send out a response. A complete list of pre-defined system events is provided in the following table.

To create additional system event types, see [Create a custom filter](#).

- 2 Click **OK** or the next tab you want to modify.

Note: For descriptions of the audit severities, see [Understanding audit event severities](#).

Table 47 Description of pre-defined system events

Event	Description
Access Control List	Detects all ACL audit events.
ACL Allow	Detects when a connection is allowed by a rule in the active policy.
All Audit	Detects all attack and system events, regardless of characteristics.
Config Change	Detects when the firewall's configuration changes.
Error	Detects all system events identified as AUDIT_T_ERROR in the audit stream.
HA Failover	Detects when a failover IP address changes because a High Availability cluster failed over to its secondary/standby.
Hardware Software Failure	Detects when a hardware or software component fails.
Host License Exceeded	Detects when the number of hosts protected by the firewall exceeds the number of licensed hosts.
IPsec Error	Detects when traffic generates IPsec errors.
License Expiration	Detects when a licensed feature is about to expire.
Log Overflow	Detects when the log partition is close to filling up.
Network Traffic	Detects all connections that successfully pass through the firewall.
Not Config Change	Detects all attack and system events that are not configuration changes.
Power Failure	Detects when an Uninterruptible Power Supply (UPS) device detects a power failure and the system is running on UPS battery power.
Syslog	Detects all audit attacks and system events created via syslog.
System All	Detects all system events of all severities, including power failures, hardware and software failures, failover events, license expiration, host license exceeded, log overflows, and IPsec errors.
System Critical	Detects all critical system events, including power failures, hardware failures, critical software failures, and failover events. Critical system events indicate that a component or subsystem stopped working, that the system is going down (expectedly or unexpectedly), or that the system is not expected to work again without intervention.
System Critical And Severe	Detects critical and severe system events including power failures, hardware failures, critical and severe software failures, failover events, license expiration, log overflows, and IPsec errors. Critical system events indicate a component or subsystem stopped working, that the system is going down (expectedly or unexpectedly) or that the system is not expected to work again without intervention. Severe attacks indicate something is occurring that an administrator should know.
UPS System Shutdown	Detects when a UPS is running out of battery power or has been on battery power for the estimated battery time.
VPN	Detects VPN audit events.

About the Modify System Responses: Event Frequency tab

Use this tab to modify the parameters to be met before the firewall generates a response. The options are:

- **Always respond** – Select this option to have the firewall respond each time the event type specified on the Event tab occurs.
- **Limit responses** – Select this option to respond only when the event's pattern matches the parameters set here:
 - **Respond if x events in y seconds** where:
 - valid values for x are between 2 and 100000. The firewall responds when the x event occurs.
 - valid values for y are between 1 and 100000. This represents the last y seconds, so the firewall checks the current time - y.
 - **Reset event count to zero after responding** – After x events, the firewall zeroes out its event counter and waits until another x events occur in y seconds. If this option is not selected, each subsequent system event that occurs in y seconds will generate a response.

For example, if you want to respond to 5 events in 30 seconds, the firewall constantly checks the past 30 seconds. When the firewall receives 5 system events in that time frame, it responds according to the Event Response tab settings. If it zeroes out after responding, it waits until 5 more events occur in a 30 second time period before responding again.

About the Modify System Response: Event Response tab

Use this tab to configure how the firewall should respond when the event matches the parameters on the Event Frequency tab. The firewall can send an alert using an e-mail, an SNMP trap, or both. The options are:

- **Configure an alert.** The firewall can send an alert using an e-mail, an SNMP trap, or both.
 - **Send e-mail to:** Select this option and select a group from the drop-down list to send an e-mail to each address in the selected group.

You can create different groups to receive e-mails for different types of events. Create groups of e-mail addresses from the main System Response window. Additional information is available in [Configuring the e-mail settings](#).
 - **Send SNMP trap:** Select this option to send an SNMP trap to the location(s) configured for the snmpd server. (Configure the SNMP server at **Services Configuration > Servers > snmpd**. Additional information is available in [About Sidewinder SNMP traps](#))
- [Conditional] If configuring an alert, specify how long the firewall should wait before sending the next e-mail or SNMP trap for the same system event by using the **Time to wait between alerts (seconds)** option. Valid values are between 0 and 65535.

For example, suppose you configure an alert to trigger when 10 or more IPsec errors occur in a 60 second period, and you instruct the firewall to wait 300 seconds (five minutes) between alerts.


In this configuration, if the firewall detects 10 errors in a 60 second period, a response is triggered. However, if it detects 5 more IPsec errors during the next 30 seconds, the firewall will not send another alert.

After five minutes, if the threshold is again reached, another alert will be triggered.

Configuring the e-mail settings

To view, add, modify, or delete the e-mail addresses that will receive alerts, click **Response Settings**, in the System Responses main window's lower right corner. The following window appears:

Figure 221 System Responses: Response Settings window



Use this window to configure groups of e-mail addresses that will receive alerts. The groups you create here can be selected in the Event Response tab. For every triggered system event response that is set to send an e-mail alert, the selected group of e-mail addresses will receive an alert.

You can configure entries by using the buttons described here:

- **New** – Click this button to define a new group of e-mail addresses to receive system event alerts.
- **Modify** – Select an entry and click this button to modify an existing group of e-mail addresses.
- **Delete** – Select a group and click this button to delete that group of e-mail addresses.

Ignoring network probe attempts

If a host on the network attempts to connect to the firewall for a service that is not running, an audit record is generated and may trigger an alarm. An ignore list can be set up to ignore unimportant network probe audit events, but save the audit to keep track of the probe attempts. However, if connection attempts are frequent and are coming from a trusted network, then it may be desirable to ignore them completely and not audit the connection attempt by configuring the appropriate filter rules.

To ignore network probes (commonly referred to *netprobes*), you can create filter rules to deny connection requests for specific ports. For example, if you have problems with NetBios generating netprobes on the firewall, you can discard them and prevent audit events by creating a packet filter service and rule with the following key values:

- For the service, set the Agent field to **TCP/UDP Packet Filter** and set the UDP ports field to **137**. See the following figure.

Figure 222 Example of how to configure a service that can be used to deny NetBIOS netprobes

The screenshot shows the 'Service properties' dialog for the 'TCP/UDP Packet Filter' agent. The 'UDP ports' field is set to '137'. Under 'Stateful packet inspection', the following options are checked: 'Enable stateful packet inspection', 'Enable stateful session failover', and 'Reset TCP connections after connection timeout'. The 'TCP connection timeout' is set to 15 seconds, 'TCP idle timeout' to 7200 seconds, and 'UDP idle timeout' to 300 seconds. Under 'Advanced Options', 'Require UDP checksums' and 'Restrict source port' are unchecked. The 'Bi-directional' checkbox is also unchecked.

- For the rule, set the Action to **Deny** or **Drop**, set Audit to **Errors only (least)**, set the source and destination burbs to **internal**, and the endpoints to **<Any>**. See the following figure.

Figure 223 Example of how to configure a rule that can be used to deny NetBIOS netprobes

The screenshot shows the 'Rule' configuration dialog. In the 'General' tab, the 'Action' is set to 'Drop', the 'Service' is 'netprobe netBIOS (TCP/UDP Packet Filter)', and the 'Audit' is set to 'Errors only (least)'. In the 'Effective Times' section, 'Time period' is set to '<Any>'. In the 'Source' section, 'Burb' is 'internal' and 'Endpoint' is '<Any>'. In the 'Destination' section, 'Burb' is 'internal' and 'Endpoint' is '<Any>'. The 'NAT' is set to 'localhost (Host)' and 'Preserve source port' is unchecked.

Sidewinder SNMP traps

An SNMP trap is an alert message (also known as an alarm message) that is sent as an unsolicited transmission of information from a managed node (router, Sidewinder, etc.) to a management station. The firewall gives you the option of sending audit alert SNMP traps when an audit event, such as an IPS attack event or a system event, triggers a response. Pre-defined (default) alert events are shown in [Table 48 on page 360](#). You also have the option to create custom traps; refer to [Table 48](#).

- For instructions on creating a custom trap, see the `snmptrap` man page.
- To configure the firewall to send the following pre-defined traps, refer to [About the Modify Attack Response: Attack Response tab](#) and [About the Modify System Response: Event Response tab](#).

These traps can also be used in customized audit filters. See [Create a custom filter](#) and [Create a custom filter](#) for more information.

Table 48 SNMP traps

Number	Trap
Default traps	
201	NETWORK_TRAFFIC – This trap is sent when the number of traffic audit events written by the various proxies (WWW, Telnet, FTP, etc.) going through the firewall exceeds a specified number in a specified time period. This information can be useful for monitoring the use of the Sidewinder services by internal users. Note: Network traffic thresholds are reported as number of events per second, and not as number of bytes per second.
202	ATTACK_ATTEMPT – This trap is sent when an attack attempt (that is, any suspicious occurrence) is identified by one of the services on the firewall. For example, if the Network Services Sentry (NSS) detects a suspicious IP address on an incoming connection, it will issue an attack attempt trap.
203	TE_VIOLATION – This trap is sent when an unauthorized user or process attempts to perform an illegal operation on a file on the firewall.
204	ACCESS_CONTROL – This trap is sent when the number of denied access attempts to services exceeds a specified number. For example, you may set up your system so that internal users cannot FTP to a certain Internet address. If a user tried to connect to that address, the attempt would be logged as a denial.
205	BAD_PROXY_AUTH – This trap occurs when a user tries to get authenticated to the telnet or FTP proxy and enters invalid data.
206	PROBE_ATTEMPT – This trap is sent when network probe attempts are detected. A network probe is any time a user attempts to connect or send a message to a TCP or UDP port that either has no service associated with it or it is associated with an unsupported service. To ignore network probe attempts, create a filter deny rule to discard probes coming from recognized offenders. See Ignoring network probe attempts for key values to configure.
207	FILTER_FAILURE – This trap occurs when the number of mail messages or HTTP messages that failed the keyword filter exceed a specified threshold in a specified time period.
208	IPSEC_FAILURE – The trap occurs when the IPsec subsystem detects a failure in authentication or encryption of network traffic. This can be caused by a number of things ranging from key configuration errors, ISAKMP problems, interoperability issues, and network attacks.
209	FAILOVER_EVENT – This trap is sent any time a Sidewinder changes its status in an HA cluster from secondary to primary, or from primary to secondary.
210	LOG_FILE_OVERFLOW – This trap is sent when the Sidewinder audit logs are close to filling the partition.
211	SYN_FLOOD_ATTACK – This trap is sent when the firewall encounters a SYN attack.
212	UPS_POWER_FAILURE – This trap is sent when a UPS device detects a power failure and the system is running on UPS battery power.
213	UPS_SYSTEM_SHUTDOWN – This trap is sent when a UPS is running out of battery power or has been on battery power for the estimated batter time.
214	LICENSE_EXCEEDED – This trap is sent when users are denied access through the firewall due to a user license cap violation.
226	CRITICAL_COMPONENT_FAILURE – This trap is sent when the firewall detects that a critical component has failed. For example, this trap occurs when daemond detects a software module has failed.
227	VIRUS_MIME_FAILURE – This trap occurs when the number of mail or HTTP messages that failed the MIME/Virus/Spyware filter exceeds a specified threshold in a specified time period.

Table 48 SNMP traps <Comment>(continued)

Number	Trap
Custom traps	
15	USER_DEFINED_DEFAULT
16	USER_DEFINED_1
17	USER_DEFINED_2
18	USER_DEFINED_3
19	USER_DEFINED_4
20	USER_DEFINED_5
21	USER_DEFINED_6
22	USER_DEFINED_7
23	USER_DEFINED_8
24	USER_DEFINED_9
25	USER_DEFINED_10

14 Network Defenses

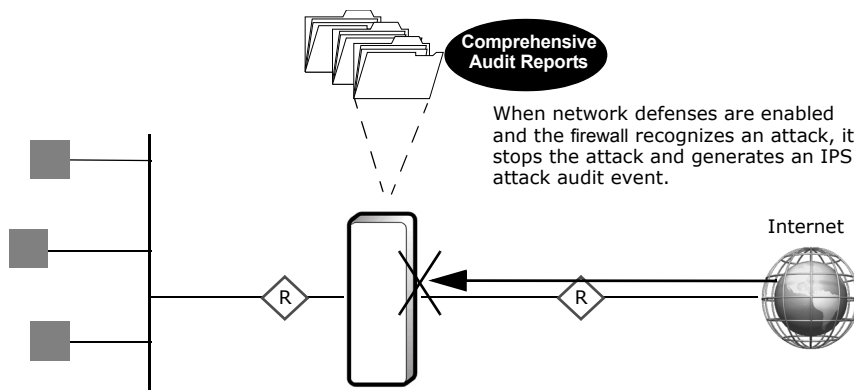
Contents

- [Viewing Network Defense information](#)
- [Configuring the TCP Network Defense](#)
- [Configuring the IP Network Defense](#)
- [Configuring the UDP Network Defense](#)
- [Configuring the ICMP Network Defense](#)
- [Configuring the ARP Network Defense](#)
- [Configuring the IPsec Network Defense tab](#)
- [Configuring the IPv6 Network Defense tab](#)

Viewing Network Defense information


Network defenses allow you to control the audit output for suspicious traffic at the data link, network, and transport layers that is detected by the Forcepoint Sidewinder when the firewall automatically prevents that traffic from entering the firewall. Some traffic is stopped because a packet, or sequence of packets, resembles a known attack. Other traffic is stopped because a packet does not comply with its protocol's standards. If network defenses are enabled, the audit reports provide detailed information on the denied traffic.

Figure 224 What happens when a network defense is enabled



If network defenses are not enabled, the firewall still stops suspicious traffic but does not generate audit.

Figure 225 What happens when a network defense is disabled



Once you decide that you want to view these denied packets' audit, you can configure the following options:

- Audit packets that the firewall determines to be part of an identifiable attack based on attack description (bad header length, bad redirect, etc.).
- Audit packets that are not specifically identified as a potential attack yet are not compliant with their protocol standards at the following levels:
 - All packets that do not comply with their protocol's standards.
 - Packets that do not comply with their protocol's standards and have been identified as a severe or moderate risk to your network.
 - Do not generate audit when the firewall stops a packet because it does not comply to its protocol's standard.

Network defenses represent one element of the firewall's audit capabilities. Information about additional auditing tools can be found in the following chapters:

- [Chapter 10, The Dashboard](#)
- [Chapter 11, Auditing](#)
- [Chapter 13, IPS Attack and System Event Responses](#)

To view the Network Defenses, select **Policy > Network Defenses**. The Network Defenses window displays with the TCP tab displayed, as shown in [Figure 226](#). All tabs are similar in appearance and function.

Figure 226 Network Defense window (TCP)

The screenshot shows the 'Network Defenses' window with the 'TCP' tab selected. At the top, there is a message: 'Select the attacks and protocol compliance issues to audit. Secure Firewall defends against all attacks and compliance issues, regardless of the audits you select.' and a 'Restore Defaults' button. Below this is a tab bar with 'TCP', 'IP', 'UDP', 'ICMP', 'ARP', and 'IPsec'. The 'TCP Audits' section is divided into two columns. The left column, 'Audit the selected TCP attacks:', contains a list of attacks with checkboxes: 'aborted connection attempt' (unchecked), 'crafted packet probe' (checked), 'forged source address' (checked), 'invalid offset' (checked), 'invalid sequence on SYN-ACK' (checked), 'LAND DOS attack' (checked), 'out of window RST' (unchecked), 'RST with no data transferred' (checked), 'SYN flood' (checked), 'syn hijack on active connection' (checked), 'SYN with FIN scan' (unchecked), and 'SYN-ACK probe' (checked). Below this list are 'Select All' and 'Deselect All' buttons. The right column, 'Audit the selected TCP compliance issues:', contains four radio buttons: 'Audit all TCP compliance issues' (unchecked), 'Audit severe and moderate TCP compliance issues' (unchecked), 'Audit severe TCP compliance issues' (checked), and 'Do not audit any TCP compliance issues' (unchecked). At the bottom, the 'TCP Audit Frequency' section has two radio buttons: 'Limit auditing (recommended)' (checked) and 'Always audit' (unchecked). Below the 'Limit auditing' option, there is a text field 'Audit the first' with the value '1', followed by 'occurrence(s) every' and another text field with the value '1', and the unit 'seconds'.

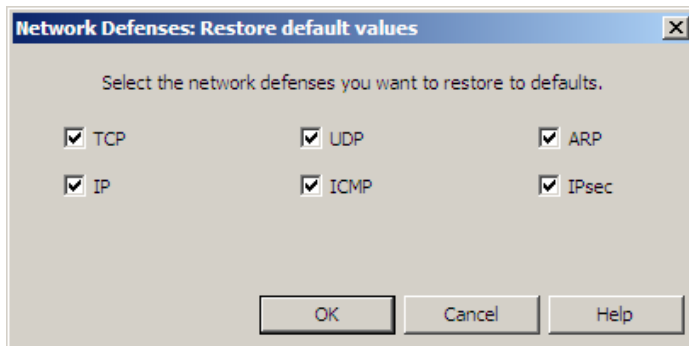
The Network Defenses tabs allows you to configure what audit the firewall generates for each of the specified protocols and how frequently to generate that audit.

For information on configuring a specific Network Defense, see the following:

- [Configuring the TCP Network Defense](#)
- [Configuring the IP Network Defense](#)
- [Configuring the UDP Network Defense](#)
- [Configuring the ICMP Network Defense](#)
- [Configuring the ARP Network Defense](#)
- [Configuring the IPsec Network Defense tab](#)

If you want to return the Network Defense settings to their defaults, click **Restore Defaults**. The following window appears:

Figure 227 Network Defenses: Restore default values window



This window allows you to restore the network defenses' attack and protocol compliance issue settings to their system defaults. When the window appears, all network defenses are selected.

- If you want to restore the defaults for all network defenses, click **OK**.
- If you want to restore the defaults for selected network defenses, clear the check box next to the network defenses that need to keep their current settings. After clearing the appropriate check box(es), click **OK**.

The selected network defenses now display and enforce their default settings.

Configuring the TCP Network Defense

The TCP Network Defense allows you to customize audit output for TCP attacks and compliance issues stopped by the firewall. To configure the TCP Network Defense, select **Policy > Network Defenses > TCP**. The following window appears:

Figure 228 Network Defenses: TCP tab

TCP | IP | UDP | ICMP | ARP | IPsec

TCP Audits

Audit the selected TCP attacks:

- ☐ aborted connection attempt
- ☒ crafted packet probe
- ☒ forged source address
- ☒ invalid offset
- ☒ invalid sequence on SYN-ACK
- ☒ LAND DOS attack
- ☐ out of window RST
- ☒ RST with no data transferred
- ☒ SYN flood
- ☒ syn hijack on active connection
- ☐ SYN with FIN scan
- ☒ SYN-ACK probe

Select All | Deselect All

Audit the selected TCP compliance issues:

- ☐ Audit all TCP compliance issues
- ☐ Audit severe and moderate TCP compliance issues
- ☒ Audit severe TCP compliance issues
- ☐ Do not audit any TCP compliance issues

TCP Audit Frequency

☒ Limit auditing (recommended)

Audit the first occurrence(s) every seconds.

☐ Always audit

Use this tab to configure what audit to generate for TCP attack and compliance issues. The firewall automatically stops all listed attacks; selecting or clearing a check box only affects whether or not this behavior is audited.

1 In the **Audit the selected TCP attacks** section, select the attacks for which you want the firewall to generate audit.

2 In the **Audit the selected TCP compliance issues** area, select how you want the firewall to audit packets that are not known attacks, but are still not compliant with the TCP standards. Options are:

- All TCP compliance issues
- Severe and moderate TCP compliance issues
- Severe TCP compliance issues
- Do not audit any TCP compliance issues

3 In the **TCP Audit Frequency** area, select how often to generate audit for TCP issues. Select one of the following:

- **Limit auditing (recommended)** – Generates an audit record for the first x occurrences for every y seconds. Other occurrences of the same audit event in that window will not be recorded. An additional audit event will be generated to record how many other audit events were suppressed.

For example, the audit is limited to generating an audit event for the first three (3) occurrences for every 60 seconds. If the firewall stopped 100 SYN-ACK probes in 60 seconds, then it generates three records for the first three denials, and then generates another audit record stating that 97 occurrences were suppressed in that 60 second window.

Limiting audit in this manner reduces system load.

- **Always audit** – Generates an audit record for every audit event.

Note: Unlimited auditing runs the risk of overflowing the log partition and creating problems for the firewall.

Options for viewing the audit output generated by these selections include:

- The Admin Console **Dashboard**
- **Monitor > Audit**
- SecurityReporter
- Third-party reporting tools

Configuring the IP Network Defense

The IP Network Defense allows you to customize audit output for IP attacks stopped by the firewall. To configure the IP Network Defense, select **Policy > Network Defenses > IP**. The following window appears:

Figure 229 Network Defenses: IP tab

The screenshot shows the 'IP' tab of the 'Network Defenses' configuration window. At the top, there are tabs for TCP, IP, UDP, ICMP, ARP, and IPsec. The 'IP' tab is selected. The main area is divided into two sections: 'IP Audits' and 'IP Audit Frequency'.

IP Audits

Audit the selected IP attacks:

- ☒ bad options
- ☒ excessively fragmented packet
- ☒ forged loopback address
- ☒ fragment attack - New Dawn
- ☒ fragment attack - Rose
- ☒ incorrect header length
- ☐ incorrect source address for interface
- ☒ invalid destination address
- ☒ missing fragment for packet
- ☒ reassembly too long
- ☒ source broadcast address
- ☒ source multicast address
- ☒ source routed packet
- ☒ unreachable destination address

Audit the selected IP compliance issues:

- ☐ Audit all IP compliance issues
- ☐ Audit severe and moderate IP compliance issues
- ☒ Audit severe IP compliance issues
- ☐ Do not audit any IP compliance issues

IP Audit Frequency

☒ Limit auditing (recommended)

Audit the first occurrence(s) every seconds.

☐ Always audit

Use this tab to configure what audit to generate for IP attack and compliance issues. The firewall automatically stops all listed attacks; selecting or clearing a check box only affects whether or not this behavior is audited.

1 In the **Audit the selected IP attacks** section, select the attacks for which you want the firewall to generate audit.

2 In the **Audit the selected IP compliance issues** area, select how you want to audit packets that are not known attacks, but are still not compliant with the IP standards. Options are:

- All IP compliance issues
- Severe and moderate IP compliance issues
- Severe IP compliance issues
- Do not audit any IP compliance issues

3 In the **IP Audit Frequency** area, select how often to generate audit for IP issues. Select one of the following:

- **Limit auditing (recommended)** – Generates an audit record for the first x occurrences for every y seconds. Other occurrences of the same audit event in that window will not be recorded. An additional audit event will be generated to record how many other audit events were suppressed.

For example, the audit is limited to generating an audit event for the first three (3) occurrences for every 60 seconds. If the firewall stopped 100 source routed packets in 60 seconds, then it generates three records for the first three denials, and then generates another audit record stating that 97 occurrences were suppressed in that 60 second window.

Limiting audit in this manner reduces system load.

- **Always audit** – Generates an audit record for every audit event.

Note: Unlimited auditing runs the risk of overflowing the log partition and creating problems for the firewall.

Options for viewing the audit output generated by these selections include:

- The Admin Console **Dashboard**
- **Monitor > Audit**
- SecurityReporter
- Third-party reporting tools

Configuring the UDP Network Defense

The UDP Network Defense allows you to customize audit output for UDP attacks stopped by the firewall. To configure the UDP Network Defense, select **Policy > Network Defenses > UDP**. The following window appears:

Figure 230 Network Defenses: UDP tab

TCP IP **UDP** ICMP ARP IPsec

UDP Audits

Audit the selected UDP attacks:

- ☒ packet too big
- ☒ zero source port

Select All Deselect All

Audit the selected UDP compliance issues:

- ☐ Audit all UDP compliance issues
- ☐ Audit severe and moderate UDP compliance issues
- ☒ Audit severe UDP compliance issues
- ☐ Do not audit any UDP compliance issues

UDP Audit Frequency

☒ Limit auditing (recommended)

Audit the first occurrence(s) every seconds.

☐ Always audit

Use this tab to configure what audit to generate for UDP attack and compliance issues. The firewall automatically stops all listed attacks; selecting or clearing a check box only affects whether or not this behavior is audited.

1 In the **Audit the selected UDP attacks** section, select the attacks for which you want the firewall to generate audit.

2 In the **Audit the selected UDP compliance issues** area, select how you want the firewall to audit packets that are not known attacks, but are still not compliant with the UDP standards. Options are:

- All UDP compliance issues
- Severe and moderate UDP compliance issues
- Severe UDP compliance issues
- Do not audit any UDP compliance issues

3 In the **UDP Audit Frequency** area, select how often to generate audit for UDP issues. Select one of the following:

- **Limit auditing (recommended)** – Generates an audit record for the first x occurrences for every y seconds. Other occurrences of the same audit event in that window will not be recorded. An additional audit event will be generated to record how many other audit events were suppressed.

For example, the audit is limited to generating an audit event for the first three (3) occurrences for every 60 seconds. If the firewall stopped 100 zero source port UDP attacks in 60 seconds, then it generates three records for the first three denials, and then generates another audit record stating that 97 occurrences were suppressed in that 60 second window.

Limiting audit in this manner reduces system load.

- **Always audit** – Generates an audit record for every audit event.

Note: Unlimited auditing runs the risk of overflowing the log partition and creating problems for the firewall.

Options for viewing the audit output generated by these selections include:

- The Admin Console **Dashboard**
- **Monitor > Audit**
- SecurityReporter
- Third-party reporting tools

Configuring the ICMP Network Defense

The ICMP Network Defense allows you to customize audit output for ICMP attacks stopped by the firewall. To configure the ICMP Network Defense, select **Policy > Network Defenses > ICMP**. The following window appears:

Figure 231 Network Defenses: ICMP tab

The screenshot shows the 'ICMP' tab selected in a configuration window. At the top, there are tabs for TCP, IP, UDP, ICMP, ARP, and IPsec. The main area is divided into two sections: 'ICMP Audits' and 'ICMP Audit Frequency'.

ICMP Audits

Audit the selected ICMP attacks:

- ☒ invalid redirect

Audit the selected ICMP compliance issues:

- ☐ Audit all ICMP compliance issues
- ☐ Audit severe and moderate ICMP compliance issues
- ☒ Audit severe ICMP compliance issues
- ☐ Do not audit any ICMP compliance issues

Buttons: Select All, Deselect All

ICMP Audit Frequency

☒ Limit auditing (recommended)

Audit the first occurrence(s) every seconds.

☐ Always audit

Use this tab to configure what audit to generate for ICMP attack and compliance issues. The firewall automatically stops all listed attacks; selecting or clearing a check box only affects whether or not this behavior is audited.

1 In the **Audit the selected ICMP attacks** section, select the attacks for which you want the firewall to generate audit.

2 In the **Audit the selected ICMP compliance issues** area, select how you want the firewall to audit packets that are not known attacks, but are still not compliant with the ICMP standards. Options are:

- All ICMP compliance issues
- Severe and moderate ICMP compliance issues
- Severe ICMP compliance issues
- Do not audit any ICMP compliance issues

3 In the **ICMP Audit Frequency** area, select how often to generate audit for ICMP issues. Select one of the following:

- **Limit auditing (recommended)** – Generates an audit record for the first x occurrences for every y seconds. Other occurrences of the same audit event in that window will not be recorded. An additional audit event will be generated to record how many other audit events were suppressed.

For example, the audit is limited to generating an audit event for the first three (3) occurrences for every 60 seconds. If the firewall stopped 100 invalid redirect ICMP attacks in 60 seconds, then it generates three records for the first three denials, and then generates another audit record stating that 97 occurrences were suppressed in that 60 second window.

Limiting audit in this manner reduces system load.

- **Always audit** – Generates an audit record for every audit event.

Note: Unlimited auditing runs the risk of overflowing the log partition and creating problems for the firewall.

Options for viewing the audit output generated by these selections include:

- The Admin Console **Dashboard**
- **Monitor > Audit**
- SecurityReporter
- Third-party reporting tools

Configuring the ARP Network Defense

The ARP Network Defense allows you to customize audit output for ARP attacks stopped by the firewall. To configure the ARP Network Defense, select **Policy > Network Defenses > ARP**. The following window appears:

Figure 232 Network Defenses: ARP tab

The screenshot shows the ARP tab in the Network Defenses configuration window. At the top, there is a navigation bar with tabs for TCP, IP, UDP, ICMP, ARP (selected), and IPsec. The main content area is divided into two sections: "ARP Audits" and "ARP Audit Frequency".

ARP Audits

Audit the selected ARP compliance issues:

- ☐ Audit all ARP compliance issues
- ☐ Audit severe and moderate ARP compliance issues
- ☒ Audit severe ARP compliance issues
- ☐ Do not audit any ARP compliance issues

ARP Audit Frequency

☒ Limit auditing (recommended)

Audit the first occurrence(s) every seconds.

☐ Always audit

Use this tab to configure what audit to generate for ARP compliance issues. The firewall automatically stops all listed attacks; selecting or clearing a check box only affects whether or not this behavior is audited.

1 In the **Audit the selected ARP compliance issues** area, select how you want the firewall to audit packets that are not known attacks, but are still not compliant with the ARP standards. Options are:

- All ARP compliance issues
- Severe and moderate ARP compliance issues
- Severe ARP compliance issues
- Do not audit any ARP compliance issues

2 In the ARP **Audit Frequency** area, select how often to generate audit for ARP issues. Select one of the following:

- **Limit auditing (recommended)** – Generates an audit record for the first x occurrences for every y seconds. Other occurrences of the same audit event in that window will not be recorded. An additional audit event will be generated to record how many other audit events were suppressed.

For example, the audit is limited to generating an audit event for the first three (3) occurrences for every 60 seconds. If the firewall stopped 100 ARP attacks in 60 seconds, then it generates three records for the first three denials, and then generates another audit record stating that 97 occurrences were suppressed in that 60 second window.

Limiting audit in this manner reduces system load.

- **Always audit** – Generates an audit record for every audit event.

Note: Unlimited auditing runs the risk of overflowing the log partition and creating problems for the firewall.

Options for viewing the audit output generated by these selections include:

- The Admin Console **Dashboard**
- **Monitor > Audit**
- SecurityReporter
- Third-party reporting tools

Configuring the IPsec Network Defense tab

The IPsec Network Defense allows you to customize audit output for IPsec attacks stopped by the firewall. Unlike the other network defenses, it also allows you to control non-malicious failure audits. To configure the IPsec Network Defense, select **Policy > Network Defenses > IPsec**. The following window appears:

Figure 233 Network Defenses: IPsec tab

Use this tab to configure what audit to generate for IPsec attacks, non-malicious failures, and compliance issues. The firewall automatically stops all listed attacks; selecting or clearing a check box only affects whether or not this behavior is audited.

Note: The IPsec Network Defense allows you to directly control audit output for some non-malicious failures because IPsec tends to have more of these types of failures than other protocols.

- 1** In the **Audit the selected IPsec attacks** section, select the attacks for which you want to generate audit.
 - 2** In the **Audit the selected IPsec compliance issues** area, select how you want to audit packets that are not known attacks, but are still not compliant with the IPsec standards. Options are:
 - All IPsec compliance issues
 - Severe and moderate IPsec compliance issues
 - Severe IPsec compliance issues
 - Do not audit any IPsec compliance issues
 - 3** In the **IP Audit Frequency** area, select how often to generate audit for IPsec issues. Select one of the following:
 - **Limit auditing (recommended)** – Generates an audit record for the first x occurrences for every y seconds. Other occurrences of the same audit event in that window will not be recorded. An additional audit event will be generated to record how many other audit events were suppressed.

For example, the audit is limited to generating an audit event for the first three (3) occurrences for every 60 seconds. If the firewall stopped 100 decryption failures in 60 seconds, then it generates three records for the first three denials, and then generates another audit record stating that 97 occurrences were suppressed in that 60 second window.

Limiting audit in this manner reduces system load.
 - **Always audit** – Generates an audit record for every audit event.
- Note:** Unlimited auditing runs the risk of overflowing the log partition and creating problems for the firewall.

Options for viewing the audit output generated by these selections include:

- The Admin Console **Dashboard**
- **Monitor > Audit**
- SecurityReporter
- Third-party reporting tools

Configuring the IPv6 Network Defense tab

Note: The IPv6 tab appears only if IPv6 is enabled on your firewall.

The IPv6 Network Defense allows you to customize audit output for IPv6 attacks stopped by the firewall. To configure the IPv6 Network Defense, select **Policy > Network Defenses > IPv6**. The following window appears:

Figure 234 Network Defenses: IPv6 tab

TCP IP UDP ICMP ARP IPsec IPv6

IP Audits

Audit the selected IP6 attacks:

- ☒ address scope conflict
- ☒ header too small
- ☒ hop-by-hop header malformed
- ☒ hop-by-hop jumbo option malformed
- ☒ invalid address
- ☒ invalid version
- ☒ link layer broadcast
- ☒ loopback address spoofed
- ☒ malformed option
- ☒ netprobe
- ☒ option header length incorrect
- ☒ too many headers
- ☒ unknown hop-by-hop option

Select All Deselect All

Audit the selected IP6 compliance issues:

- ☐ Audit all IP6 compliance issues
- ☐ Audit severe and moderate IP6 compliance issues
- ☒ Audit severe IP6 compliance issues
- ☐ Do not audit any IP6 compliance issues

IP Audit Frequency

☒ Limit auditing (recommended)

Audit the first occurrence(s) every seconds.

☐ Always audit

Use this tab to configure what audit to generate for IPv6 attacks and compliance issues. The firewall automatically stops all listed attacks; selecting or clearing a check box only affects whether or not this behavior is audited.

- 1 In the **Audit the selected IPv6 attacks** section, select the attacks for which you want to generate audit.
- 2 In the **Audit the selected IPv6 compliance issues** area, select how you want to audit packets that are not known attacks, but are still not compliant with the IPv6 standards. Options are:

- All IPv6 compliance issues
- Severe and moderate IPv6 compliance issues
- Severe IPv6 compliance issues
- Do not audit any IPv6 compliance issues

- 3 In the **IP Audit Frequency** area, select how often to generate audit for IPv6 issues. Select one of the following:

- **Limit auditing (recommended)** – Generates an audit record for the first x occurrences for every y seconds. Other occurrences of the same audit event in that window will not be recorded. An additional audit event will be generated to record how many other audit events were suppressed.

For example, the audit is limited to generating an audit event for the first three (3) occurrences for every 60 seconds. If the firewall stopped 100 decryption failures in 60 seconds, then it generates three records for the first three denials, and then generates another audit record stating that 97 occurrences were suppressed in that 60 second window.

Limiting audit in this manner reduces system load.

- **Always audit** – Generates an audit record for every audit event.

Note: Unlimited auditing runs the risk of overflowing the log partition and creating problems for the firewall.

Options for viewing the audit output generated by these selections include:

- The Admin Console **Dashboard**
- **Monitor > Audit**
- SecurityReporter
- Third-party reporting tools

15 The SNMP Agent

Contents

[Understanding SNMP options](#)

[Overview of Sidewinder as a managed node](#)

[Setting up the SNMP agent on Sidewinder](#)

[Sending SNMP traffic through Sidewinder](#)

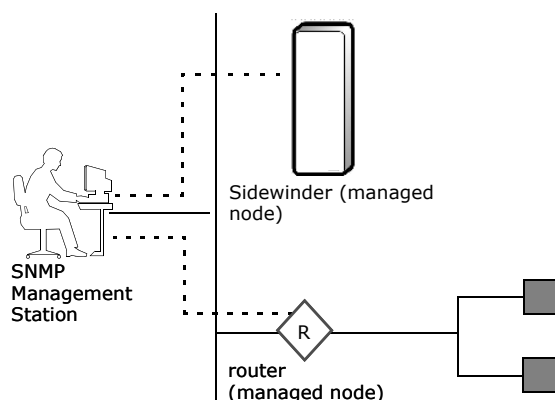
Understanding SNMP options

This section introduces Simple Network Management Protocol (SNMP) concepts and explains how to configure the Forcepoint Sidewinder SNMP agent. It also explains what needs to be done to allow the firewall to send or route SNMP messages to remote systems in an external network.

SNMP is the industry standard for network management. The Sidewinder supports SNMP v1, SNMP v2c, and SNMP v3. The firewall can participate in SNMP management in two different ways:


- You can set up SNMP agent software that allows the firewall to be an SNMP-managed node. A node is monitored by SNMP-compliant network management stations located on one of the firewall's burbs.

Figure 235 Managing distributed systems using SNMP



- Using the SNMP proxy, you can configure the firewall to route SNMP messages from a management station through the firewall to an SNMP agent on a system in an external network.

Figure 236 Managing distributed systems using SNMP



Note: If you want your firewall to simultaneously act as an SNMP agent and pass SNMP Managing distributed systems using SNMP traffic in the same burb, you will need to use a TCP/UDP Packet Filter service to pass the SNMP traffic. See [Create and modify services](#).

Overview of Sidewinder as a managed node

The following sections describe how the Sidewinder interacts with SNMP management stations:

- [Communicating with an SNMP management station](#)
- [About Sidewinder SNMP traps](#)
- [About Sidewinder SNMP MIBs](#)
- [About the management station](#)


Communicating with an SNMP management station

A network that is managed using SNMP involves two primary components: a manager (management station) and a number of managed nodes. The management station is typically a PC or UNIX workstation running network management software such as Hewlett-Packard's OpenView® or the freeware Multi-Router Traffic Grapher (MRTG). Managed nodes are networking devices such as routers or firewalls that contain an SNMP agent. [Figure 237](#) shows a management station communicating with an SNMP node to obtain network configuration information.

The management station uses the management software to display a graphical representation of a network's topology. In general, network managers can monitor SNMP nodes (including Sidewinder) by clicking icons that represent each node in the network's topology.

A management station in an internal or external network can request information from a managed node's SNMP agent. The SNMP management station sends a managed node **get** and **getnext** SNMP messages to retrieve node-specific parameters and variables, called *objects*. The message response from the managed system provides the SNMP administrator with information on a node's device names, status, network connections, etc.

Figure 237 Communication between a management station and a managed node



Note: SNMP agents typically allow **Get**, **GetNext**, and **Set** requests from the management station. However, the firewall SNMP agent does not support **Set** requests. This prevents a management system from sending commands to change variables or parameters in the firewall.

Each managed node can send an unsolicited event notification message, called a *trap*, to a management station when it detects certain system events. For example, you can configure the firewall audit system to issue a trap whenever an unauthorized user tries to read, write, or execute a protected file on the firewall. (Refer to [Sidewinder SNMP traps](#) for a list of all firewall-supported traps.)

- When setting up SNMP management for SNMP v1 or SNMP v2c, a network administrator assigns the management station and the nodes it will manage a *community name*. As shown in the following figure, the community name is in the authentication header in each SNMP message exchanged between a management station and a managed node.

Figure 238 Community name within an SNMP message

Version	Community Name	SNMP command: Get, GetNextRequest, etc.
---------	----------------	-----------------------------------------

The SNMP agent treats the community name like a password to validate the identity of a management station. For example, suppose a management station sends a **get** request to retrieve information from a managed node's SNMP agent. If the community name within the **get** request is not also used by the SNMP agent, the agent will not return information to the management station.

Caution: To increase security on your network, **do not** use common default names such as "public" or "private," which can be easily guessed.

- SNMP management in SNMP v3 requires a user name and password. The password is encrypted in SNMP messages, increasing the security.

Note: The SNMP v3 password is used as the encryption password.

Both the management station and the managed node also contain Management Information Bases (MIBs) that store information about the managed objects. Currently, the SNMP agent on the firewall supports standard MIB II objects, the Host Resources MIB (RFC 1514), and the firewall-specific MIB objects. MIBs are discussed in greater detail in [About Sidewinder SNMP MIBs](#).

Note: Sidewinder MIB files are located in `/secureos/etc/snmp` on the firewall's file system.

If you need more information on SNMP, an excellent source is *Managing Internetworks with SNMP* by Mark A. Miller, P.E. (M&T Books).

About Sidewinder SNMP traps

An SNMP trap is an alert message that is sent as an unsolicited transmission of information from a managed node (router, firewall, etc.) to a management station. Most management stations can be configured to either: (1) display received traps in a pop-up window, or (2) automatically dial a phone number, such as a pager number.

The Sidewinder SNMP agent supports a basic trap, called the **ColdStart trap**, that is sent whenever the SNMP agent is enabled. It is also sent if the Admin Console modifies the SNMP configuration file (`/secureos/etc/snmp/snmpd.conf`). You cannot disable the ColdStart trap.

You also have the option to configure the firewall to send audit alert SNMP traps when an audit event triggers a response. Additional information about requesting and configuring SNMP traps is available in [Sidewinder SNMP traps](#).

About Sidewinder SNMP MIBs

Management Information Bases (MIBs) are associated with both the management station and the Sidewinder SNMP agent. The SNMP agent supports two MIB structures (as well as a Host MIB).

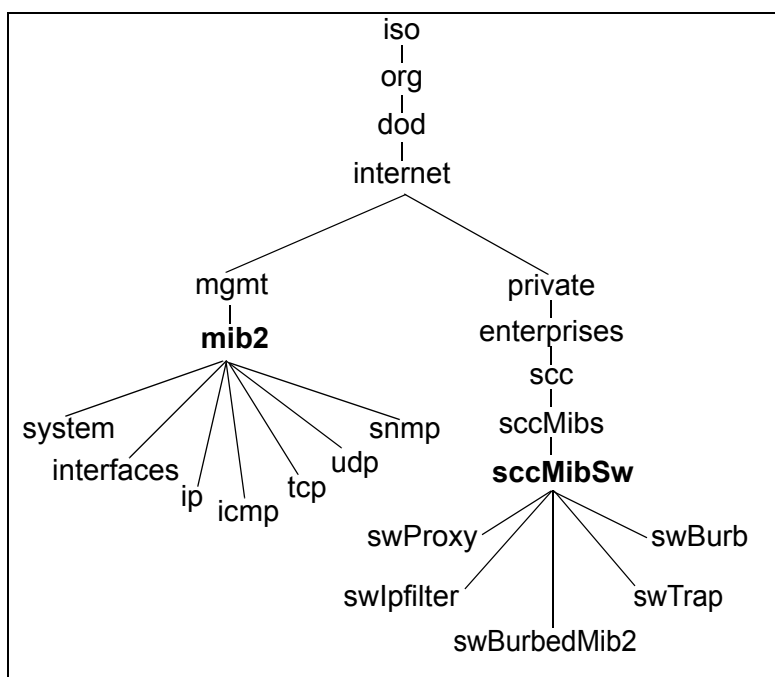
- **mib2** – This is a standard SNMP MIB as defined in RFC 1213.
- **sccMibSw** – This is a firewall-specific MIB. [Figure 239](#), located on the following page, shows the location of the firewall MIB structures within the SNMP root hierarchy.

Note: Sidewinder MIB files are located in /secureos/etc/snmp on the firewall's file system.

Individual objects (parameters and variables) managed by an SNMP management station are part of an object group within an MIB. For example, the **swProxy** group stores information about currently-defined proxies on the system. The information might include the proxy name and the current status of the proxy.

When a management station requests information from the Sidewinder SNMP agent, the SNMP agent may or may not associate the returned information with a specific *burb*.

Figure 239 MIBs supported by the Sidewinder SNMP agent



About the management station

The administrator of the SNMP management station should be made aware of the following in order to retrieve information from the Sidewinder SNMP agent:

- **Sidewinder host name or IP address**

This is needed to set up communication with the firewall. Note:

- If the burb in which the SNMP agent is running contains more than one interface, specify the address of the first interface in the burb. The SNMP agent will only respond to the first interface in the burb.
- If you are using High Availability (HA), specify the shared HA common IP address or host name, not the actual interface address or host name.

- **Community names configured in the Sidewinder SNMP agent**

If you are using SNMP v1 or SNMP v2c, this is needed to allow the management station to retrieve MIB objects from the SNMP agent.

- **SNMP v3 user name and password**

If you are using SNMP v3, you must configure the firewall with the user names and passwords established on the management station.

Note: The SNMP v3 password is also used as the encryption password.

- **MIB information**

This may be needed to properly translate the object identifications. Inform the administrator that the firewall supports the Host Resources MIB.

Forcepoint Sidewinder MIB files are located in */secureos/etc/snmp* on the firewall's file system. The files can be accessed directly on the firewall or downloaded from the Internet via an FTP client or web browser. The MIB files are *SCC-MIB.txt* and *SCC-SW-MIB.txt*.

- To retrieve the files by FTP, from your FTP client log into <ftp://sidewinder.downloads.forcepoint.com/>. The files are located in */pub/mibs*.
- To retrieve the files using a web browser, point the browser to <ftp://sidewinder.downloads.forcepoint.com/pub/mibs>.

Setting up the SNMP agent on Sidewinder

This section explains how to configure the SNMP agent on Sidewinder. It involves the following steps:


- Configure the SNMP agent (**Policy > Rule Elements > Services > SNMP Agent**).
- Create a rule allowing access from the management station to the Sidewinder SNMP agent (**Policy > Rules**).
Note: If you are configuring SNMP on firewall that is part of an HA cluster, all firewall queries must use the HA cluster address.
- Send custom traps (for example, from shell scripts) using the `snmptrap` command. See [Sidewinder SNMP traps](#) and the `snmptrap` man page.
- Use the IPS Attack and System Event Responses (**Monitor > IPS Attack Responses/System Event Responses > Response** tab) to manage when the firewall sends SNMP traps to its management station. See [Chapter 13, IPS Attack and System Event Responses](#).

All of these steps play an important role in providing your SNMP management station with information.

Configuring the SNMP agent

To set up the SNMP agent, select **Policy > Rule Elements > Services**, then double-click `snmpd` and click **Properties**. The SNMP Agent Configuration window appears.

Figure 240 SNMP Agent Configuration window



The image shows the 'Modify Service: SNMP Agent Configuration' window. It contains several sections for configuring the SNMP agent. On the left, there are fields for 'Location' and 'Contact', and a radio button for 'Enable authentication failure trap' (set to 'No'). Below these is a section for 'Allowed Protocols' with checkboxes for 'SNMP v1', 'SNMP v2c', and 'SNMP v3', all of which are checked. At the bottom left is a table for 'Allowed get communities' with one entry 'public'. On the right, there is a section for 'SNMP v3 users' with a table for 'User' and 'Minimum security level'. Below this is a 'Trap version' dropdown set to 'v2c' and a 'v3 settings' button. At the bottom right is a table for 'Trap destinations' with columns for 'Host' and 'Community'. The window has 'New', 'Modify', and 'Delete' buttons for each table, and 'OK', 'Cancel', and 'Help' buttons at the bottom.

Use this window to enter configuration information for the SNMP agent.

To set up the SNMP agent:

- 1 [Optional] In the **Location** field, type a description of the physical location of your firewall.
- 2 [Optional] In the **Contact** field, type an identifying name, such as your firewall administrator user name or e-mail address.

- 3 In the **Enable Authentication Failure Trap** field, select **Yes** to enable authentication failure traps, or **No** to disable authentication failure traps. If you click **Yes**, the firewall will send authentication failure traps to all configured management stations whenever it detects an unauthenticated **Get** command.
- 4 In the Allowed Protocols area, select the versions of SNMP that incoming SNMP requests are allowed to use. SNMP message with versions that are not allowed are ignored.
- 5 From the **Trap version** drop-down list, select the SNMP version that the firewall should use when sending traps.

Note: This is a global setting that will affect all components that originate traps.

- 6 [Conditional] If you select trap version v3, click **v3 settings** and configure the security settings to use when sending traps:
 - **Username** and **Password** – Enter the user name and password to use when sending traps. All trap destinations will use the same SNMP user when using SNMP v3. Enter the password again to confirm.
 - **Security level** – From the drop-down list, select whether authentication and encryption should be used when sending traps:
 - **noAuth** – No authentication or encryption is required.
 - **authNoPriv** – A password is required. Payload encryption is not used.
 - **authPriv** – A password and payload encryption are required.
- 7 Click **OK** to return to the SNMP Agent Configuration window.
- 8 Use the SNMP v3 users list to view, create, and manage SNMP v3 users who can issue requests to the Sidewinder SNMP agent.

To configure SNMP v3 users who can issue requests to the Sidewinder SNMP agent, click **New** and enter the appropriate information:

- **Username** – Enter the user name established on the SNMP management station.
 - **Description** – Optionally enter a description to easily identify this user.
 - **Password** – Enter the password established on the SNMP management station. Enter the password again to confirm.
- Note:** The SNMP v3 password is used as the encryption password.
- **Minimum security level** – From the drop-down list, select whether authentication and encryption should be used when issuing requests:
 - **noAuth** – Any security level can be used.
 - **authNoPriv** – A password is required. Payload encryption is optional.
 - **authPriv** – A password and payload encryption are required.
- 9 Click **OK** to return to the SNMP Agent Configuration window.

10 In the Allowed Get Communities list, you can view all of the community names authorized to retrieve MIB information. The community name is part of the authentication header in all SNMP messages. The Sidewinder SNMP agent checks the community name in all v1 and v2c SNMP messages it receives to verify the identity of a manager.

To add, modify, or delete communities, use the **New**, **Modify**, and **Delete** buttons located directly beneath the list.

- The SNMP agent will not start unless a community name is specified. By default, if you do not specify an **Allowed Get Community** name, then only **Allowed Get Community** is “public.”
- Communities are ignored in SNMP v3.

11 In the Trap Destinations list, you can view all of the hosts that will receive traps generated by the Sidewinder SNMP agent.

To add, modify, or delete trap destinations, use the **New**, **Modify**, and **Delete** buttons located directly beneath the list.

- By default, if you do not specify a trap destination community name, the firewall uses the community name “public.”
- If the trap version selected is v3, the community name in Trap destinations is ignored.

12 Click **OK** to return to the SNMP Agent Configuration window.

Be sure to save your changes when you return to the main Services window. Once you create an enabled rule with the SNMP agent as the service, a ColdStart trap is issued to all configured trap destinations.

Creating a rule to allow access to the SNMP agent

You must create a rule that allows SNMP queries to reach the Sidewinder SNMP agent. (For information on creating rules, see [Creating, modifying, and duplicating rules](#).)

- If the management station is in a trusted, internal burb, create the following rule to allow traffic between the management station and the Sidewinder SNMP agent:

Table 49 Key features in the SNMP agent rule

Rule area	Value
Service	SNMP agent
Source Burb	Must be a single burb and must match the destination burb. The SNMP agent will be enabled in the burb selected here.
Destination Burb	Must match the source burb.

- The SNMP agent can only be enabled in one burb. If you have management stations in other burbs that must reach the SNMP agent, see [Sending SNMP traffic through Sidewinder](#) for information on creating rules for those situations.

Sending SNMP traffic through Sidewinder


You can route (or forward) SNMP messages between a management station behind the firewall and any SNMP managed node on the other side of the firewall. If your management station is in an untrusted burb, or you have multiple management stations in different burbs, you can also allow access to the Sidewinder SNMP agent using the SNMP proxy in a rule. This section describes three scenarios that use SNMP and provides guidance on how to set up the necessary rules.

The Sidewinder SNMP proxy sends SNMP requests and messages via UDP port 161. That proxy sends SNMP traps to an external management station via UDP port 162.

The following figure displays the following three scenarios:

- 1 Passing traffic from an internal SNMP management station through the firewall via the SNMP proxy to an external managed node (SNMP agent). Set the rule's service to the SNMP proxy. The source and destination burbs will be different (for example, internal to external).
- 2 Passing traffic from a management station to the Sidewinder SNMP agent when both are located on the same burb. This scenario does not use the SNMP proxy. Set the rule's service to SNMP agent. The source and destination burb must be the same (for example, internal to internal).
- 3 Passing traffic across burb boundaries to the SNMP agent. Although only one SNMP agent is allowed to operate on the firewall, access through other burbs is supported using the SNMP proxy. To allow SNMP management stations that reside in other burbs to connect to the SNMP agent, you must create an allow rule using the SNMP proxy. The source for this rule should consist of a network object group that contains only SNMP management station IP addresses. The destination should specify the destination IP address for the burb in which SNMP is running. Using redirection is common in this scenario (for example, external to external with redirection to the internal interface).

Figure 241 Sidewinder serving as an SNMP agent for internal or external management station



SECTION 4

Networking

Chapter 16, Burbs, Interfaces, and Quality of Service

Chapter 17, Routing

Chapter 18, DNS (Domain Name System)

Chapter 19, E-mail

Chapter 20, Virtual Private Networks

16 Burbs, Interfaces, and Quality of Service

Contents

[Configuring burbs](#)

[Configuring interfaces](#)

[Configuring Quality of Service](#)


Configuring burbs

A burb is a type enforced network area used to isolate network interfaces from each other.

- An internal burb and an external burb are defined in your Forcepoint Sidewinder during the installation process.
- You create, modify, and delete burbs in the Burb Configuration window.
- You select these burbs as Source and Destination burbs when creating a rule in the Rules window.

To create, modify, and delete burbs, select **Network > Burb Configuration**. The Burb Configuration window appears.

Figure 242 Burb Configuration window



Name	Groups	Description
external - Internet burb		
internal		

Burb name: ID: 1

Description:

Groups

Connection Options


- ☐ Honor ICMP redirect
- ☐ Respond to ICMP echo and timestamp
- ☒ Hide port unreachables

The upper pane lists the existing burbs and burb groups. When you select a burb or burb group in the table, the properties appear in the lower pane.

- You can configure a maximum of 63 burbs on a Sidewinder.

- The *Internet* burb cannot be deleted. The Internet burb has pre-defined attributes, both configurable and non-configurable, to supply a secured connection to the internet.
- At least two burbs (in addition to the *Firewall* burb) must exist at all times.
- A virtual burb is a burb that does not contain a network interface.
 - Virtual burbs are used to apply security policy to VPN traffic.
 - Virtual burbs do not support ICMP.

Figure 243 Burb Configuration toolbar



Use the toolbar to perform these actions:

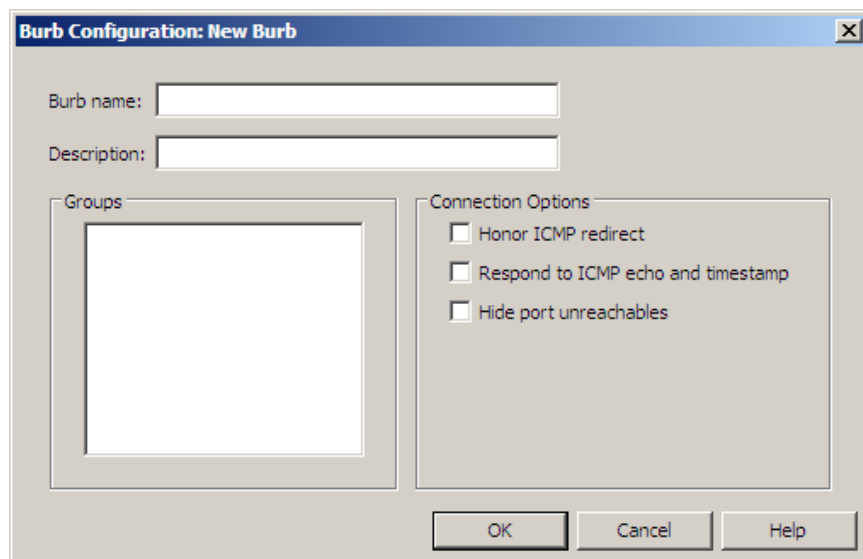
Table 50 Burb Configuration toolbar

Icon	Action
New Burb	Create a new burb by clicking New Burb and entering burb information in the pop-up window. See Creating or modifying a burb for details.
New Burb Group	Create a new burb group by clicking New Burb Group and entering information in the pop-up window. When configured, the group appears in the Groups list in the lower pane. See Creating or modifying a burb group for details.
Modify	Modify a burb or burb group by selecting it in the upper pane, and then clicking Modify . Modify the settings in the pop-up window. (Read-only administrators can click View to view a burb or burb group.) You can also select a burb or burb group and modify the settings in the lower pane.
Delete	Delete a burb or burb group by selecting the burb in the upper pane and clicking Delete . You cannot delete a burb or burb group that is currently referenced elsewhere on the system (for example, a rule or interface configuration). To determine whether a burb or burb group is currently being referenced, select it and click Usage .
Usage	View all areas where a burb or burb group is currently being used by selecting the burb in the upper pane and clicking Usage . The Burb Usage window appears listing every area in which the burb is currently used.
Internet Burb	Designate the internet burb by selecting the appropriate burb from the Internet Burb drop-down list.

Creating or modifying a burb

Use this window to create or modify a burb.

Figure 244 New/Modify Burb window



To create or modify a burb:

- 1 Type a name in the **Burb name** field. This is the name you will see in the Burb drop-down list in the Rules window.
 - Do not use “Firewall” or “firewall” as a burb name, as this name is already used elsewhere in the Sidewinder.
 - Case matters in burb names. For example, if you create a burb named *Joe* and another burb named *joe*, they are separate burbs.
 - If you are modifying a burb, you cannot change the name.
- 2 [Optional] Type a more detailed description of the burb.
- 3 Select connection options for the burb:
 - **Honor ICMP redirect** – ICMP messages are used to optimize the routes for getting IP traffic to the proper destination. On a trusted network, honoring ICMP redirects can improve the throughput of the system. On an untrusted network, ICMP redirects can be used by hackers to examine, reroute, or steal network traffic. Enabling this parameter allows the firewall to honor ICMP redirects.
 - **Respond to ICMP echo and timestamp** – ICMP echo and timestamp messages (also known as *ping* messages) are used to test addresses on a network. The messages are a handy diagnostic tool, but can also be used by hackers to probe for weaknesses. Enabling this parameter allows the firewall to respond to these messages.
 - **Hide port unreachable's** – If this parameter is enabled, the firewall will give no response if a node on the network attempts to connect to a port on which the firewall is not listening. This increases security by not divulging configuration information to potential hackers.


Note: Do not select this option for a heartbeat burb in an HA cluster.
- 4 [Optional] In the Groups list, select a burb group or burb groups for this burb to belong to.
- 5 Click **OK** and save your changes.

Creating or modifying a burb group

Use this window to create or modify a burb group.

Burb groups are a way to apply a rule to multiple burbs. If you select a burb group in the Source and Destination areas in a rule, that rule will apply to each burb in the burb group.

Figure 245 New/Modify Burb Group window



The image shows a dialog box titled "Burb Configuration: New Burb Group". It contains three input fields: "Group name:", "Description:", and "Burbs:". The "Burbs:" field has two checkboxes, "external" and "internal", both of which are currently unchecked. At the bottom of the dialog box are three buttons: "OK", "Cancel", and "Help".

To create or modify a burb group:

- 1 Type a name in the **Group name** field. This is the name you will see in the Burb drop-down list in the Rules window. (If you are modifying the burb group, you cannot change the name.)
- 2 [Optional] Type a more detailed description of the burb group.
- 3 In the Burbs list, select which burbs belong to this group.
- 4 Click **OK** and save your changes.

Configuring interfaces

This section provides information on configuring and modifying interfaces on Sidewinder.

For information about how interfaces are used on Sidewinder, see the following:

- [About interfaces](#)
- [About IPv6 addresses on Sidewinder](#)

For instructions on creating and modifying interfaces on your firewall, see the following:

- See [About the Interfaces: Interface Configuration tab](#) to configure and modify interfaces.
- See [About the Interfaces: NIC and NIC Group Configuration tab](#) to configure the physical NICs and to create NIC groups for redundant NICs.
- See [Creating interfaces](#) for procedures to configure different types of interfaces and interface elements.

About interfaces

A Sidewinder interface is a logical representation of network interface hardware. Network configuration settings are applied to the interface, which is associated with a NIC or NIC group (the network interface card hardware). The relationship between interfaces and NICs allows you to easily move your network configuration to different network hardware by assigning the interface to a different NIC or NIC group.

Table 51 Configurable properties of NICs and interfaces

NIC configurable properties	Interface configurable properties
<ul style="list-style-type: none">• Media type• Media capabilities	<ul style="list-style-type: none">• IP address(es)• VLAN ID• NIC or NIC Group• Burb• Quality of Service• MTU size (Maximum Transmission Unit)

NICs are configured separately from the interface.

- You can modify the media type and media capabilities of a NIC.
- You select an available NIC for an interface when creating the interface.
- You can create NIC groups to use for the redundant NIC function: if the primary NIC in a group stops working or is disconnected, the standby NIC starts passing the traffic.
- Note the following interface–NIC association rules:
 - A NIC can be referenced by only one enabled non-VLAN interface.
 - A NIC can be referenced by multiple enabled VLAN interfaces.
 - A NIC cannot be referenced by enabled VLAN and non-VLAN interfaces simultaneously.

Note: These rules do not apply to disabled interfaces. For example, multiple interfaces can reference the same NIC as long as only one of those interfaces is enabled at a time.

The internal and external network interfaces of the Sidewinder are defined during the initial configuration. These interfaces have IPv4 addresses.

- An interface can have IPv4 addresses, IPv6 addresses, or both.
- By using VLANs, you can create up to:
 - 512 interfaces on a standalone firewall.
 - 255 interfaces on a High Availability cluster.

About IPv6 addresses on Sidewinder

Sidewinder offers a limited implementation of IPv6 facilities.

IPv6 addressing is enabled and addresses are entered on the Interface Properties window. See [Enter an IPv6 address on an interface](#) for information.


IPv6 addresses are supported in the following features:

- **Rules** – All policy management is done through packet filter rules. Note the following:
 - All packet filter rule types (TCP, UDP, ICMP, FTP, and Other) are supported.
 - Create IPv6 network objects to use in a rule. Endpoints in a rule must have the same type of address.
 - An IPv4 source can connect only to an IPv4 destination.
 - An IPv6 source can connect only to an IPv6 destination.

Sidewinder can pass both kinds of traffic using dual stack architecture as shown in [Figure 246](#) below.

- Network address translation (NAT) is not supported.
- Redirect endpoint is not supported; redirect port is supported.
- Proxy rules are not supported.
- **DNS** – Single, unbound DNS configuration is supported.
- **VPN** – Manually keyed VPNs are supported.

Figure 246 Dual stack architecture




Administrators should be familiar with IPv6 addressing conventions and uses before enabling IPv6 on their Sidewinder. See the following resources for more information on IPv6 addressing:

- RFC 2460 [IPv6 Specification](#)
- RFC 4193 [Unique Local IPv6 Unicast Addresses](#)
- RFC 4291 [IPv6 Addressing Architecture](#)
- *IPv6 Essentials*, 2nd ed. by Silvia Hagen (O'Reilly)

About the Interfaces: Interface Configuration tab

To create and modify interfaces, select **Network > Interfaces**. The Interface Configuration tab appears.

Figure 247 Interfaces: Interface Configuration tab



Use this tab to configure and modify interfaces. You can create an unlimited number of interfaces. The table lists all configured interfaces as well as each NIC or NIC group that is not in use.

Figure 248 Interface Configuration tab toolbar



Use the toolbar to perform these actions:

Table 52 Interface Configuration tab toolbar

Icon	Action
New	Create a new interface by clicking New and entering network link information in the pop-up window. See Creating interfaces for details.
Modify	Modify an interface by selecting it and then clicking Modify . Modify the settings in the pop-up window. See Creating interfaces for details.
Swap Parameters	Switch the configuration settings between two interfaces by selecting both interfaces (press and hold the Ctrl key while selecting the interfaces) and clicking Swap Parameters . This action essentially swaps the names of the selected interfaces. The NIC is still associated with the same IP address, burbs, and other attributes.
Delete	Delete an interface by selecting it and clicking Delete . Note: This deletes the link data for the interface. It does not delete the NIC. A NIC must be physically removed to remove it from the list.
Rename	To rename an interface, select it and click Rename and type a new name in the pop-up window.
Find	Search for a specific element(s) in the list using the Find field. Type your search criteria, and interfaces with matching elements will appear in the list. Clear this field to see the full list again.
Show Status	To view the status of interfaces and their associated NICs, click Show Status . You can also restart NICs and ping addresses in the pop-up window.

About the Interfaces: NIC and NIC Group Configuration tab

To create and modify hardware parameters for NICs:

Select **Network > Interfaces**. The Interface Configuration tab appears.

Click the **NIC and NIC Group Configuration** tab.

Figure 249 Interfaces: NIC and NIC Group Configuration tab

NIC and NIC Group Configuration				
NIC Or NIC Group ▲	MAC Address	Used By	Media Type	Media Capabilities
em0	00:04:23:bc:4d:8c	external_network	autoselect	rxcsu txcsu
em1	00:04:23:bc:4d:8d	internal_network	autoselect	rxcsu txcsu
em2	00:04:23:bc:4d:8e		autoselect	rxcsu txcsu
em3	00:04:23:bc:4d:8f		autoselect	rxcsu txcsu

Use this tab to modify hardware parameters for Network Interface Cards (NICs) and to create NIC groups used for redundant NICs.

Note: To delete an individual NIC, you must physically remove it.

The table lists each NIC and NIC group.

Figure 250 NIC and NIC Group Configuration tab toolbar



Use the toolbar to perform these actions:

Table 53 NIC and NIC Group Configuration tab toolbar

Icon	Action
New NIC Group	Create a NIC group by clicking New NIC Group and selecting and ordering two available NICs in the pop-up window. No more than two NICs can be part of a NIC group. A group can contain a single NIC.
Modify	Modify a NIC or NIC group by selecting it and then clicking Modify . Modify the settings in the pop-up window.
Swap Parameters	Switch the configuration settings between two NICs, two NIC groups, or a NIC and a NIC group by selecting both items in the list (press and hold the Ctrl key while selecting) and clicking Swap Parameters . Swapping parameters changes the IP address, burbs, aliases, and other configured attributes associated with the NIC or NIC group, and different rules are applied. If you swap a NIC and a NIC group, the interface that used the single NIC will now use the NIC group. Caution: Swapping NIC or NIC group parameters after you have initially configured your firewall could have unexpected results.
Delete	Delete a NIC group by selecting it and clicking Delete . <ul style="list-style-type: none"> You cannot delete a NIC group that is referenced by an interface. The firewall automatically detects NICs. To delete an individual NIC, you must physically remove it.
Usage	To view how a NIC or NIC group is being referenced, select it in the list and click Usage .
Find	Search for a specific element(s) in the list using the Find field. Type your search criteria, and NIC and NIC groups with matching elements will appear in the list. Clear this field to see the full list again.
Show Status	To view the status of interfaces and their associated NICs, click Show Status. You can also restart NICs and ping addresses in the pop-up window.

Creating interfaces

The internal and external network interfaces of the Sidewinder are defined during the initial configuration. These interfaces have IPv4 addresses.

An interface can have IPv4 addresses, IPv6 addresses, or both.

By using VLANs, you can create up to:

- 512 interfaces on a standalone firewall.
- 255 interfaces on a High Availability cluster.

You can configure the following types of interfaces and interface elements:

- [*Create a standard interface*](#)
- [*Create a VLAN interface*](#)
- [*Create a DHCP interface*](#)
- [*Create a transparent interface*](#)
- [*Enter an IPv4 address and aliases on an interface*](#)
- [*Enter an IPv6 address on an interface*](#)
- [*Configure redundant NICs*](#)
- [*Create or modify a High Availability interface*](#)
- [*Restart a NIC*](#)
- [*Send a ping to test connectivity*](#)

Create a standard interface

Figure 251 Standard interface properties

Interface Configuration: Interface Properties

Interface name: ☐ Enable interface

Description:

NIC or NIC Group: ☐ VLAN id: (1 - 4094)

MTU Size (Bytes): ☒ Standard (1500) ☐ Jumbo (9000) ☐ Custom (576 - 1500)

Address Configuration: ☐ Obtain an IPv4 address automatically via DHCP ☐ Enable IPv6 on this interface

IPv4 addresses:

Type	Address/Mask
primary	x.x.x.x/24

IPv6 addresses:

Address/Mask

Quality of Service Profile: ☐ QoS profiles only limit the bandwidth of traffic leaving an interface. QoS profiles are not supported on VLANs.

Interface id: ☐ The interface id only applies to interfaces where IPv6 is enabled.

OK Cancel Help

- 1 Select **Network > Interfaces**. The Interfaces window appears.
- 2 On the Interface Configuration tab, click **New**. The Interface Properties window appears.
- 3 Enter a name and description for the interface. The name can contain alphanumeric characters, dashes (-), underscores (_), and spaces ().
- 4 From the **NIC or NIC Group** drop-down list, select the NIC that will be associated with this interface.
You can click **Modify NIC or NIC group** to make changes to the selected NIC's hardware properties.
- 5 Select the size of the Maximum Transmission Unit (MTU) for *outgoing* packets:
 - **Standard (1500)** – Select this option to use the standard MTU.
 - **Jumbo (9000)** – Select this option to allow jumbo frames. This option is only available on NICs that support jumbo frames.
 - **Custom (576–9000 for IPv4/1280–9000 for IPv6)** – Select this option if you need to specify a custom MTU. If the NIC does not support jumbo frames, the range for this option will be 576–1500.
- 6 From the **Burb** drop-down list, select the burb that the interface is in. You can click **New burb** to create a new burb.

- 7 Enter the appropriate IP addresses and aliases to be associated with this interface. You can enter an IPv4 address and aliases, IPv6 addresses, or both.

See the following for details:

- [Enter an IPv4 address and aliases on an interface](#)
- [Enter an IPv6 address on an interface](#)

- 8 [Optional] From the **Quality of Service** drop-down list, select a Quality of Service profile to allocate available bandwidth to traffic leaving this interface.
- 9 Click **OK** and save your changes.

Create a VLAN interface

A VLAN is a virtual interface that allows administrators to segment a LAN into different broadcast domains regardless of the physical location.

- VLANs might not work on some older NICs.
- You must use a network switch or router that can decipher VLAN traffic to use VLANs.
- VLANs are supported in a High Availability (HA) configuration. For best results, configure VLANs before configuring HA.
- To filter traffic for a VLAN, use the following syntax:
 - For a NIC – `tcpdump -pni nic vlan vlanID`
 - For a NIC group – `tcpdump -pni nic_group vlan vlanID`

Figure 252 VLAN interface properties

The screenshot shows the 'Interface Configuration: Interface Properties' dialog box. The 'Interface name' field is set to 'vlan_network' and the 'Enable interface' checkbox is checked. The 'Description' field is set to 'Internal VLAN interface'. The 'NIC or NIC Group' dropdown is set to 'em2'. The 'VLAN id' is set to '2' (range 1 - 4094). The 'MTU Size (Bytes)' section has 'Standard (1500)' selected. The 'Address Configuration' section has 'Burb' set to 'internal'. The 'Obtain an IPv4 address automatically via DHCP' checkbox is unchecked. The 'IPv4 addresses' table has one entry: 'primary' with 'x.x.x.x/24'. The 'IPv6 addresses' section is empty. The 'Quality of Service Profile' dropdown is set to '<None>'. The 'Interface id' field is empty. There are two information icons: one for QoS profiles and one for the interface id.

Interface Configuration: Interface Properties

Interface name: ☒ Enable interface

Description:

NIC or NIC Group: ☒ VLAN id: (1 - 4094)

MTU Size (Bytes): ☒ Standard (1500) ☐ Jumbo (9000) ☐ Custom (576 - 9000)

Address Configuration: ☐ Obtain an IPv4 address automatically via DHCP

IPv4 addresses:

Type	Address/Mask
primary	x.x.x.x/24

IPv6 addresses:

Quality of Service Profile:

Interface id:

OK Cancel Help

To create a VLAN interface:

- 1** Select **Network > Interfaces**. The Interfaces window appears.
- 2** On the Interface Configuration tab, click **New**. The Interface Properties window appears.
- 3** Enter a name and description for the interface. The name can contain alphanumeric characters, dashes (-), underscores (_), and spaces ().
- 4** From the **NIC or NIC Group** drop-down list, select the NIC that will be associated with this interface.
You can click **Modify NIC or NIC group** to make changes to the selected NIC's hardware properties.
- 5** Select **VLAN id**.
- 6** In the **VLAN id** field, specify a numeric ID for this VLAN.
 - Valid values are 2–4094. (1 is reserved for special configurations.)
 - VLAN IDs must be unique across all VLAN interfaces tied to the same physical NIC or NIC group.
- 7** Select the size of the Maximum Transmission Unit (MTU) for *outgoing* packets:
 - **Standard (1500)** – Select this option to use the standard MTU.
 - **Jumbo (9000)** – Select this option to allow jumbo frames. This option is only available on NICs that support jumbo frames.
 - **Custom (576–9000 for IPv4/1280–9000 for IPv6)** – Select this option if you need to specify a custom MTU. If the NIC does not support jumbo frames, the range for this option will be 576–1500.
- 8** From the **Burb** drop-down list, select the burb that the interface is in. You can click **New burb** to create a new burb.
- 9** Enter the appropriate IP addresses and aliases to be associated with this interface. You can enter an IPv4 address and aliases, IPv6 addresses, or both.
See the following for details:
 - [Enter an IPv4 address and aliases on an interface](#)
 - [Enter an IPv6 address on an interface](#)
- 10** Click **OK** and save your changes.

Create a DHCP interface

An interface using Dynamic Host Configuration Protocol (DHCP) allows you to centrally manage IP addresses within your network.

Note that:

- Only one DHCP interface can be enabled at a time.
- You can enter IPv6 addresses on an interface that is using DHCP for IPv4 addresses.
- DHCP interfaces are not allowed on an HA cluster.

Figure 253 DHCP interface properties

Interface Configuration: Interface Properties

Interface name: ☒ Enable interface

Description:

NIC or NIC Group

☐ VLAN id: (1 - 4094)

MTU Size (Bytes)

☒ Standard (1500)
☐ Jumbo (9000)
☐ Custom (576 - 9000)

Address Configuration

Burb:

☒ Obtain an IPv4 address automatically via DHCP

☐ Enable IPv6 on this interface

IPv6 stateless auto address configuration:
☒ Static ☐ Host mode ☐ Router mode

IPv4 addresses:

Type	Address/Mask
primary	x.x.x.x/24

IPv6 addresses:

Address/Mask

Quality of Service Profile

QoS profiles only limit the bandwidth of traffic leaving an interface. QoS profiles are not supported on VLANs.

Interface id:

The interface id only applies to interfaces where IPv6 is enabled.

OK Cancel Help

To create a DHCP interface:

- 1 Select **Network > Interfaces**. The Interfaces window appears.
- 2 On the Interface Configuration tab, click **New**. The Interface Properties window appears.
- 3 Enter a name and description for the interface. The name can contain alphanumeric characters, dashes (-), and underscores (_).
- 4 From the **NIC or NIC Group** drop-down list, select the NIC that will be associated with this interface.
You can click **Modify NIC or NIC group** to make changes to the selected NIC's hardware properties.

Note: You can use redundant NICs for a DHCP interface. See [Configure redundant NICs](#) for information.

- 5 Select the size of the Maximum Transmission Unit (MTU) for *outgoing* packets:
 - **Standard (1500)** – Select this option to use the standard MTU.
 - **Jumbo (9000)** – Select this option to allow jumbo frames. This option is only available on NICs that support jumbo frames.
 - **Custom (576–9000 for IPv4/1280–9000 for IPv6)** – Select this option if you need to specify a custom MTU. If the NIC does not support jumbo frames, the range for this option will be 576–1500.
- 6 In the **Address Covariation** area, select **Obtain an IPv4 address automatically** via DHCP.
 - The internet burb is automatically selected in the **Burb** field and cannot be modified.
 - The IPv4 addresses area is disabled. You cannot add an IPv4 address or aliases.
- 7 [Optional] Enter the appropriate IPv6 addresses. See [Enter an IPv6 address on an interface](#) for details.
- 8 [Optional] From the **Quality of Service** drop-down list, select a Quality of Service profile to allocate available bandwidth to traffic leaving this interface.
- 9 Click **OK** and save your changes.

Create a transparent interface

A transparent interface is made up of two bridged interfaces. You can use a transparent interface to separate a single network into two burbs. This allows you to enforce security policy on traffic that passes through your firewall's transparent interface without re-addressing the network around the firewall. For more information, see [Sidewinder deployment options on page 23](#).

[Table 54](#) shows the default Sidewinder interface configuration. These interfaces, or any other two interfaces, can be used to configure a transparent interface.

Table 54 Standard interfaces

User defined interface name	NIC or NIC Group	Burb name
external_network	em0	external
internal_network	em1	internal

[Table 55](#) shows a transparent interface configured using the default interfaces. Note that bridge0 is made up of em0 and em1.

Table 55 Transparent interface

User defined transparent interface name	NIC or NIC Group
bridged_network	bridge0 (em0, em1)

If you configure a transparent interface, you cannot enable or configure:

- Split DNS
- High Availability
- Sendmail
- Dynamic routing
- DHCP on the transparent interface
- DHCP Relay agent
- VPN termination in a transparent burb
- IPv6 addresses on the transparent interface

Note: A transparent interface passes traffic at layer two, similar to a bridge. Because Sidewinder does not run the Spanning Tree bridging protocol, enabling Spanning Tree on the switch that is connected to the firewall is not recommended.

To create a transparent interface:

- 1 Select **Network > Interfaces**. The Interfaces window appears.
- 2 On the Interface Configuration tab, click **New > Transparent Interface**. The Interface Properties window appears.

Figure 254 Transparent interface properties

Interface Configuration: Interface Properties (Transparent)

Interface name: ☒ Enable interface

Description:

Bridged interfaces

Find:

Use	Interface	Burb	VLAN ID	NIC or NIC Group
<input type="checkbox"/>	external_network	external	--	em0
<input type="checkbox"/>	internal_network	internal	--	em1
<input type="checkbox"/>	<Not in use>	--	--	em2

MTU Size (Bytes)

☒ Standard (1500)
☐ Jumbo (9000)
☐ Custom (576 - 9000)

Address Configuration

Burb:

☐ Obtain an IPv4 address automatically via DHCP

IPv4 addresses:

Type	Address/Mask
primary	x.x.x.x/24

IPv6 stateless auto address configuration

☒ Static ☐ Host mode ☐ Router mode

IPv6 addresses:

Address/Mask

Quality of Service Profile

QoS profiles only limit the bandwidth of traffic leaving an interface. QoS profiles are not supported on VLANs.

Interface id:

The interface id only applies to interfaces where IPv6 is enabled.

- 3 Enter a name and description for the interface. The name can contain alphanumeric characters, dashes (-), underscores (_), and spaces ().

- 4** In the **Bridged interfaces** area, select the two interfaces that will be members of this transparent interface. Note that:
- The member interfaces can be standard, VLAN, or redundant interfaces.
 - If the member interfaces have IP addresses assigned to them, these addresses will be removed when the transparent interface is created.
 - Before being added to a transparent interface, the member interfaces must be:
 - Assigned a name
 - Associated with a NIC or NIC group
 - Assigned to a unique burb

If either or both of the interfaces that you want to bridge are not yet configured, do so now by clicking **New** in the Bridged interfaces toolbar. See [Creating interfaces](#) for more information.

- 5** Select the size of the Maximum Transmission Unit (MTU) for *outgoing* packets:
- **Standard (1500)** – Select this option to use the standard MTU.
 - **Jumbo (9000)** – Select this option to allow jumbo frames. This option is only available on NICs that support jumbo frames.
 - **Custom (576–9000 for IPv4/1280–9000 for IPv6)** – Select this option if you need to specify a custom MTU. If the NIC does not support jumbo frames, the range for this option will be 576–1500.
- 6** In the **IPv4 addresses** area, enter the appropriate IP addresses and aliases to be associated with this interface. For details, see [Enter an IPv4 address and aliases on an interface](#).
- 7** Click **OK** and save your changes.

Enter an IPv4 address and aliases on an interface

The first IP address in the **IPv4 addresses** area is labeled as the primary address for this interface. All subsequent addresses added to this interface are aliases.

Figure 255 IPv4 addresses section of the Interface Properties window

Interface Configuration: Interface Properties

Interface name: ☒ Enable interface

Description:

NIC or NIC Group:

☐ VLAN id: (1 - 4094)

MTU Size (Bytes): ☒ Standard (1500) ☐ Jumbo (9000) ☐ Custom (576 - 9000)

Address Configuration:

☐ Obtain an IPv4 address automatically via DHCP

☐ Enable IPv6 on this interface

IPv6 stateless auto address configuration: ☒ Static ☐ Host mode ☐ Router mode

IPv4 addresses:

Type	Address/Mask
primary	192.168.0.1/24
alias	192.168.0.20/24

IPv6 addresses:

Address/Mask

Quality of Service Profile:

QoS profiles only limit the bandwidth of traffic leaving an interface. QoS profiles are not supported on VLANs.

Interface id:

The interface id only applies to interfaces where IPv6 is enabled.

1 Enter an IPv4 address:

- Click the **x.x.x.x/24** field and type an IPv4 address that will be associated with this interface.
- [Optional] Modify the network mask. Valid values are 0–32. The network mask is used to identify the significant portion of the IP address.
- Press **Enter**.

Note: To delete an IPv4 address, select it in the Address/Mask list and click **Delete**.

2 Enter alias IP addresses.

The first IP address in the **IPv4 addresses** area is the primary IPv4 address for this interface. All subsequent addresses are aliases.

Alias IP addresses are used in Multiple Address Translation (MAT). Adding alias IP addresses to a network interface can be used for purposes such as the following:

- Specific logical networks connected to one interface can be consistently mapped to specific IP aliases on another interface when using address hiding.
- The interface can accept connection requests for any defined alias.
- The interface can communicate with more than one logical network without the need for a router.
- The interface can have more than one address on the same network and have DNS resolve different domains to each host address.

To enter an alias IP address:

- a** In the **IPv4 addresses** area, click **New**.
- b** Click the **x.x.x.x/24** field and type an alias IP address that will be associated with this interface IP address.
- c** [Optional] Modify the network mask. The network mask is used to identify the significant portion of the IP address.
- d** Press **Enter**.

Note: To delete an alias IP address, select it in the Address/Mask list and click **Delete**.

3 Order the addresses.

- You can change the order of the addresses in the **IPv4 addresses** area. The first address in the list is selected as the outgoing address when sending data.
- You can swap the primary address and alias addresses and you can change the order of the aliases. The top address in the list is labeled as the primary address.

To change address locations in the list, select an address and click the **Move up** and **Move down** arrows.

Enter an IPv6 address on an interface

You should understand IPv6 addresses and how they are implemented on Sidewinder before enabling them on your firewall.

- For information on IPv6 addressing, see the following resources:
 - RFC 2460 [IPv6 Specification](#)
 - RFC 4193 [Unique Local IPv6 Unicast Addresses](#)
 - RFC 4291 [IPv6 Addressing Architecture](#)
 - *IPv6 Essentials*, 2nd ed. by Silvia Hagen (O'Reilly)
- For information about the IPv6 implementation on Sidewinder, see [About IPv6 addresses on Sidewinder](#).

IPv6 must be enabled on your Sidewinder before you can pass IPv6 traffic. After IPv6 is enabled, it cannot be reversed except by restoring a previous configuration backup. A configuration backup is automatically created when you enable IPv6.

When IPv6 is enabled, two default network objects are available for rules to distinguish between endpoints of **<Any>**: **<Any V4>** and **<Any V6>**. The first time IPv6 is enabled on your firewall, a wizard will prompt you to select how current rules with a source or destination endpoint of **<Any>** will be handled.

- You can choose to convert existing **<Any>** rules to **<Any V4>**. The source or destination endpoint of **<Any>** on all current rules will be changed to **<Any V4>**. Only IPv4 traffic will match the **<Any V4>** endpoint.
- You can choose to leave the **<Any>** rules as they are. The source or destination endpoint of **<Any>** on all current rules will remain **<Any>**. IPv4 and IPv6 traffic will match the **<Any>** endpoint. Some Sidewinder facilities do not currently support IPv6 but may in the future, so IPv6 traffic will match those rules when they are supported.

See the procedure below for instructions on enabling IPv6 and entering IPv6 addresses.

Figure 256 IPv6 addresses section of the Interface Properties window

Interface Configuration: Interface Properties

Interface name: ☐ Enable interface

Description:

NIC or NIC Group:

☐ VLAN id: (1 - 4094)

MTU Size (Bytes):
☒ Standard (1500)
☐ Jumbo (9000)
☐ Custom (1280 - 1500)

Address Configuration

Burb:

☐ Obtain an IPv4 address automatically via DHCP

IPv4 addresses:

Type	Address/Mask
primary	x.x.x.x/24

IPv6 addresses:

Address/Mask
::xxxx/64

Quality of Service Profile:

☒ Enable IPv6 on this interface

IPv6 stateless auto address configuration:
☒ Static ☐ Host mode ☐ Router mode

Interface id:

To enable IPv6 and enter addresses:

- 1 In the **IPv6 addresses** area, select the **Enable IPv6 on this interface** check box.

Note: When IPv6 is first enabled on a firewall, a wizard will prompt you to select how current rules with a source or destination endpoint of **<Any>** will be handled. Follow the wizard instructions to make your selection.

- 2 Select a stateless auto-address configuration.

Static configuration is the most suitable configuration for most firewalls. Host mode and router mode should be used only if you want to use autoconfiguration. Using these modes can cause unexpected results, for example:

- A firewall with an interface configured in host mode can automatically add new IPv6 addresses to the interface that the user might not expect.
- A firewall with an interface configured in router mode with static IPv6 addresses can, if the *rtadvd.conf* file is not modified, advertise prefixes derived from the static IPv6 addresses. This can result in unexpected addresses being added to IPv6 devices in the same network operating in host mode.

Make a selection:

- **Static** – The interface is assigned the link-local address plus any static addresses you enter. The link-local address is automatically created whenever an interface becomes enabled.
- **Host mode** – The interface is assigned the link-local address plus any static addresses you enter. It is also assigned autoconfigured addresses derived by combining any prefixes received in router advertisements with the interface ID.
- **Router mode** – The interface is assigned the link-local address plus any static addresses you enter. The firewall sends out router advertisements either with prefixes in the *rtadvd.conf* file or with prefixes derived from the static addresses on the interface.

- 3 Enter an IPv6 address:

- a In the **IPv6 addresses** area, click **New**.
- b Click the **xxxx** field and type an IPv6 address that will be associated with this interface.
- c [Optional] Modify the mask length. Valid values are 0–128.
- d Press **Enter**.

Note: To delete an IP address, select it in the Address/Mask list and click **Delete**.

- 4 Order the addresses.

You can change the order of the addresses in the **IPv6 addresses** area.

To change address locations in the list, select an address and click the **Move up** and **Move down** arrows.

- 5 [Optional] Modify the interface ID.

The 16-hexadecimal ID in the **Interface id** field is automatically created. By default it is derived from the NIC or NIC group's MAC address and is used to generate the link-local address for the interface.

Note: Create a default route that forwards IPv6 traffic with no known route to its destination address. See [Configuring static routes](#) for information.

Configure redundant NICs

The redundant NIC function is configured using NIC groups. A NIC group contains two NICs: a primary and a standby. If the primary NIC loses its link or is disconnected, the standby NIC becomes the primary and starts passing traffic.

- The Sidewinder verifies a link at the physical layer (layer 1). The Sidewinder inspects the carrier detect status on the primary NIC in the NIC group.
 - If the link is active, the primary NIC is used to pass traffic.
 - If the link is not active, a failover event occurs and the standby NIC starts passing traffic.

When the link for the primary NIC is active again, a failback event automatically occurs and the primary NIC starts passing traffic.

- The firewall does not verify communication at the network layer with the next device. A failure in this part of the connection does not trigger a failover event.
- There can be a delay before the standby NIC starts passing traffic while the switch or router recognizes the change and selects the appropriate port.
- The NIC group uses the MAC address of the primary NIC no matter which NIC is actively passing traffic. The MAC address is used for communication at the data-link layer.

To create a NIC group:

- No more than two NICs can be part of a NIC group. A group can contain a single NIC.
- Both NICs must have the same media capabilities enabled.
- A NIC can be a member of multiple NIC groups, but it can be referenced by only one enabled interface at a time.

1 Open the NIC Group Properties window in either of these ways:

- In the NIC or NIC Group area of the Interface Configuration tab, click **Create new NIC group**.
- On the NIC and NIC Group Configuration tab, click **New NIC Group**.

A group name is automatically assigned. You cannot change a NIC group name.

2 [Optional] In the **Description** field, type a description of the NIC group to further identify it. This entry appears in the Description column of the NIC and NIC Group Configuration table.

3 Select NICs for the NIC group:

- a** In the left pane list of available NICs, select the appropriate NIC. You can select two NICs by pressing and holding the **Ctrl** key while selecting the NICs.
- b** Click the right arrow to move the select NIC(s) to the right pane. The NICs in this pane are members of the NIC group.

4 Use the up and down arrows to order the selected NICs. The first NIC in the list is the primary.

5 Click **OK**.

You can use redundant NICs on a standard, VLAN, DHCP, or transparent interface. Select the NIC group when you create the interface.

Create or modify a High Availability interface

Note that:

- If you make any configuration changes to an HA cluster interface, both cluster firewalls must be restarted. See [Restarting an HA cluster](#).
- You cannot use IPv6 addresses in an HA cluster.

If a firewall is part of an HA cluster, the Interface Properties window has two IP addresses:

- **Cluster IP address** – This is the IP address of the interface in the HA cluster. This address also appears on the High Availability: Common Parameters window. If you modify the cluster IP address in one window, it is automatically updated in the other window.

Three interfaces in each HA cluster member should have a cluster IP address: the heartbeat burb and one interface for each cluster member. You can create additional interfaces for private or management purposes without assigning a cluster IP addresses to them.

- **Primary IP address** – This is the IP address of the interface before joining the HA cluster.

If you modify any of these attributes in an interface, that same modification is automatically made in the corresponding interface of the other cluster member:

- Burb
- Quality of Service profile
- Alias address
- MTU

Figure 257 HA interface properties

Interface Configuration: Interface Properties

Interface name: ☒ Enable interface

Description:

NIC or NIC Group: ☐ VLAN id: (1 - 4094)

MTU Size (Bytes): ☒ Standard (1500) ☐ Jumbo (9000) ☐ Custom (576 - 9000)

Address Configuration: ☐ Obtain an IPv4 address automatically via DHCP

IPv4 addresses:

Type	Shared	Address/Mask
cluster	<input checked="" type="checkbox"/>	10.65.248.20/24
primary	<input type="checkbox"/>	10.65.248.2/24

IPv6 addresses:

Quality of Service Profile: ☐ QoS profiles only limit the bandwidth of traffic leaving an interface. QoS profiles are not supported on VLANs.

Interface id:

The interface id only applies to interfaces where IPv6 is enabled.

OK Cancel Help

Restart a NIC

Perform this procedure to restart a down NIC.


- 1 Select **Network > Interfaces**. The Interfaces window appears.
- 2 On the Interface Configuration tab or the NIC and NIC Group Configuration tab, click **Show Status**. The Interface and NIC Status window appears.
- 3 In the table, select the down NIC.
- 4 Click **Restart**.

Send a ping test connectivity

Perform this procedure to test connectivity for an interface or NIC.

- 1 Select **Network > Interfaces**. The Interfaces window appears.
- 2 On the Interface Configuration tab or the NIC and NIC Group Configuration tab, click **Show Status**. The Interface and NIC Status window appears.
- 3 Click **Ping**. The Ping Test window appears.

Figure 258 Ping Test window



- 4 In the **IP address to ping** field, enter an IP address or fully qualified domain name that the ping will be sent to. To find the IP address for a host name, type the name and click **DNS Lookup**.
- 5 [Optional] In the **Commandline flags** field, enter command line parameters for the ping test. You can enter alphanumeric characters, dashes (-), and underscores (_).
- 6 Click **Start Ping**. The button changes to **Stop Ping** and the ping results appear in the window.
- 7 Click **Stop Ping** to stop the test.
- 8 Click **Close**.

Configuring Quality of Service

Quality of Service (QoS) guarantees a certain level of performance for a data flow by using different priorities and queuing mechanisms to allocate available bandwidth. QoS is beneficial for networks with limited bandwidth that must pass latency-sensitive or bandwidth-intensive traffic.

From the Quality of Service window, you can create QoS profiles that can be applied to the network interfaces of the Sidewinder. Each QoS profile contains one or more queues that allow you to prioritize network performance based on network traffic type. All queues are assigned a priority value, allocated a percentage of available bandwidth, and can be allowed to borrow bandwidth from other queues. When a queue is full, any additional packets matching that queue are dropped. Queues are applied to network traffic based on the services that are selected.

When QoS policy is applied to a network interface, only outgoing traffic on that interface is controlled by QoS—packets arriving on that interface are not affected. If you require traffic for a particular service to be controlled in both directions, that service must be present in the QoS policy of both interfaces where traffic for that service leaves the firewall. Consider the following QoS configurations and their effect on a connection between an internal client and external web server:

- The external interface's QoS profile includes HTTP – Traffic sent from the internal client to the external web server is affected by QoS.
- The internal interface's QoS profile includes HTTP – Traffic sent from the web server to the internal client is affected by QoS.
- Both internal and external interface QoS profiles include HTTP – All traffic between the client and web server is affected by QoS.

QoS is applied to network traffic at the IP and transport layers based on the service(s) selected in each queue. Protocols that use dynamic ports negotiated at the application layer like FTP or VoIP will not match QoS queues using those services, since QoS does not examine the application layer when processing packets.

Consider the case in which a QoS queue has been created with the FTP proxy service selected. QoS is applied to the control connection (tcp port 21) but not the data connection (high random tcp port or tcp port 20). Since the control connection is made on the port defined in the service, QoS policy is applied to it. However, QoS is not applied to the data connection because it is made on a port negotiated at the application layer between the client and server.

Note: To apply QoS to protocols that employ dynamic ports, create a service that includes the range of dynamic ports, and select this service on the QoS queue.

To apply QoS to a network interface:

- 1 Create a QoS profile.
- 2 Add QoS queues to the profile.
- 3 Apply the QoS profile to a network interface under **Network > Interfaces**.

Note: QoS cannot be configured on VLANs.

Select **Network > Quality of Service**.

The Quality of Service window consists of two panes:

- Profiles (upper pane) – Use this pane to configure QoS profiles.
- Queues (lower pane) – Use this pane to configure QoS queues for the profile selected in the Profiles pane.

Also see [Example QoS scenarios](#).

Figure 259 Quality of Service window

The screenshot shows the Quality of Service window with two main panes. The top pane, titled 'Profiles', contains a table with two columns: 'Profile Name' and 'Description'. It lists two profiles: 'inside' and 'outside'. The bottom pane, titled 'Queues', contains a table with five columns: 'Queue Name', 'Priority', 'Allocated Bandwidth', 'Services', and 'Can Borrow'. It lists three queues: 'default', 'mail', and 'web'. The 'Profile name' field is set to 'inside' and the 'Remaining bandwidth' is 0%.

Profile Name	Description
inside	
outside	

Profile name: Remaining bandwidth:

Description:

Queue Name	Priority	Allocated Bandwidth	Services	Can Borrow
default	0	20%		Yes
mail	0	30%	imap, pop, smtp	No
web	0	50%	http, https	Yes


Configuring QoS Profiles

QoS profiles contain QoS policy that can be assigned to a particular network interface. They behave as containers for QoS queues that make up the QoS policy.

Each profile contains a default queue that cannot be deleted or renamed. The default queue processes all packets that do not match any queues you have explicitly defined.

Use the toolbar to perform the actions described in this section.

Figure 260 Quality of Service profile toolbar



Use the toolbar to perform these actions.

Table 56 QoS profile toolbar

Button	Action
New	Click New to create a new profile. The profile name must be seven characters or less.
Modify	To modify a profile, select it in the Profile pane and configure its attributes in the Queues pane (alternately, select the profile, then click Modify). See "Configuring QoS queues" below for more information.
Delete	Select the profile, then click Delete .
Duplicate	Click Duplicate to create a copy of an existing profile. Type a name and [Optional] description in the Modify Profile window that appears.
Rename	To change the name of an existing profile, click Rename . Enter the new name, then click OK .
Simulate Bandwidth Allocation	Click Simulate Bandwidth Allocation to simulate various settings for the profile. For complete information, see About the Bandwidth Allocation Simulator window .
Usage	To show which network interfaces are using a profile, click Usage .
Search	To find a profile, enter all or part of the name. When the system finds a match, it appears highlighted in the pane. If the system does not find a match, the pane appears blank. Use the Backspace key to find partial matches or delete the search term to return to the main window.
Load QoS Policy	Click Load QoS Policy to reapply the policy. You may want to do this if the Quality of Service does not seem to be performing as expected. Use the <code>cf qos show</code> command to determine if the bandwidth information is correct.
QoS Status	Click QoS Status to view QoS filter rules, and queue statistics. Statistics are reset when any QoS policy change is made. Queue names are presented in the format <code>queuename_profilename</code> .

Configuring QoS Queues

Use QoS queues to allocate available bandwidth based on traffic type. Queues make up the policy in QoS profiles—each queue in a profile is assigned a priority value and dedicated a percentage of available bandwidth.

Figure 261 QoS Queues pane

Profile name: Remaining bandwidth:

Description:

| ABC | Find:

Queue Name	Priority	Allocated Bandwidth	Services	Can Borrow
default	0	20%		Yes
mail	0	30%	imap, pop, smtp	No
web	0	50%	http, https	Yes

To create QoS policy, select the profile you want to modify in the profile pane, then use the Queue pane to make policy changes.

To prioritize bandwidth usage within a profile, configure the following attributes of each queue in the profile:

- **Priority** – A value between 0–7 (lowest–highest) that determines the order the queue is processed relative to the other queues in the profile. Higher priority queues are processed first, resulting in lower latency for them.
- **Allocated Bandwidth** – The percentage of available bandwidth to be dedicated to the queue. The available bandwidth for a QoS profile is determined by the link speed of the network interface it is associated with.
- **Services** – The types of traffic the queue applies to.
- **Can Borrow** – If enabled, allows the queue to borrow bandwidth from the other queues in the profile when it exhausts its allocated bandwidth.

Each profile contains a default queue that cannot be deleted or renamed. The default queue processes all packets that do not match any queues you have explicitly defined. Edit the **Priority**, **Bandwidth**, and **Can Borrow** attributes of the default queue to control how QoS allocates bandwidth for services that are not included in custom queues.

Figure 262 Quality of Service queue toolbar





Table 57 QoS queue toolbar

Button	Action
New	Click New to create a new queue. See About the New/Modify Queue window .
Modify	Double-click the queue you want to change (alternately, select the queue, then click Modify). See About the New/Modify Queue window .
Delete	Select the queue, then click Delete .
Rename	To change the name of an existing queue, click Rename . Enter the new name, then click OK .
Update bandwidth	Click Update Bandwidth to change the queue settings. See About the Adjust Bandwidth window .
Find	To find a queue, enter all or part of the name. When the system finds a match, it appears highlighted in the pane. If the system does not find a match, the pane appears blank. Use the Backspace key to find partial matches or delete the search term to return to the main window.

About the Bandwidth Allocation Simulator window

Use the this window to simulate how QoS will allocate bandwidth based on the QoS queues you have defined for a profile.

Figure 263 Bandwidth Allocation Simulator window



When the Bandwidth Allocation Simulator window opens, the simulated demand for each queue defaults to the percentage of bandwidth allocated to it.

To change the simulated demand for a queue:

- 1 Double-click the value in the Simulated Demand column.
- 2 Type the percentage demand you would like to simulate.
- 3 Click outside the Simulated Demand column.

To simulate the worst case scenario in which each queue is at 100% demand, click **Worst Case**. For examples, see [Example QoS scenarios](#).

About the New/Modify Queue window

Use this window to create new queues or modify existing queues.

Figure 264 New/Modify Queue window

Name:

Description:

Priority: (0-7) Bandwidth: 0% remaining ☒ Can borrow

7 is the highest priority and 0 is the lowest.

Available Services: Find:

Use	Name	Agent	Ports	Summary	Description
<input type="checkbox"/>	netprobe netBIOS	TCP/UDP Packe	137/udp	Bi-directional=Off, Statef	
<input type="checkbox"/>	aol	Generic Proxy	5190/tcp	Fast path=On, TCP idle ti	aol proxy
<input type="checkbox"/>	dns	DNS Proxy	53 (tcp and udp)	TCP idle timeout=3600, U	DNS proxy
<input type="checkbox"/>	finger	Generic Proxy	79/tcp	Fast path=On, TCP idle ti	finger proxy
<input type="checkbox"/>	ftp	FTP Proxy	21/tcp	Connection type=transpe	FTP proxy
<input type="checkbox"/>	fwregisrtp	Cluster Registr	9010/tcp	TCP idle timeout=7200	Cluster Regis
<input type="checkbox"/>	gopher	Generic Proxy	70/tcp	Fast path=On, TCP idle ti	gopher proxy
<input type="checkbox"/>	h323	H323 Proxy	1720/tcp, 1719/udp	TCP idle timeout=7200, U	H323 proxy
<input type="checkbox"/>	http	HTTP Proxy	80/tcp	Connection type=transpe	HTTP proxy
<input type="checkbox"/>	https	HTTPS Proxy	443/tcp	Connection type=transpe	HTTPS proxy
<input type="checkbox"/>	ica	Citrix Proxy	1494/tcp	Fast path=On, TCP idle ti	ICA proxy
<input type="checkbox"/>	ident	Generic Proxy	113/tcp	Fast path=On, TCP idle ti	ident proxy
<input type="checkbox"/>	iiop	IIOP Proxy	683/tcp	Fast path=Off, TCP idle ti	IIOP proxy
<input type="checkbox"/>	imap	Generic Proxy	143/tcp	Fast path=On, TCP idle ti	imap proxy
<input type="checkbox"/>	irc	Generic Proxy	6667/tcp	Fast path=On, TCP idle ti	irc proxy
<input type="checkbox"/>	ironmail-admin	HTTPS Proxy	10443/tcp	Connection type=transpe	IronMail web
<input type="checkbox"/>	ironmail-support	Generic Proxv	20022/tcn	Fast path=On, TCP idle ti	IronMail soft

New View Select All Deselect All

OK Cancel Help

To create or modify a QoS queue:

- 1 In the **Name** field, type a name for the new queue (Use **Rename** on the queue toolbar to rename an existing queue).

Note: The queue name must be seven characters or less.

- 2 [Optional] In the **Description** field, type a more detailed description of the queue.

- 3 In the **Priority** field, type the priority value (0–7) for this queue.

- 4 In the **Bandwidth** field, type the percentage of bandwidth to be allocated for this queue. This value is limited to the amount of bandwidth not already allocated to other queues. Bandwidth cannot be set to 0.

- 5 To allow this queue to borrow bandwidth from the other queues, select the **Can borrow** box.

Note: The Can borrow option is selected by default. Unless you want to allow this queue to appropriate bandwidth from queues with equal or lower priority, disable this option.

- 6 In the **Available Services** pane, select the service(s) that you want to associate with this queue. If you want to select a service that is not listed, you can create a new one by clicking **New**.

Note: QoS queue policy is applied to packets that match the protocol and port of the selected service(s).

- 7 Click **OK** to finish configuring the queue.

Repeat this procedure for each additional queue you wish to add for this profile.

About the Adjust Bandwidth window

Use this window to adjust the Priority and Bandwidth for all queues in a profile from a central location.

Figure 265 Adjust Bandwidth window

New Profile: Adjust Bandwidth

Profile:

Total allocated bandwidth:

Queue Name	Priority	Bandwidth
default	0	20%
mail	0	30%
web	0	50%

To change the Priority or Allocated Bandwidth for a queue:

- 1 Double click the value in the **Priority** or **Bandwidth** column that you wish to change.
- 2 Type the desired value:
 - For **Priority**, type a value between 0–7 (lowest–highest).
 - For **Bandwidth**, type a value between 1–100 representing the percentage of bandwidth to allocate to this queue.
- 3 Click outside the modified cell.
- 4 Click **OK** and save your changes.

Note: The total allocated bandwidth cannot exceed 100%.

Example QoS scenarios

The interaction between multiple QoS queues with differing priorities, allocated bandwidth, and borrowing can be complex. Use the following example scenarios to familiarize yourself with QoS in practice.

In the examples below, two queues are configured—ssh and http. No other traffic is flowing, although other queues may be defined.

Case 1

SSH is allocated 10% of bandwidth at priority 7 with no borrowing allowed, and HTTP is allocated 10% of bandwidth at priority 7 with no borrowing allowed.

At congestion levels, exactly 10% of available bandwidth is allocated to each of the queues.

Case 2

SSH is allocated 10% of bandwidth at priority 0 with no borrowing allowed, and HTTP is allocated 10% of bandwidth at priority 7 with no borrowing allowed.

At congestion levels, exactly 10% of available bandwidth is allocated to each of the queues; however, HTTP traffic is processed before SSH traffic and hence experiences lower latency.

Case 3

SSH is allocated 30% of bandwidth at priority 7 with no borrowing allowed, and HTTP is allocated 10% of bandwidth at priority 7 with no borrowing allowed.

At congestion levels, exactly 30% of available bandwidth is allotted to the SSH queue with 10% going to the HTTP queue.

Case 4

SSH is allocated 30% of bandwidth at priority 7 with borrowing allowed, and HTTP is allocated 10% of bandwidth at priority 7 with borrowing allowed.

At congestion levels, a proportionally larger percentage of available bandwidth is allotted to the SSH queue, with the remaining traffic going to the HTTP queue. (Since SSH is allocated a larger portion of the bandwidth than HTTP, it gets more weight at the time of borrowing since they are of the same priority.)

Case 5

SSH is allocated 10% of bandwidth at priority 7 with borrowing allowed, and HTTP is allocated 10% of bandwidth at priority 7 with borrowing allowed.

At congestion levels, the two queues share the borrowed bandwidth equally (40% each).

Case 6

SSH is allocated 10% of bandwidth at priority 0 with borrowing allowed, and HTTP is allocated 10% of bandwidth at priority 7 with borrowing allowed.

At congestion levels, the HTTP queue commandeers all of the bandwidth since it is the highest priority queue and it is allowed to borrow.

Case 7

SSH is allocated 30% of bandwidth at priority 7 with no borrowing allowed, and HTTP is allocated 10% of bandwidth at priority 7 with borrowing allowed.

At congestion levels, the SSH queue uses 30% of available bandwidth and the HTTP queue commandeers all of the remaining bandwidth.

Summary

- If multiple queues have the same priority and borrowing is allowed, each queue borrows a percentage of available bandwidth. The amount of bandwidth each queue can borrow is determined by its allocated bandwidth in proportion to the allocated bandwidth of the other queues.
- If a queue with higher priority is allowed to borrow, it will starve lower priority queues, but not vice versa.
- If borrowing is not allowed, queues share available bandwidth per their allocated bandwidth value. Higher priority queues are serviced first, resulting in reduced latency for them at the expense of the lower priority queues.

17 Routing

Contents

[About routing on Sidewinder](#)

[Configuring static routes](#)

[RIP on Sidewinder](#)

[OSPF on Sidewinder](#)

[OSPF IPv6 on Sidewinder](#)

[BGP on Sidewinder](#)

[PIM-SM on Sidewinder](#)

[Dynamic routing in HA clusters](#)

[Troubleshooting dynamic routing issues](#)

About routing on Sidewinder

Traffic between machines on different networks or subnets requires routing. This routing information can be input manually using *static routes* and learned automatically using *dynamic routing*.

Each computer in your network also designates a specific route as its *default route*, to use when the computer cannot find an explicit route to the destination. This default gateway is generally a router that allows access to distant subnets. You can configure an alternate default route to act as a redundant route. If your primary default route becomes inaccessible, an alternate default route begins forwarding traffic.

The Forcepoint Sidewinder can participate in routing using information from static routes, and can act as a default gateway for your network.

The firewall supports four dynamic routing protocols:


- Routing Information Protocol (RIP)
- Open Shortest Path First (OSPF) protocol
- OSPF IPv6 protocol
- Border Gateway Protocol (BGP)
- Protocol Independent Multicast - Sparse Mode (PIM-SM) protocol

This chapter provides a brief overview of how each protocol works, and guidelines and scenarios for configuring the dynamic routing protocols and servers on the Sidewinder. The Sidewinder implementation of these protocols and their respective servers (ripd, ospfd, bgpd, and pimd) are based on the Quagga implementation. Any administrator planning on configuring RIP, OSPF, BGP, or PIM-SM on a Sidewinder is strongly encouraged to use the online help that is available when connected to a Sidewinder-hosted routing server using a command line interface, and the Quagga documentation available at <http://www.quagga.net/docs/quagga.pdf>.

Configuring static routes

Static routes are based on a fixed forwarding path. To create and modify static routes, select **Routing > Static Routing**. The Static Routing window appears.

Figure 266 Static Routing window



The screenshot shows the Static Routing window with a table of static routes. The table has five columns: Route Type, Destination, Netmask, Gateway, and Description. The first three rows are highlighted in blue. The first row is 'Primary Default' with destination 'default' and gateway '10.65.249.202'. The second row is 'Alternate Default' with destination 'backup' and gateway '<None>'. The third row is 'IPv6 Default' with destination 'v6_default' and gateway '<None>'. The rest of the table is empty.


Route Type ▲	Destination	Netmask	Gateway	Description
Primary Default	default		10.65.249.202	
Alternate Default	backup	<None>	<None>	
IPv6 Default	v6_default	<None>	<None>	

Use this window to manage default routes and to create and manage other static routes.

The table lists static routes configured on this firewall. Primary Default and Alternate Default appear automatically in the table. If IPv6 is enabled on the firewall, IPv6 Default also appears.

- The primary default route is created when you initially configure the firewall.
- The alternate default route is a placeholder. You must configure the alternate default route to enable default route failover.

Figure 267 Static Routing toolbar



Use the toolbar to perform these actions:

Table 58 Static routing toolbar

Icon	Action
New	Create a static route by clicking New and entering information in the Host/Network Route Properties pop-up window. See Configure other static routes for details.
Modify	<ul style="list-style-type: none"> Modify the default route by selecting Primary Default and clicking Modify. Modify the single static default route settings in the Default Route Properties pop-up window. See Configure default routes for details. Configure default route failover by selecting Primary Default or Alternate Default and clicking Modify. Configure the primary and alternate default routes in the Default Route Properties pop-up window. See Configure default routes for details. Modify an existing static route by selecting it and then clicking Modify. Modify the settings in the Host/Network Route Properties pop-up window. See Configure other static routes for details.
Delete	Delete a static route by selecting it and clicking Delete . Note: The values of the primary and alternate default routes are deleted. The placeholders for the default routes remain in the table.
Find	Search for a specific element(s) in the list using the Find field. Type your search criteria, and routes with matching elements will appear in the list. Clear this field to see the full list again.
Status	To view the status information of the routes configured for the firewall, click Status . You can also view route failover status and route failover audit, and you can reset the default route when it becomes accessible. See Check route status and reset the default route for more information.

Configure default routes

Use the Default Route Properties window to modify the default route or to configure an alternate route to use for default route failover.

- A default route, also known as the gateway of last resort, is the device the firewall sends traffic to if no other known route exists for the destination address.
- The alternate default route is a redundant route. If your primary default route becomes inaccessible, the firewall begins forwarding traffic to an alternate device.

Figure 268 Default Route Properties window: Single default route view

Static Routes: Default Route Properties

Current default route: 10.65.249.202

☒ Use single static default route

IP address: 10.65.249.202

Description:

☐ Use alternate default routes

Primary default route:

IP address:

Description:

Ping addresses:

IP Address

Ping interval: (seconds) Failures allowed:

Alternate default route:

IP address:

Description:

Ping addresses:

IP Address

Ping interval: (seconds) Failures allowed:

OK Cancel Help

Modify the primary default route

- 1 Select **Use single static default route**.
- 2 In the **IP address** field, enter the address of the device the firewall forwards traffic to if there is no known route for the destination address. This is usually the IP address of a device that forwards packets to your Internet Service Provider.
 - To find the IP address for a host name, enter the host name in the field and click **DNS Lookup**.
 - To return the IP address to the currently configured default route, click **Use current default route value**.
- 3 [Optional] In the **Description** field, enter information that will help identify the route in the Static Routing window.
- 4 Click **OK** and save your changes.

Note: If you have a DHCP interface configured and enabled, the default route is assigned dynamically. The dynamic address supersedes a single static default route configured in this window. The dynamically assigned default route appears in the read-only **Current default route** field. Click **Refresh** to update this field.

Configure default route failover

To configure redundant default routes, you define an alternate default route and ping addresses for the default routes.

- The Sidewinder continuously pings the default route IP address and any other ping addresses that you define.

- If all configured ping addresses fail, the alternate default route becomes the acting default route.
- Use the Static Route Status window to reset the primary default route when it is active again.

The current default route is shown in a read-only field. Click **Refresh** to update this field.

Figure 269 Default Route Properties window: Default route failover view

The screenshot shows the 'Static Routes: Default Route Properties' window. At the top, the 'Current default route' is displayed as '10.65.249.202' with a refresh icon. Below this, there are two radio buttons: 'Use single static default route' (unselected) and 'Use alternate default routes' (selected). Under 'Use single static default route', there are fields for 'IP address' and 'Description'. Under 'Use alternate default routes', there are two panels: 'Primary default route' and 'Alternate default route'. The 'Primary default route' panel has fields for 'IP address' (containing '10.65.249.202'), 'Description', and 'Ping addresses' (containing '10.65.249.202'). The 'Alternate default route' panel has fields for 'IP address', 'Description', and 'Ping addresses'. At the bottom of each panel are 'Ping interval' and 'Failures allowed' fields. The 'Primary default route' panel has 'Ping interval: 30 (seconds)' and 'Failures allowed: 3'. The 'Alternate default route' panel has 'Ping interval: ' and 'Failures allowed: '. At the bottom of the window are 'OK', 'Cancel', and 'Help' buttons.

1 Select Use alternate default routes.

2 [Optional] Configure the primary default route IP address. The currently configured default route information appears automatically.

- In the **IP address** field, enter the address of the device the firewall forwards traffic to if there is no known route for the destination address. This is usually the IP address of a device that forwards packets to your Internet Service Provider.
 - To find the IP address for a host name, enter the host name in the field and click **DNS Lookup**.
 - To return the IP address to the currently active default route, click **Use current default route value**.
 - To use dynamic addressing for the primary default route, type **dhcp**. You must have a DHCP interface enabled.
- In the **Description** field, enter information that will help identify the route in the Static Routing window.

- 3** [Optional] In the Ping addresses area, configure the IP addresses that the firewall will ping to confirm that the primary default route is accessible.

The primary default route IP address appears automatically. We recommend using an IP address upstream from the alternate default route.

To configure additional ping addresses:

- a** Click **New**, then click the **Specify IP Address** field and type an IP address that the firewall will ping.
- b** In the **Ping interval** field, specify how often (in seconds) the firewall will ping the configured IP addresses to ensure that the path is accessible.
- c** In the **Failures allowed** field, specify the number of failed ping attempts that must occur before the alternate default route takes over as the primary.

Failures are counted in increments and decrements rather than successively. This means that a failed ping adds to the failure total, and a successful ping subtracts from the failure total. The failure total is never less than zero and it is never more than the configured failures allowed.

For example, if the configured failures allowed is **3**, this is how the failure count is tallied based on the ping results:

Table 59 Sample failed ping attempt tally

Ping result:	failure	success	success	failure	failure	success	failure	failure	Failover event occurs
Failure total:	1	0	0	1	2	1	2	3	

- To modify a ping IP address, double-click the address in the list and make the change.
 - To delete a ping IP address, select it in the list and click **Delete**.
- 4** Configure the alternate default route IP address.
- In the **IP address** field, enter the address of the device the firewall forwards traffic to if there is no known route for the destination address. This should be a different route than the primary default route, or to a different ISP.
 - To find the IP address for a host name, enter the host name in the field and click **DNS Lookup**.
 - To return the IP address to the currently active default route, click **Use current default route value**.
 - To use dynamic addressing for the alternate default route, type **dhcp**. You must have a DHCP interface enabled.
 - In the **Description** field, enter information that will help identify the route in the Static Routing window.
- 5** In the Ping addresses area, configure the IP addresses that the firewall will ping to confirm that the alternate default route is accessible.
- a** Click **New**, then click the **Specify IP Address** field and type an IP address that the firewall will ping. We recommend using an IP address upstream from the alternate default route.
 - b** In the **Ping interval** field, specify how often (in seconds) the firewall will ping the configured IP addresses to ensure that the path is accessible.
 - c** In the **Failures allowed** field, specify the number of failed ping attempts that must occur before the alternate default route is considered inaccessible.

Failures are counted in increments and decrements rather than successively. This means that a failed ping adds to the failure total, and a successful ping subtracts from the failure total. The failure total is never less than zero and it is never more than the configured failures allowed.

For example, if the configured failures allowed is **3**, this is how the failure count is tallied based on the ping results:

Table 60 Sample failed ping attempt tally for the alternate default route

Ping result:	failure	success	success	failure	failure	success	failure	failure	Alternate stops forwarding
Failure total:	1	0	0	1	2	1	2	3	

6 Click **OK** and save your changes.

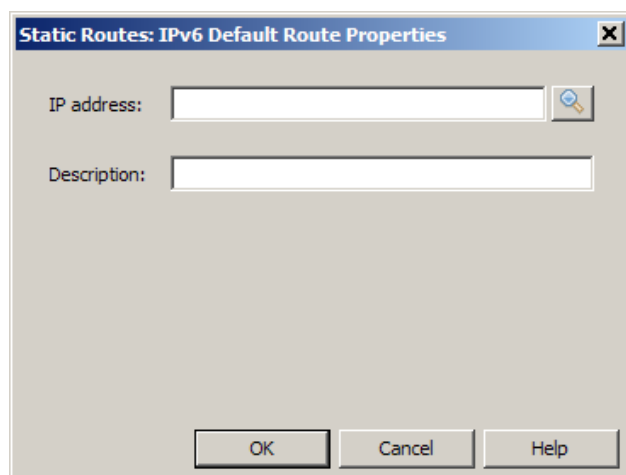
- To modify a ping IP address, double-click the address in the list and make the change.
- To delete a ping IP address, select it in the list and click **Delete**.

Configure the IPv6 default route

Use the IPv6 Default Route Properties window to modify the default route for IPv6.

A default route, also known as the gateway of last resort, is the device the firewall sends traffic to if no other known route exists for the destination address.

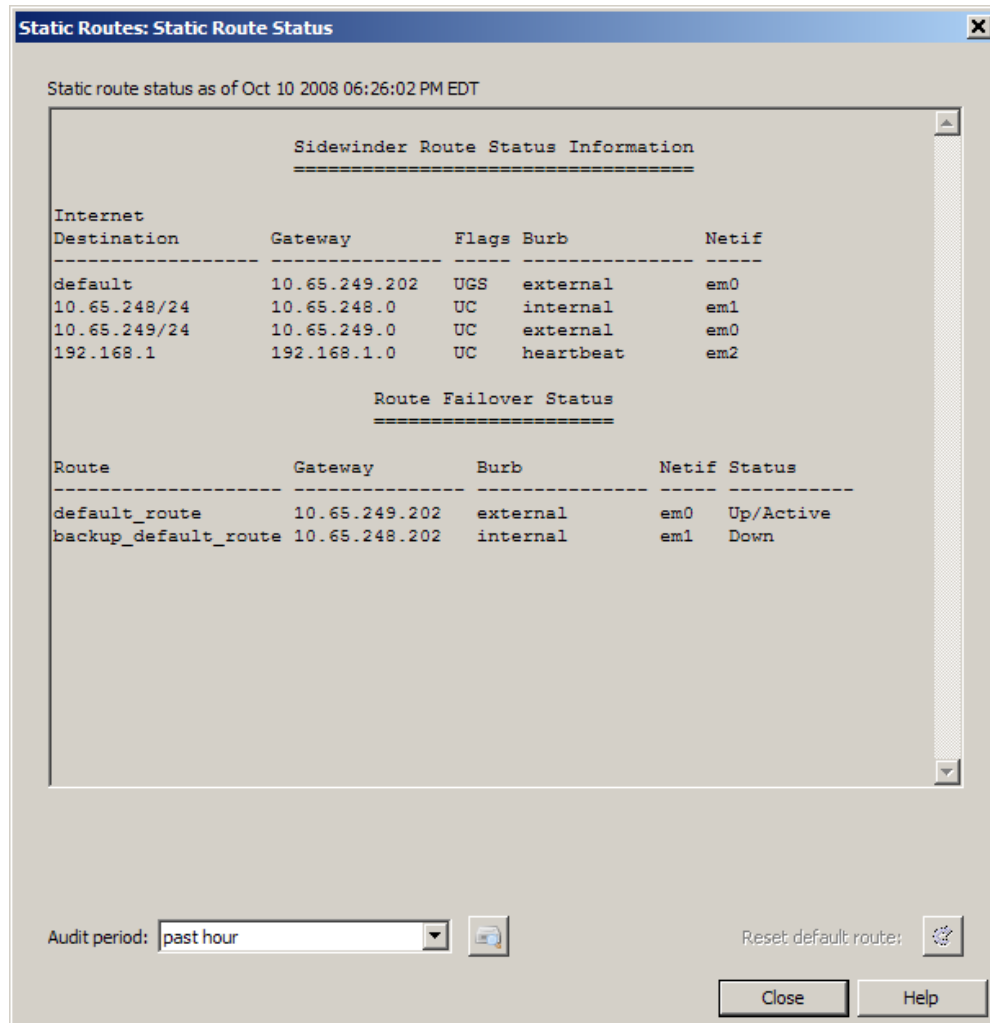
Figure 270 IPv6 Default Route Properties window



- 1** In the **IP address** field, enter the address of the device the firewall forwards traffic to if there is no known route for the destination address. This is usually the IP address of a device that forwards packets to your Internet Service Provider.
To find the IP address for a host name, enter the host name in the field and click **DNS Lookup**.
- 2** [Optional] In the **Description** field, enter information that will help identify the route in the Static Routing window.
- 3** [Conditional] If the static route is a link-local address (begins with **fe80**), you must select an interface from the **Interface** drop-down list.
- 4** Click **OK** and save your changes.

Check route status and reset the default route

Figure 271 Static Route Status window




Use the Static Route Status window to do the following:

- View the status information of the routes configured for the firewall.
- View the status of the failover routes configured for the firewall. A message appears if a default route failover has occurred and traffic is being forwarded to the secondary default route.
- View the audit of any route failover activity: Select a time period from the **Audit period** drop-down list and click **Show route failover audit**.
- Return to using the primary default route: If the route failover status or audit shows that the primary default route is active, click **Reset default route**.

Configure other static routes

Use the Host/Network Route Properties window to create or modify static routes.

Figure 272 Host/Network Route Properties window



- 1 Select the route type:
 - **Host** – Select this option if your destination is a specific IP address.
 - **Network** – Select this option if your destination is a network.
- 2 In the **Description** field, enter information that will help identify the route in the Static Routing window.
- 3 In the **Destination** field, enter the host IP address or subnet address of your end target.
To find the IP address for a host name, enter the host name in the field and click **DNS Lookup**.
- 4 [Network only] Make the appropriate entry.
 - For an IPv4 address, in the **Netmask** field, type the network mask that will be used for this route.
 - For an IPv6 address, in the **Prefix length** field, enter the mask length. Valid values are 0–128
- 5 In the **Gateway** field, type the gateway address that the route will use to pass traffic on to the destination. The gateway address must be reachable by the Sidewinder.
To find the IP address for a host name, enter the host name in the field and click **DNS Lookup**.
- 6 [Conditional] If an IPv6 static route is a link-local address (begins with *fe80*), you must enter a valid interface in the **Interface** field.
- 7 Click **OK** and save your changes.

RIP on Sidewinder

The Routing Information Protocol (RIP) passes dynamic routing information to be used by routers and servers performing routing functions. A router passing RIP traffic can be configured to receive routing information, install routes in its local routing table, and advertise routing information. A router uses this information to determine the shortest available path between networks. By default, routing information is exchanged every 30 seconds and when a router receives updates. rpd operates by listening for UDP broadcasts on port 520. It sets a timer to send a RIP packet advertising its routing information every 30 seconds. When a RIP broadcast is received, the rpd server updates the local routing table with any new routes. When the 30 second timer expires, the rpd server reads and updates its local routing table, and then advertises its local routing information.

ripd can be enabled in two ways:

- unbound – Automatically broadcasts routing information to all burbs
- bound to a burb – Learns routes without broadcasting routing information

Note: Only one ripd method can be enabled in a burb.


The following sections contain scenarios that explain the general concept of RIP processing, some considerations when using RIP on a single burb, and some considerations when using RIP on multiple burbs.

For information on configuration RIP processing, see [Configuring RIP \(ripd\)](#).

Security Alert: In general, dynamic routing is less secure than static routing. If your network requires dynamic routing using RIP, We recommend using RIP v2, which is more secure than RIP v1 and also offers authentication. By default, the Forcepoint implementation of ripd uses v2 without authentication. See the Quagga documentation for enabling authentication.

This example describes how RIP processing aids in routing IP packets through a network that has a redundant routing architecture. [Figure 273](#) illustrates this redundant architecture.

Figure 273 Dynamic routing a with standard IP route



In this example, the Telnet server has a static route to router A, and the Telnet client has a static route to router B. The Telnet client has two different possible paths of reaching the server: (1) via B to A, and (2) via D to C to A. The routing table on router B has two possible routes to the Bizco network: one with a hop count equal to two (through router A), and the other with a hop count to three (through router D). All routers are using RIP to advertise, create, and receive routing information from the other routers.

Typically, when the Telnet client needs to connect to the Telnet server, it sends a connection request to router B (the client's default route). B then forwards the request to router A, because that is the shortest route (two hops versus three hops). Router A then forwards the request to the Telnet server in the Bizco network, which uses the same route to respond to the request.

The dynamic routing capability of RIP can be seen when the link between router A and router B is lost. As soon as B notices that it is no longer receiving RIP updates from A, it updates its local routing table hop count for that route to 16 (route unreachable) and broadcasts this to others on its local network (this is to notify router D).

Next, the Telnet client sends another packet to the server via router A, unaware that the route between A and B has been lost. Router B looks at its local routing table and discovers there are two routes: one is unreachable and the other goes through router D. Because D is on the same network as the client, router B sends an `ICMP Redirect` back at the client stating that it can reach the Telnet server network through router D. The client updates its local routing table to point that host at router D. The client then re-sends its last packet to router D. Router D receives the packet and forwards it on to router C, which forwards it on to router A, etc. The session continues on through router D without interruption. When the link between A and B is re-established, the Telnet client will receive an `ICMP Redirect` from router D pointing it back at router A. The session will again continue without interruption.

Configuring RIP (ripd)

RIP processing is done via a Sidewinder server process called **ripd**. To implement RIP processing, a ripd server process must be configured and there must be an active rule that allows RIP broadcasts; ripd is then enabled in that rule's source burb for ripd bound to a burb, or multiple burbs for ripd unbound service. RIP packets are UDP datagrams with destination port 520. For RIP version 1, the destination address is a network broadcast address such as 10.10.10.255. For RIP version 2, all the routers multicast the address 224.0.0.9. Each burb will have no more than a single ripd instance to handle the network traffic for all interfaces assigned to the burb.

These are the high level steps to set up RIP on the Sidewinder.

- 1 Sketch a diagram showing your planned Sidewinder configuration (similar to the diagrams in [RIP on Sidewinder](#)). Include the following items on your diagram:
 - configuration of the routers to which the firewall connects
 - RIP network
 - the Sidewinder interfaces (burbs)
- 2 Define one or more netgroups for the routers to which the firewall connects. See [Creating network objects](#).
- 3 Configure one or more rules for the RIP traffic. See [Create a rule for ripd-unbound](#).
- 4 Configure the appropriate RIP parameters. See [RIP processing options](#).

See the following sections for details on these high level steps.

Using RIP in your network is a two-step process: First you must create a rule that allows ripd to pass traffic. Then you must configure ripd with the appropriate network information and processing options.

Create a rule for ripd-unbound

To pass RIP traffic in more than one burb, you must run the same ripd instance in more than one burb. To do this, configure the ripd unbound service. Create a rule with the Service field set to **ripd-unbound**; using ripd-unbound in an enabled rule automatically enables the ripd-unbound server in the rule's source burb. (You cannot access the ripd-unbound configuration files using the CLI until this rule is created and enabled.) You can disable the server by disabling or deleting all rules that use ripd-unbound as a service, and by disabling the ripd-unbound server in its configuration file.

To create a ripd rule:

- 1 Select **Policy > Rules**.
- 2 Click **New Rule**.
- 3 Enter a name and description that quickly identified this as the rule that provides access to the ripd-unbound server.
- 4 In the **Service** field, select **ripd-unbound**.
- 5 Configure the Source and Destination fields as necessary to enforce your RIP security policy.
Note: The same burb cannot be used in both a rule using the ripd-unbound service and a rule using the ripd service.
- 6 Save your changes.

Create a rule for ripd bound to a burb

To pass RIP traffic bound to a burb, the firewall needs a rule with the Service field set to **ripd**. The source and destination burbs must be the same, and should be set to the burb on which you intend to receive RIP packets. The source endpoint represents who you want to accept RIP traffic from, such as a single router or a netgroup of routers and/or hosts. The destination endpoint will usually be set to **Any**, since the destination is the broadcast address that corresponds to the source and destination burb.

Using ripd in an enabled rule automatically enables the ripd server in the rule's source burb. (You cannot access the ripd configuration files using the CLI until this rule is created and enabled.) You can disable the server by disabling or deleting all rules that use ripd as a service, and by disabling the ripd server in its configuration file.

Note: Use of the ripd service binds the ripd server to a single burb. Since no routes can be shared between ripd servers, ripd learns routes only in that burb.

To create a ripd rule:

- 1 Select **Policy > Rules**.
- 2 Click **New Rule**.

- 3 Enter a name and description that quickly identified this as the rule that provides access to the ripd server.
- 4 In the Service field, select **ripd**.
- 5 Set the Source Burb and the Destination Burb fields to the same burb. This enables ripd in that burb.
 - The same burb cannot be used in both a rule using the ripd-unbound service and a rule using the ripd service.
 - You can enable ripd in multiple burbs. There is one configuration file per burb, and each file must be edited separately.
- 6 Configure the other Source and Destination fields as necessary to enforce your RIP security policy.
- 7 Save your changes.

For the firewall to pass RIP traffic, you now need to configure the ripd configuration file with the settings appropriate for your security policy. See the following section for the preferred method for enabling and disabling the ripd server.

Configure basic ripd processing

There are several ways to configure ripd on the Sidewinder. They are:

- Using Telnet to connect to the ripd server on the firewall.
- Using the Admin Console File Editor to edit the ripd configuration file.
- Using a different file editor, such as vi, to edit the ripd configuration file.

Because the CLI method provides ripd help and validates commands as they are entered, the following sections focus on this method. The same commands and functionality described here are valid when using the other methods, but require different formatting. Be sure that you are familiar with ripd formatting conventions before using those methods.

For additional documentation on RIP processing, see the official Quagga web site at www.quagga.net.

To enable basic ripd processing using a CLI:

- 1 Using a command line session, log into the firewall and switch to the Admn domain by entering:

```
srole
```

- 2 Telnet into the Sidewinder ripd server on localhost by entering the appropriate command:

- unbound – `telnet localhost ripd`
- bound to burb – `telnet localhost_n ripd`

where *n* = the burb index of the burb used as the source burb in the enabled ripd rule.

Tip: Use `cf burb query` to look up a burb's index. It is also listed on the **Network > Burb Configuration** window as the ID.

A password prompt appears.

- 3 Enter `zebra`.

A `ripd>` prompt appears.

- 4 Enable the full command set by entering:

```
ripd>en
```

The prompt changes to `ripd#` to indicate that the full command set is enabled.

- 5 Enable configuration mode by entering:

```
(config)#conf t
```

The prompt changes to `ripd(config)#` to indicate that configuration mode is enabled.

- 6 Enable ripd and configure it to advertise routes, receive updates, and install routes in the local routing table by entering the following commands:

```
(config)#router rip
(config-router)#network X.X.X.X/mask
```

where *X.X.X.X/mask* is the subnet and network mask of the interface on which you are enabling RIP. You can enter multiple network statements.

- 7 [Optional] To make changes persistent across reboots, write the changes to the configuration file by entering:

```
(config)#write
```

ripd is now enabled and is sending, receiving, and creating routing information. See the following section for information on other configuration options.

To disable ripd, follow [Step 1](#) through [Step 5](#) in the previous procedure, and then enter:

```
(config)#no router rip
```

ripd is now disabled and will not participating in routing.

RIP processing options

The following is a list of common RIP configurations and the commands to implement these configurations. Only administrators who are experienced with routing in general, and RIP dynamic routing in particular, should configure ripd. These commands are presented as they are entered at a command line interface. They also assume that you have entered the appropriate network statements when you first accessed the ripd server. Another option is to configure these options by using the Admin Console File Editor or other file editor to edit the configuration file directly. If you chose to modify the file directly, pay close attention to formatting. See the Quagga documentation at www.quagga.net for formatting assistance.

Tip: Use the ripd online help, available when using the CLI, for details on modifying the commands given here as well as other supported configurations. To access the ripd online help, enter a mode (such as router rip or route-map) and then enter **?** or **list**. You must be currently running a mode to see its documentation.

- **Receive and create routes, but do not advertise routes**

This configuration enables RIP on all interfaces that are on the specified subnet. In this option, ripd receives updates and creates routes in the local routing table, but does not advertise routes.

Use these commands to configure this option:

```
(config)#router rip
(config-router)#passive-interface if_name
```

where *if_name* is the interface name of the burb that is to learn routes, but does not advertise routes. Use **default** instead of an interface name to set this configuration on all interfaces.

- **Advertise routing information, but do not receive or create routes**

This configuration enables ripd to send RIP updates that advertise local routing information available within the current burb. RIP ignores received updates and does not create routes in the local routing table.

Use these command to configure this option:

```
(config)#ip prefix-list name seq n deny x.x.x.x/mask
(config)#router rip
(config-router)#distribute-list prefix name in|out
```

where:

- *name* is the name of the prefix-list
- *n* indicates the order of the prefix-list. Sequence numbers are generally multiples of 5.
- *x.x.x.x/mask* is the IP address and netmask that identified the route. To include all routes, use **any**.
- use **in** to filter routes received by this burb and **out** to filter routes sent by this burb.

For example, you would create an `ip prefix-list` named **none** with a seq 5 that denies all routes. The second command uses `distribute-list` to filter out all received (inbound) updates.

- **Advertise as the default route**

This configuration enables ripd to advertise the default route prefix.

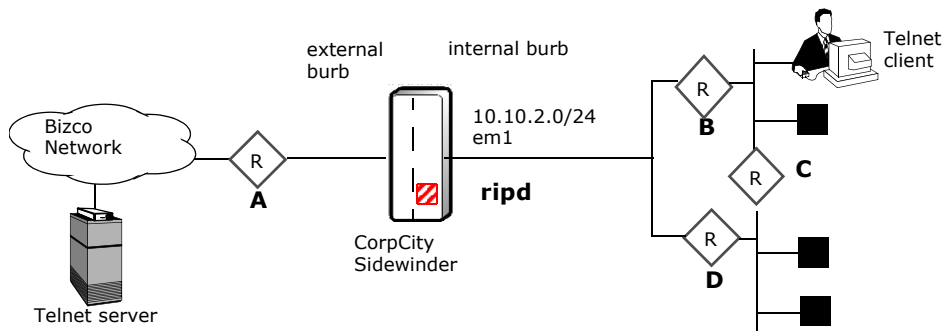
Use this command to configure this option:

```
(config)#router rip
(config-router)#default-information originate
```

Enabling ripd on a single Sidewinder burb

A simple implementation of RIP on the Sidewinder is to enable ripd in a single burb. This configuration is useful when the firewall has a burb that is connected to a network with a redundant routing topology and the firewall needs to participate in that routing infrastructure, but does not need to share that information with other burbs.

Figure 274 Using RIP in a single Sidewinder burb



In this scenario, the company security policy calls for ripd to participate in dynamic routing internally without sharing routing information with any other burbs. To achieve this goal, an administrator enables ripd on the internal burb. If any of the internal routers (B, C, or D) becomes unreachable, ripd receives this information, updates its routing table accordingly, and then advertises the change. For example, if the Telnet client was using router B and it goes down, the client's host machine gets an update for the Sidewinder ripd and reroutes its request through router C and D. When router B is available, the client's host machine receives that update and begins using router B again. On the external burb, the firewall maintains a static route with router A.

To implement this policy, the administrator configures the following ripd options on the *internal* burb:

- Advertise routing information to the internal burb
- Distribute a default route
- Receive routing information from other routers on the internal burb
- Does not send or receive information from any other burbs

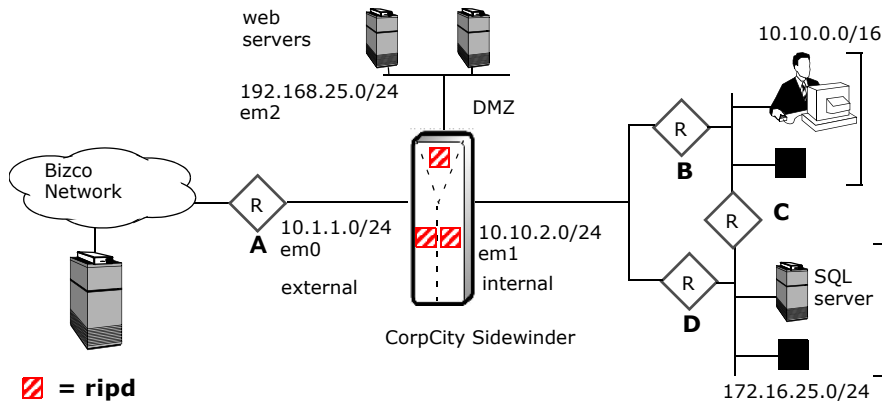
The configuration file for this policy would be similar to the following:

```
!ripd.conf.internal for internal burb
router rip
  network 10.10.2.0/24
  default-information originate
```

Enabling RIP processing on multiple Sidewinder burbs

Using ripd in multiple Sidewinder burbs involves more options than using it in a single burb. You can make decisions about what information to share and what information to filter out.

Figure 275 Using unbound RIP in multiple Sidewinder burbs



In this scenario, the company security policy calls for using unbound RIP to share routing information between the external burb and the DMZ burb, while passing routing information from the internal burb. The administrator must configure the ripd-unbound server to pass routing information between the DMZ and the external burb, and advertise the subnet containing the company's SQL servers, but filter out the routing information for the subnet hosting the employees' workstations.

To implement this policy, the administrator configures the ripd-unbound service for the external burb and the DMZ burb to share all information, but only advertise the SQL subnet information from the internal burb.

The configuration file for the *external* burb and the *DMZ* burb would be similar to the following:

```
!ripd.conf for dmz and internal burb

router rip#
  network 196.168.25.0/24
  network 10.1.1.0/24
  route 172.16.25.0/24
```

The administrator then configures the ripd service with these options on the *internal* burb:

- Advertise routing information within the internal burb
- Distribute a default route to the internal burb

The configuration file for this policy would be similar to the following:

- ripd service bound to internal burb

```
!ripd.conf.internal for internal burb


router rip
  network 10.10.2.0/24
  default-information originate
```

Viewing and comparing ripd configurations

The Admin Console provides tools to help you manage your RIP configuration. You can use these tools to quickly view the entire configuration file, compare different states of the configuration file, or list items such as the RIP neighbors and routes. You can also use the RIP area to edit the configuration file using the File Editor and to manually overwrite the configuration to be used the next time the ripd restarts.

To use these tools, select **Network > Routing > Dynamic Routing > RIP** or **Policy > Rule Elements > Services > ripd**. The following window appears:

Figure 276 The Dynamic Routing > RIP window



Use this window to view and compare versions of the configuration file. The different versions are:

- **Starting configuration** – This is the version that is used when ripd server restarts. The following events update this version:
 - An administrator makes changes using the CLI and saves the changes using the **write** command.
 - An administrator uses the Admin Console File Editor to save the configuration file or uses the **Overwrite** button.
 - An administrator saves changes using a file editor, such as vi, and then restarts ripd.
- **Running configuration** – This is the version currently being used by the firewall. This may differ from the configuration file and the starting configuration if an administrator logs into the server using the CLI and makes changes, but does not issue the **write** command.
- **Configuration file** – This is the most recent saved configuration. If an administrator makes changes using a non-Admin Console file editor but does not restart ripd, this version will be different from starting configuration file.

On this window, you can do the following:

- **Determine which configuration file to view and edit.** The unbound RIP option and each burb have a separate configuration file. Select an option from the **Burb** drop-down list to determine which configuration file to manage.
Note: In addition to editing a ripd configuration file, you must create a rule before RIP traffic can be passed. See [Create a rule for ripd-unbound](#) and [Create a rule for ripd bound to a burb](#).
- **Edit a configuration file.** Click **Edit** to open the selected burb's configuration file using the Admin Console File Editor. Edit the file as needed and then save your changes. The firewall automatically restarts the ripd server. See [Configuring RIP \(ripd\)](#) for more information.
- **View and compare files.** Select an option from the list and then click **Retrieve**. A pop-up window appears displaying the requested information. Close this pop-up to return to the main RIP window.
- **Save the running configuration to the configuration file.** Click **Overwrite** to save the running configuration. The running configuration and the starting configuration are now the same.

OSPF on Sidewinder

The Open Shortest Path First (OSPF) protocol passes link-state information about the internal routers in a given network. All routers communicating using OSPF use an algorithm to calculate the shortest path among the routers. On the Sidewinder, OSPF processing is done via a Sidewinder server process called `ospfd`. To implement OSPF processing, an `ospfd` server process must be configured and there must be an active rule that allows OSPF broadcasts. Unlike `ripd` which is burb-specific, `ospfd` automatically advertises its routing information to all burbs on the firewall. OSPF runs as its own protocol (protocol 89) at the IP layer. OSPF uses 224.0.0.5 and 224.0.0.6 as broadcast addresses.

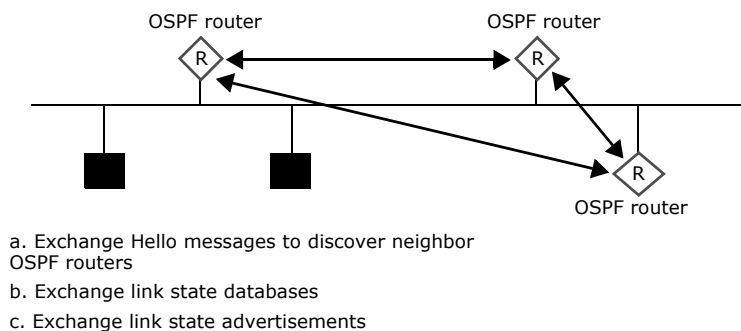
OSPF multicasts information frequently. When a host detects a change to a routing table or a change in the network topology, it immediately multicasts the information to all other hosts in the network. Unlike the RIP in which the entire routing table is sent, the host using OSPF sends only the part that has changed. With RIP, the routing table is sent to neighboring hosts every 30 seconds. OSPF multicasts updated information only when a change occurs.

Rather than counting the number of hops, OSPF bases its path descriptions on *link states* that factor in additional network information. Also, OSPF lets you assign cost metrics to a given host router so that some paths are given preference.

There are three phases to the OSPF protocol:

- 1 Routers discover neighboring OSPF routers by exchanging Hello messages. The Hello messages also determine which routers are to act as the *Designated Router* (DR) and *Backup Designated Router* (BDR). These messages are exchanged periodically to ensure connectivity between neighbors still exists.
- 2 Routers exchange their *link state databases*. *Link state* means the information about a system's interfaces, such as its IP address, network mask, the cost for using that interface, and whether it is up or down.
- 3 The routers exchange additional information via a number of different type of *Link State Advertisements* (LSAs). These supply the information needed to calculate routes. Some reasons for generating LSAs are interfaces going up or down, distant routes changing, static routes being added or deleted, etc.

Figure 277 Three OSPF protocol phases



At this point, all routers should have a full database. Each database contains consistent (not identical) information about the network. Based upon this information, routes are calculated via the “Dijkstra” algorithm. This algorithm generates the set of shortest routes needed to traverse the network. These routes are then enabled for use by IP.

All OSPF routers on a network do not exchange OSPF data—this limits network overhead. Instead, they communicate with the DR (and BDR), which are then responsible for updating all other routers on the network. Election of the DR is based upon the priority of that router. OSPF multicasts using the *AllSPFRouters* (224.0.0.5) and *AllIDRouters* (224.0.0.6) addresses. The DR and the BDR receive packets on the second address.

Note: Since the Sidewinder performs many other functions, We recommend that customers should not configure the firewall to become DR (or BDR) unless forced to by network topology.


OSPF is considered an Interior Gateway Protocol (IGP). An IGP limits the exchange of routes to a *domain of control*, known as an Autonomous System (AS). An AS is a large network created under a central authority running a consistent routing policy that includes different routing protocols, such as the networks commonly run by ISPs. RIP V1 and V2 are also IGPs.

Routers on the edge of the AS generate special LSAs (AS-External-LSAs) for the rest of the AS. There is also an address-forwarding mechanism that allows an OSPF router to obtain a route from a specified location. This feature allows a customer to introduce static routes for their network from a central router.

Autonomous Systems can be large. It is not necessary for the whole AS to know everything about all routes. Each AS may be broken down into areas. All routing information must be identical within an area. Routing between areas goes through a *backbone*. All routers on a backbone have to be able to communicate with each other. Since they belong to the same area (area 0 of a particular AS), they also all have to agree. *Area Border Routers* (ABRs) have one interface defined to run in the backbone area. Other interfaces can then be defined to run in a different area.

The following figure is a sample configuration of OSPF areas. [Figure 278](#) shows a large internal network and backbone terminating at a router.

Figure 278 OSPF areas



For additional documentation on OSPF processing, see the official Quagga web site at www.quagga.net.

Tip: You should use OSPF only if you have identified that your routing topology is too complicated to use only static routing or the Routing Information Protocol (RIP). OSPF is a complex IP routing protocol and deploying OSPF should involve discussions between routing subject matter experts and security subject matter experts.

To implement OSPF processing on the Sidewinder, you must create an enabled rule with `ospfd` selected in the Service field and the Source and Destination Burbs set to Any. You can control which routers `ospfd` can communicate with by managing the source and destination endpoints in the `ospfd` rule. Each burb will have no more than a single `ospfd` instance to handle the network traffic for all interfaces assigned to the burb.

The OSPF implementation on the Sidewinder supports all of the standards specified in RFC 2328.

Configuring OSPF (ospfd)

See the following section for information on configuring OSPF processing.

These are the high level steps to set up OSPF on the Sidewinder.

- 1 Sketch a diagram showing your planned Sidewinder configuration (similar to the diagrams in [RIP on Sidewinder](#)). Include the following items on your diagram:
 - configuration of the routers to which the firewall connects
 - OSPF areas in the network(s)
 - the Sidewinder interfaces (burbs)
- 2 Define one or more netgroups for the routers to which the firewall connects. See [Creating network objects](#).
- 3 Configure one or more rules for the OSPF traffic. See [Create a rule for ospfd](#).
- 4 Configure the appropriate OSPF parameters. See [OSPF processing options](#).

Using OSPF in your network is a two-step process: First you must create a rule that allows `ospfd` traffic. Then you must configure `ospfd` with the appropriate network information and processing options.

Create a rule for ospfd

To enable access to the `ospfd` configuration file:

- 1 Select **Policy > Rules**.
- 2 Click **New Rule**.
- 3 Enter a name and description that quickly identified this as the rule that provides access to the `ospfd` server.
- 4 In the Service field, select **ospfd**.
- 5 Set both the Source Burb and the Destination Burb fields to **Any**.
- 6 Configure the other Source and Destination fields as necessary to enforce your OSPF security policy.

7 Save your changes.

For the firewall to pass OSPF traffic, you now need to configure the ospfd configuration file with the settings appropriate for your security policy. See the following section for the preferred method for enabling and disabling the ospfd server.

Configure basic ospfd processing

There are several ways to configure ospfd on the Sidewinder. They are:

- Telneting into the ospfd server on the firewall and using a command line interface (CLI).
- Using the Admin Console File Editor to edit the ospfd configuration file.
- Using a different file editor, such as vi, to edit the ospfd configuration file.

Because the CLI method provides ospfd help and validates commands as they are entered, the following sections focus on this method. The same commands and functionality described here are valid when using the other methods, but require different formatting. Be sure that you are familiar with ospfd formatting conventions before using those methods.

For additional documentation on OSPF processing, see the official Quagga web site at www.quagga.net.

To enable basic ospfd processing using a CLI:

- 1** Using a command line session, log into the firewall and switch to the Admn domain by entering:

```
srole
```

- 2** Telnet into the Sidewinder ospfd server by entering:

```
telnet localhost ospfd
```

A password prompt appears.

- 3** Enter **zebra**.

A ospfd> prompt appears.

- 4** Enable the full command set by entering:

```
ospfd>en
```

The prompt changes to ospfd# to indicate that the full command set is enabled.

- 5** Enable configuration mode by entering:

```
(config)#conf t
```

The prompt changes to ospfd(config)# to indicate that configuration mode is enabled.

- 6 Enable ospfd and configure it to advertise routes, receive updates, and install routes in the local routing table by entering the following commands:

```
(config)#router ospf  
(config-router)#network X.X.X.X/mask area n.n.n.n
```

where

- *X.X.X.X/mask* is the subnet and network mask of the interface on which you are enabling OSPF. You can enter multiple network statements.
- *n.n.n.n* is the area within the AS, such as 0.0.0.0 for the backbone area.

- 7 [Optional] To make changes persistent across reboots, write the changes to the configuration file by entering:

```
(config)#write
```

ospfd is now enabled and is advertising, receiving, and creating routing information. See the following section for information on other configuration options.

To disable ospfd, follow [Step 1](#) through [Step 5](#) in the previous procedure, and then enter:

```
(config)#no router ospf
```

ospfd is now disabled and will not participate in routing.

OSPF processing options

As with RIP, only administrators who are experienced with routing in general, and OSPF dynamic routing in particular, should configure ospfd.

These commands are presented as they are entered at a command line interface. They also assume that you have entered the appropriate network and area statements when you first accessed the ospfd server. Another option is to configure these options by using the Admin Console File Editor or other file editor to edit the configuration file directly. If you chose to modify the file directly, pay close attention to formatting. See the Quagga documentation at www.quagga.net for formatting assistance.

Tip: Use the ospfd online help, available when using the CLI, for details on modifying the commands given here as well as other supported configurations. To access the ospfd online help, enter a mode (such as router ospf or route-map) and then enter **?** or **list**. You must be currently running a mode to see its documentation.

In general, the OSPF configuration options are similar to the RIP configuration options, particularly the route-map, prefix-list, and redistribution commands. See [RIP processing options](#) for details. However, the servers' implementation differences of the `passive-interface` command is worth noting.

For both servers, the `passive-interface` command enables the routing protocol on all interfaces that are on the specified subnet. For ripd, the server receives updates and creates routes in the local routing table, but does not advertise routes. For ospfd, the server passively advertises the local interface information, but does not form adjacency with other routers over the specified interface.

For OSPF, use these commands to configure this option:

```
(config)#router ospf  
(config-router)#passive-interface if_name
```

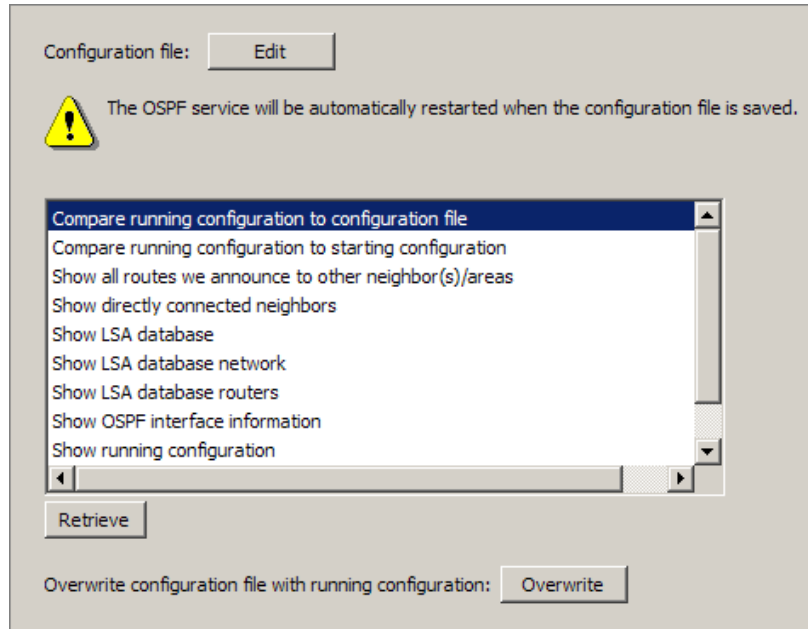
where *if_name* is the interface name of the burb that is to learn routes, but does not send HELLOs to other routers. Use **default** to set this configuration on all interfaces.

Viewing and comparing OSPF configurations

The Admin Console provides tools to help you manage your OSPF configuration. You can use these tools to quickly view the entire configuration file, compare different states of the configuration file, or list items such as the OSPF neighbors and routes. You can also use the OSPF area to edit the configuration file using the File Editor and to overwrite the configuration to be used the next time the ospfd restarts.

To use these tools, select **Network > Routing > Dynamic Routing > OSPF** or **Policy > Rule Elements > Services > ospfd**. The following window appears:

Figure 279 The Dynamic Routing > OSPF window



Use this window to view and compare versions of the configuration file. The different versions are:

- **Starting configuration** – This is the version that is used when ospfd server restarts. The following events update this version:
 - An administrator makes changes using the CLI and saves the changes using the **write** command.
 - An administrator uses the Admin Console File Editor to save the configuration file or uses the **Overwrite** button.
 - An administrator saves changes using a file editor, such as vi, and then restarts ospfd.
- **Running configuration** – This is the version currently being used by the firewall. This may differ from the configuration file and the starting configuration if an administrator logs into the server using the CLI and makes changes, but does not issue the **write** command.
- **Configuration file** – This is the most recent saved configuration. If an administrator makes changes using a non-Admin Console file editor but does not restart ospfd, this version will be different from starting configuration file.

On this window, you can do the following:

- **Edit a configuration file.** Click **Edit** to open the configuration file using the Admin Console File Editor. Edit the file as needed and then save your changes. The firewall automatically restarts the ospfd server. See [Configure basic ospfd processing](#) for more information.
Note: Remember to create a rule using ospfd in the Service field before attempting to pass OSPF traffic. See [Create a rule for ospfd](#).
- **View and compare files.** Select an option from the list and then click **Retrieve**. A pop-up window appears displaying the requested information. Close this pop-up to return to the main OSPF window.
- **Save the running configuration to the configuration file.** Click **Overwrite** to save the running configuration. The running configuration and the starting configuration are now the same.

OSPF IPv6 on Sidewinder

The OSPF IPv6 protocol concepts are the same as OSPF for IPv4. Note the following differences:

- OSPF IPv6 processing is done via a Sidewinder server process called **ospf6d**.
- New LSAs have been created to carry IPv6 addresses and prefixes.
- OSPF IPv6 multicasts using the following broadcast addresses:
 - *AllSPFRouters* – This multicast address has been assigned the value FF02::5. All routers running OSPF should be prepared to receive packets sent to this address. Hello packets are always sent to this destination.
 - *AllDRouters* – This multicast address has been assigned the value FF02::6. Both the Designated Router and Backup Designated Router must be prepared to receive packets destined to this address.

See RFC 2740 for more information

Creating a rule for ospf6d

To enable access to the ospf6d configuration file:

- 1 Select **Policy > Rules**.
- 2 Click **New Rule**.
- 3 Enter a name and description that quickly identified this as the rule that provides access to the ospfd server.
- 4 In the Service field, select **ospf6d**.
- 5 Set both the Source Burp and the Destination Burp fields to **Any**.
- 6 Configure the other Source and Destination fields as necessary to enforce your OSPF IPv6 security policy.
- 7 Save your changes.

Configuring basic ospf6d processing

There are several ways to configure ospf6d on the Sidewinder. They are:

- Telneting into the ospfd server on the firewall and using a command line interface (CLI).
- Using the Admin Console File Editor to edit the ospf6d configuration file.
- Using a different file editor, such as vi, to edit the ospf6d configuration file.

Because the CLI method provides ospfd help and validates commands as they are entered, the following sections focus on this method. The same commands and functionality described here are valid when using the other methods, but require different formatting. Be sure that you are familiar with ospf6d formatting conventions before using those methods.

For additional documentation on OSPF IPv6 processing, see the official Quagga web site at www.quagga.net.

To enable basic ospf6d processing using a CLI:

- 1 Using a command line session, log into the firewall and switch to the Admin domain by entering:

```
srole
```

- 2 Telnet into the Sidewinder ospf6d server by entering:

```
telnet localhost ospf6d
```

A password prompt appears.

3 Enter **zebra**.

An `ospf6d>` prompt appears.

4 Enable the full command set by entering:

```
ospf6d>en
```

The prompt changes to `ospf6d#` to indicate that the full command set is enabled.

5 Enable configuration mode by entering:

```
(config) #conf t
```

The prompt changes to `ospf6d(config) #` to indicate that configuration mode is enabled.

6 Enable ospf6d and configure it to advertise routes, receive updates, and install routes in the local routing table by entering the following commands: -

```
(config) #router ospf6  
(config-router) router-id X.X.X.X  
(config-router) #interface XXX area n.n.n.n
```

where

- `X.X.X.X` is the value other routers will know this router by. Router IDs are the IPv4 size of 32-bits.
- `XXX` is the interface NIC on which you are enabling OSPF IPv6. You can enter multiple interfaces.
- `n.n.n.n` is the area within the AS, such as 0.0.0.0 for the backbone area.

7 [Optional] To make changes persistent across reboots, write the changes to the configuration file by entering:

```
(config) #write
```

ospf6d is now enabled and is advertising, receiving, and creating routing information.

To disable ospf6d, follow [Step 1](#) through [Step 5](#) in the previous procedure, and then enter:

```
(config) #no router ospf6
```

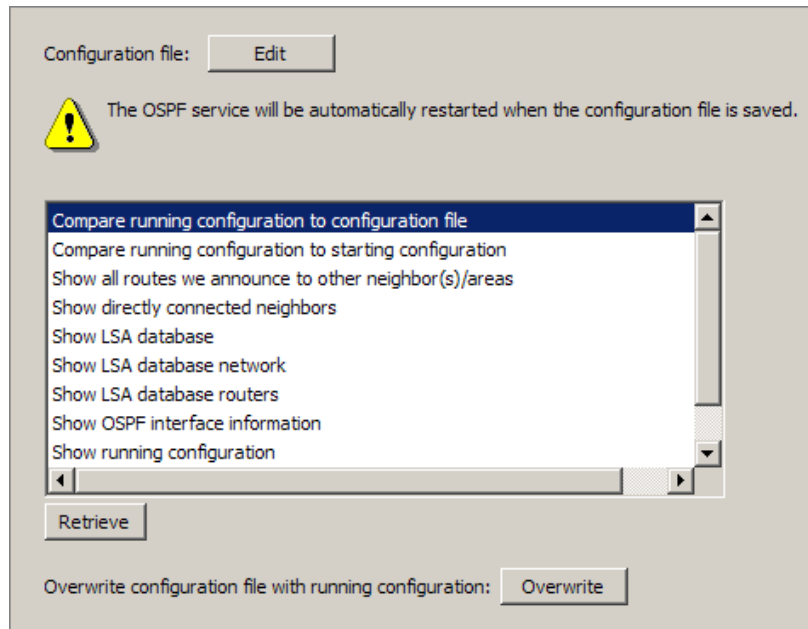
ospf6d is now disabled and will not participate in routing.

Viewing and comparing OSPF IPv6 configurations

The Admin Console provides tools to help you manage your OSPF IPv6 configuration. You can use these tools to quickly view the entire configuration file, compare different states of the configuration file, or list items such as the OSPF IPv6 neighbors and routes. You can also use the OSPF IPv6 area to edit the configuration file using the File Editor and to overwrite the configuration to be used the next time the ospf6d restarts.

To use these tools, select **Network > Routing > Dynamic Routing > OSPF IPv6** or **Policy > Rule Elements > Services > ospf6d**. The OSPF IPv6 window appears.

Figure 280 The Dynamic Routing > OSPF IPv6 window



Use this window to view and compare versions of the configuration file. The different versions are:

- **Starting configuration** – This is the version that is used when ospf6d server restarts. The following events update this version:
 - An administrator makes changes using the CLI and saves the changes using the **write** command.
 - An administrator uses the Admin Console File Editor to save the configuration file or uses the **Overwrite** button.
 - An administrator saves changes using a file editor, such as vi, and then restarts ospf6d.
- **Running configuration** – This is the version currently being used by the firewall. This may differ from the configuration file and the starting configuration if an administrator logs into the server using the CLI and makes changes, but does not issue the **write** command.
- **Configuration file** – This is the most recent saved configuration. If an administrator makes changes using a non-Admin Console file editor but does not restart ospf6d, this version will be different from starting configuration file.

On this window, you can do the following:


- **Edit a configuration file.** Click **Edit** to open the configuration file using the Admin Console File Editor. Edit the file as needed and then save your changes. The firewall automatically restarts the ospf6d server.
Note: Remember to create a rule using ospf6d in the Service field before attempting to pass OSPF IPv6 traffic. See .
- **View and compare files.** Select an option from the list and then click **Retrieve**. A pop-up window appears displaying the requested information. Close this pop-up to return to the main OSPF IPv6 window.
- **Save the running configuration to the configuration file.** Click **Overwrite** to save the running configuration. The running configuration and the starting configuration are now the same.

BGP on Sidewinder

The Border Gateway Protocol (BGP) is an Exterior Gateway Protocol (EGP) used to pass routing information between Autonomous Systems (AS). Unlike OSPF, which is an Interior Gateway Protocol (IGP), BGP is used to connect to external routers, such as your ISP. It does, however, learn information from an interior network that it then passes to an external network.

Routers using BGP are commonly located at the perimeter of an AS, as shown in [Figure 278](#).

Figure 281 BGP areas



Routers employing BGP use TCP connections to communicate with peer routers, known as *neighbors*. After a connection is established, routing information is exchanged. Traffic is passed on port 179. The connection is maintained using keep-alives that are sent by both neighbors at a default rate of every 60 seconds, with a 3 minute timeout.

On the Sidewinder, BGP processing is done via a Sidewinder server process named `bgpd`. To implement BGP processing, a `bgpd` server process must be configured and there must be an active rule that allows BGP broadcasts. You can control which routers `bgpd` can communicate with by managing the source and destination endpoints in the `bgpd` rule. Each burb will have no more than a single `bgpd` instance to handle the network traffic for all interfaces assigned to the burb.

As with the other Sidewinder dynamic routing protocols, see the Quagga documentation for a list of supported features.

Configuring BGP (bgpd)

See the following section for information on configuring BGP processing.

These are the high level steps to set up BGP on the Sidewinder.

- 1 Sketch a diagram showing your planned Sidewinder configuration (similar to the diagrams in [BGP areas](#)). Include the following items on your diagram:
 - configuration of the routers to which the firewall connects
 - BGP areas in the network(s)
 - the Sidewinder interfaces (burbs)
- 2 Define one or more netgroups for the routers to which the firewall connects. See [Creating network objects](#).
- 3 Configure one or more rules for the BGP traffic. See [Create a rule for bgpd](#).
- 4 Configure the appropriate BGP parameters. See [BGP processing options](#).

Using BGP in your network is a two-step process: First you must create a rule that allows `bgpd` traffic. Then you must configure `bgpd` with the appropriate network information and processing options.

Create a rule for bgpd

To enable access to the bgpd configuration file:

- 1 Select **Policy > Rules**.
- 2 Click **New Rule**.
- 3 Enter a name and description that quickly identified this as the rule that provides access to the bgpd server.
- 4 In the Service field, select **bgpd**.
- 5 Set both the Source Burp and the Destination Burp fields to **Any**.
- 6 Configure the other Source and Destination fields as necessary to enforce your BGP security policy.
- 7 Save your changes.

For the firewall to pass BGP traffic, you now need to configure the bgpd configuration file with the settings appropriate for your security policy. See the following section for the preferred method for enabling and disabling the bgpd server.

Configure basic bgpd processing

There are several ways to configure bgpd on the Sidewinder. They are:

- Telneting into the bgpd server on the firewall and using a command line interface (CLI).
- Using the Admin Console File Editor to edit the bgpd configuration file.
- Using a different file editor, such as vi, to edit the bgpd configuration file.

Because the CLI method provides bgpd help and validates commands as they are entered, the following sections focus on this method. The same commands and functionality described here are valid when using the other methods, but require different formatting. Be sure that you are familiar with bgpd formatting conventions before using those methods.

For additional documentation on BGP processing, see the official Quagga web site at www.quagga.net.

To enable basic bgpd processing using a CLI:

- 1 Using a command line session, log into the firewall and switch to the Admn domain by entering:

```
srole
```

- 2 Telnet into the Sidewinder bgpd server by entering:

```
telnet localhost bgpd
```

A password prompt appears.

- 3 Enter **zebra**.

A **bgpd>** prompt appears.

- 4 Enable the full command set by entering:

```
bgpd>en
```

The prompt changes to **bgpd#** to indicate that the full command set is enabled.

- 5 Enable configuration mode by entering:

```
(config)#conf t
```

The prompt changes to **bgpd(config)#** to indicate that configuration mode is enabled.

- 6 Enable bgpd and configure it to advertise routes, receive updates, and install routes in the local routing table by entering the following commands:

```
(config)#router bgp
```

```
(config-router)#network X.X.X.X/mask
```

where *X.X.X.X/mask* is the subnet and network mask of the interface on which you are enabling BGP. You can enter multiple network statements.

- 7 [Optional] To make changes persistent across reboots, write the changes to the configuration file by entering:

```
(config)#write
```

bgpd is now enabled and is advertising, receiving, and creating routing information. See the following section for information on other configuration options.

To disable bgpd, follow [Step 1](#) through [Step 5](#) in the previous procedure, and then enter:

```
(config) #no router bgp
```

bgpd is now disabled and will not participate in routing.

BGP processing options

As with RIP and OSPF, only administrators who are experienced with routing in general, and BGP dynamic routing in particular, should configure bgpd.

These commands are presented as they are entered at a command line interface. They also assume that you have entered the appropriate network and area statements when you first accessed the bgpd server. Another option is to configure these options by using the Admin Console File Editor or other file editor to edit the configuration file directly. If you chose to modify the file directly, pay close attention to formatting. See the Quagga documentation at www.quagga.net for formatting assistance.

Tip: Use the bgpd online help, available when using the CLI, for details on modifying the commands given here as well as other supported configurations. To access the bgpd online help, enter a mode (such as router bgp or route-map) and then enter **?** or **list**. You must be currently running a mode to see its documentation.


In general, the BGP configuration options are similar to the RIP and OSPF configuration options, particularly the route-map, prefix-list, and redistribution commands. See [RIP processing options](#) for details. However, instead of using interface names to identify the source and destination of routing information, BGP uses names of neighbors.

Viewing and comparing BGP configurations

The Admin Console provides tools to help you manage your BGP configuration. You can use these tools to quickly view the entire configuration file, compare different states of the configuration file, or list items such as the BGP neighbors and routes. You can also use the BGP area to edit the configuration file using the File Editor and to manually overwrite the configuration to be used the next time the bgpd restarts.

To use these tools, select **Network > Routing > Dynamic Routing > BGP** or **Policy > Rule Elements > Services > bgpd**. The following window appears:

Figure 282 The Dynamic Routing > BGP window



Use this window to view and compare versions of the configuration file. The different versions are:

- **Starting configuration** – This is the version that is used when bgpd server restarts. The following events update this version:
 - An administrator makes changes using the CLI and saves the changes using the **write** command.
 - An administrator uses the Admin Console File Editor to save the configuration file or uses the **Overwrite** button.
 - An administrator saves changes using a file editor, such as vi, and then restarts bgpd.
- **Running configuration** – This is the version currently being used by the firewall. This may differ from the configuration file and the running configuration if an administrator logs into the server using the CLI and makes changes, but does not issue the **write** command.
- **Configuration file** – This is the most recent saved configuration. If an administrator makes changes using a non-Admin Console file editor but does not restart bgpd, this version will be different from starting configuration file.

On this window, you can do the following:

- **Edit a configuration file.** Click **Edit** to open the configuration file using the Admin Console File Editor. Edit the file as needed and then save your changes. The firewall automatically restarts the bgpd server. See [Configuring BGP \(bgpd\)](#) for more information.
Note: Remember to create a rule using bgpd in the Service field before attempting to pass BGP traffic. See [Create a rule for bgpd](#).
- **View and compare files.** Select an option from the list and then click **Retrieve**. A pop-up window appears displaying the requested information. Close this pop-up to return to the main BGP window.
- **Save the running configuration to the configuration file.** Click **Overwrite** to save the running configuration. The running configuration and the starting configuration are now the same.

PIM-SM on Sidewinder

The Protocol Independent Multicast - Sparse Mode (PIM-SM) protocol is used to route traffic to multicast groups. *Multicast* is communication between a single or multiple senders and multiple receivers on a network. The Sidewinder uses a XORP routing package which contains IGMP and PIM-SM protocols to route multicast traffic:


- The Internet Group Management Protocol (IGMP) is used by hosts and adjacent routers to establish multicast group memberships. IGMP tells routers that a host wants to receive multicast traffic for the specified multicast group.
- The PIM-SM protocol sets up a multicast forwarding table in routers. Multicast traffic is directed to a rendezvous point (RP), which distributes it toward PIM-registered receivers.

When a host wants to join a multicast session, IGMP sends a join request to its gateway router for a multicast group. Since the gateway router doesn't have information about the source address, it will send a PIM join back to the rendezvous point, which will contain the source information.

The rendezvous point facilitates the route setup between the sender and receiver. The sending gateway router sends multicast data to a rendezvous point encapsulated in a unicast PIM packet.

Once a gateway router with direct connection to the receiver's network has received traffic from the source, the gateway router might start a process to build a direct path from the sender to the source.

Figure 283 Multicast routing using IGMP and PIM-SM protocols



Configuring PIM-SM (pimd)

To configure a Sidewinder to route multicast traffic using PIM-SM, you must perform the following procedures:

- 1 [Create policy rules](#) to enable the pimd service and allow multicast traffic and PIM traffic forwarding.
- 2 [Configure the pimd \(XORP server\) service.](#)
- 3 [Configure IGMP.](#)
- 4 [Configure PIM-SM.](#)
- 5 [Restart the pimd \(XORP server\) service](#) (XORP server) service.

It is recommended that you make all of these configuration changes at one time, since you must restart the pimd service to initialize your changes.

Note: When making subsequent changes to PIM-SM, there are two types of changes that require different procedures. See [Exceptions to making PIM-SM changes](#) for more information.

Create policy rules

You must create these policy rules to allow multicast routing:

Create a rule to enable the pimd (XORP server) service

- 1 Select **Policy > Rules**.
- 2 Click **New Rule**.
- 3 Enter a name and description that quickly identifies this as the rule that enables the pimd (XORP server) service.
- 4 In the Service field, select **pimd** from the drop-down list.
- 5 Set both the Source Burb and the Destination Burb fields to **Any**.
- 6 Configure the other Source and Destination fields as necessary to enforce your PIM-SM security policy.
- 7 Save your changes.

Create a rule to enable PIM traffic forwarding to rendezvous points and bootstrap routers

- 1 Create a packet filter service for the rule:
 - a Select **Policy > Rule Elements > Services**.
 - b Click **New Service**. The New Service window appears.
 - c Enter a name and description that easily identifies the service.
 - d From the Agent drop-down list, select **Other Protocol Packet Filter**.
 - e From the Protocol drop-down list, select **103 - pim**.
 - f Select **Bi-directional**.
 - g Click **Add** and save your changes.
- 2 Create a rule using the service:
 - a Select **Policy > Rules**.
 - b Click **New Rule**.
 - c Enter a name and description that quickly identifies this as the rule that enables PIM traffic forwarding.
 - d In the Service field, select the new traffic forwarding service.
 - e Set both the Source Burb and the Destination Burb fields to **Any**.
 - f Configure the other Source and Destination fields as necessary to enforce your PIM-SM security policy. Include all rendezvous points and bootstrap routers within the PIM network.
 - g Click **Add** and save your changes.

Create a rule to enable multicast traffic

- 1 Create a packet filter service for the rule:
 - a Select **Policy > Rule Elements > Services**.
 - b Click **New Service**. The New Service window appears.
 - c Enter a name and description that easily identifies the service.
 - d From the Agent drop-down list, select **TCP/UDP Packet Filter**.
 - e In the UDP ports field, select the UDP ports your multicast applications will be using.
 - f Select **Bi-directional**.
 - g Make any other changes necessary for your site's security policy.
 - h Click **Add** and save your changes.
- 2 Create a rule using the service:

- a Select **Policy > Rules**.
- b Click **New Rule**.
- c Enter a name and description that quickly identifies this as the rule that enables multicast traffic forwarding.
- d In the Service field, select the new multicast traffic service.
- e Configure the Source and Destination fields as necessary to enforce your multicast security policy. Include the multicast groups in the Destination Endpoint field.
- f Click **Add** and save your changes.

Configure the pimd (XORP server) service

- 1 Select **Network > Routing > Dynamic Routing > PIMSM**.
- 2 Click **Edit**. The xorp configuration file opens in the File Editor.
- 3 Verify that the interface names in the file are correct.
- 4 Remove the comments for these parameters:

Table 61 pimd parameters





Parameters	PIM-SM Editor window
Interfaces you want to run multicast over. <i>default-system-config</i> causes pimd to use the interface configuration from the system kernel.	 <pre> File Edit Help /secureos/etc/xorp/config.boot /*XORP Configuration File, v1.0*/ interfaces { restore-original-config-on-shutdown: false interface bce0 { description: "data interface" disable: false default-system-config } interface bce1 { description: "data interface" disable: false default-system-config } } </pre>
<i>mfea4</i> identifies which interfaces are being used for multicast traffic. <i>register_vif</i> is necessary for XORP processing.	 <pre> File Edit Help /secureos/etc/xorp/config.boot plumbing { mfea4 { disable: false interface bce0 { vif bce0 { disable: false } } interface register_vif { vif register_vif { disable: false } } traceoptions { flag all { disable: false } } } } </pre>

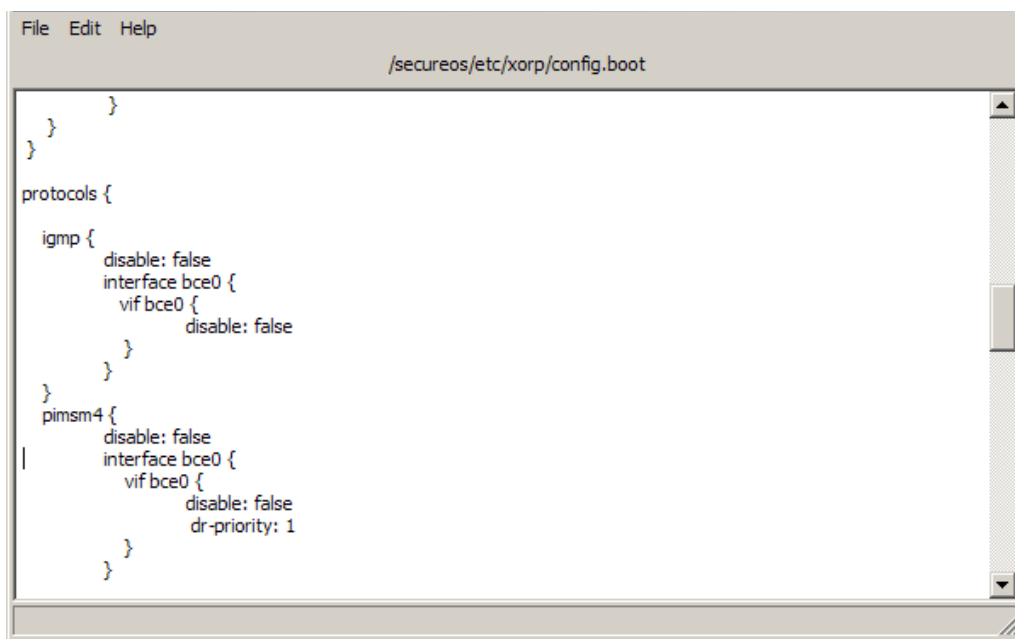
Table 61 pimd parameters <Comment>(continued)

Parameters	PIM-SM Editor window
<i>fea</i> tells pimd how to locate unicast routes.	 <pre> File Edit Help /secureos/etc/xorp/config.boot fea { unicast-forwarding4 { disable: false forwarding-entries { retain-on-startup: false retain-on-shutdown: false } } } protocols { pimsm4 { disable: false interface bce0 { vif bce0 { disable: false dr-priority: 1 } } } } </pre>
<i>fib2mrib</i> tells PIM-SM to use the unicast routing table to find a route to the rendezvous points and to the sender.	 <pre> File Edit Help /secureos/etc/xorp/config.boot * Note: fib2mrib is needed for multicast only if the unicast protocols * don't populate the MRIB with multicast-specific routes. protocols { fib2mrib { disable: false } } * See xorp/mibs/snmpdscrips/README on how to configure Net-SNMP in your host * before uncommenting the snmp section below. * Also check that the "bgp4_mib_1657.so" exists in the correct location. */ protocols { snmp { mib-module bgp4_mib_1657 { abs-path: "/usr/local/xorp/mibs/bgp4_mib_1657.so" } } } </pre>

Configure IGMP

- 1 [If necessary] Select **Network > Routing > Dynamic Routing > PIMSM** and click **Edit** to open the xorp configuration file.
- 2 Add an IGMP clause to the configuration file, specifying the interfaces to networks where hosts are receiving multicast packets. See the example below.

Figure 284 IGMP added to the xorp configuration file



Note: To disable IGMP for the network, disable the corresponding interface in the igmp section.

- 3 Save your changes.

Configure PIM-SM

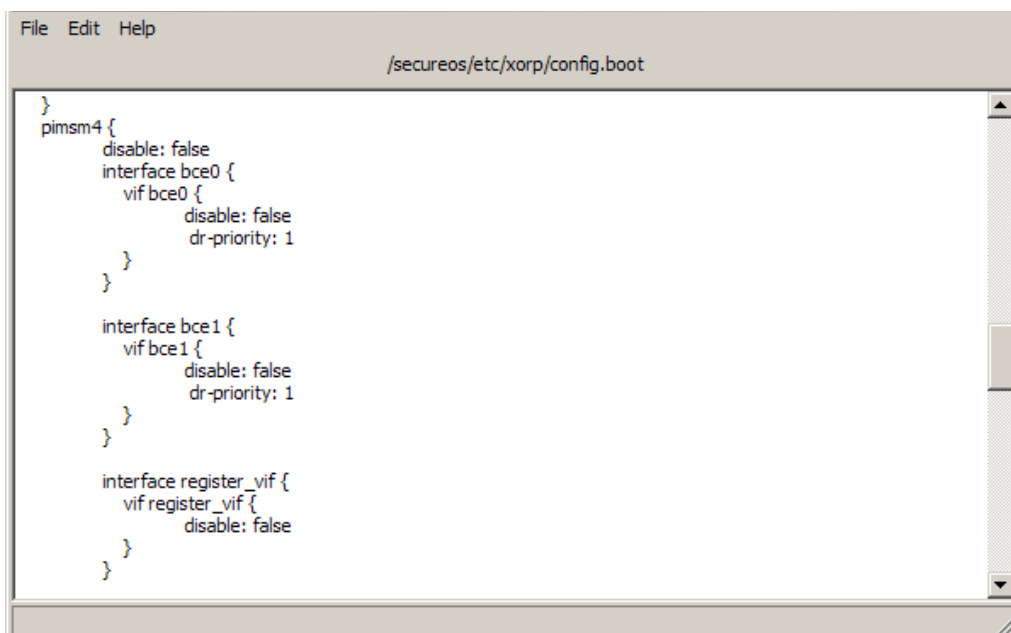
You need to perform two tasks for a dynamic PIM-SM configuration:

- Specify interfaces you expect to receive multicast traffic.
- Configure the rendezvous points for a dynamic or static configuration.

Specify interfaces

- 1 [If necessary] Select **Network > Routing > Dynamic Routing > PIMSM** and click **Edit** to open the xorp configuration file.
- 2 Configure the interfaces that will run PIM-SM.
 - For each interface, an interface statement within the *pimsm4* section of the config file must be included.
 - *register_vif* must be included.

Figure 285 Bootstrap router parameters



Configure the rendezvous points

There are two ways to configure the rendezvous point: dynamically with a bootstrap router or using static configuration.

To configure rendezvous points dynamically with a bootstrap router:


The bootstrap (dynamic) protocol is useful for large networks.

- You can have multiple rendezvous points—if one rendezvous point goes away, another one is elected.
- You can specify whether you want to be a rendezvous point, and you can specify whether you want to communicate with another router that is a rendezvous point.
- You do not have to configure rendezvous points—rendezvous points are learned. You can specify which interfaces on your firewall can learn the rendezvous points.

- 1 [If necessary] Select **Network > Routing > Dynamic Routing > PIMSM** and click **Edit** to open the xorp configuration file.

Note: You cannot change the bsr-priority (bootstrap router priority) setting in this file. If you need to change this setting, see [Change the bsr-priority setting](#) for instructions.

Figure 286 Bootstrap router parameters



2 Remove the comments from the bootstrap router and rendezvous points.

- *cand-bsr* is the bootstrap protocol that selects a bootstrap router. The bootstrap router tells all PIM-SM routers what the rendezvous points are.
- *cand-rp* tells the bootstrap router that this router is a candidate to be a rendezvous point.
- *switch-to-spt-threshold* lets you specify the data rate at which the router selects the shortest path between the sender and the receiver.
 - If you have a lot of multicast traffic and use multicast for a long time, finding a shortest path is useful.
 - If you don't have much traffic, or if you use multicast for a short time, finding a shortest path isn't necessary.

See xorp documentation for more information.

- *traceoptions* sends debug tracing to syslog.


3 Save your changes.

To configure rendezvous points using static configuration:

Static PIM-SM is a simpler configuration that is useful for smaller networks, for example, if you have only two PIM routers or if your ISP provides the rendezvous point.

- 1** [If necessary] Select **Network > Routing > Dynamic Routing > PIMSM** and click **Edit** to open the xorp configuration file.
- 2** Add static-rps clause to the configuration file, specifying the rendezvous point for a range of group prefixes. See the example below.
 - If more than one rendezvous point is specified for a group, the rendezvous point with the lowest priority is used.
 - All PIM-SM routers must be configured with the same rendezvous points.

Figure 287 Bootstrap router parameters



3 Save your changes.

Restart the pimd (XORP server) service

- 1 Select **Policy > Rules**.
- 2 Select the rule that uses the pimd service and click **Modify**.
- 3 Clear the **Enable** box.
- 4 Click **OK** and save your changes.
- 5 Select the rule that uses the pimd service and click **Modify**.
- 6 Select the **Enable** box.
- 7 Click **OK** and save your changes.

Exceptions to making PIM-SM changes

The procedures in this document explain how to configure the XORP server using the Sidewinder Admin Console. To configure the XORP server through a command line interface, you use the XORP command shell *xorpsh*.

The Admin Console's PIMSM window, *xorpsh*, and any file editor open the same config.boot file (/secureos/etc/xorp/config.boot). However, the PIMSM editor and *xorpsh* interact, which can cause conflicts.

To avoid conflicts, there are two types of changes to PIM-SM that require different procedures:

- Disabling and enabling PIM-SM
- Changing the bsr-priority setting

Disable and enable PIM-SM

You cannot use *xorpsh* to enable or disable PIM-SM. To avoid an error message, you must enable or disable the rule that uses the pimd service.

- 1 Select **Policy > Rules**.
- 2 Select the rule that uses the pimd service and click **Modify**.
- 3 Make the appropriate action:
 - To disable the pimd (XORP server) service, clear the **Enable** box.
 - To enable the pimd (XORP server) service, select the **Enable** box.
- 4 Click **OK** and save your changes.

Change the bsr-priority setting

The procedures in this document explain how to configure the XORP server using the Sidewinder Admin Console. To configure the XORP server through a command line interface, you use the XORP command shell *xorpsh*.

The Admin Console's PIMSM window, *xorpsh*, and any file editor open the same config.boot file (/secureos/etc/xorp/config.boot). However, the PIMSM editor and *xorpsh* interact, which can cause conflicts.

To avoid conflicts, you cannot change the bsr-priority (bootstrap) parameter using the Edit function on the PIMSM window. To avoid an error message, you must stop the XORP server, change the parameter, and restart the XORP server.

To change the bsr-priority parameter:

- 1 Stop the XORP server:
 - a Select **Policy > Rules**.
 - b Select the rule that uses the pimd service and click **Modify**.
 - c Clear the **Enable** box.
 - d Click **OK** and save your changes.
- 2 Change the bsr-priority parameter:
 - a Select **Maintenance > File Editor** and open the following firewall file:
/secureos/etc/xorp/config.boot
 - b Make the desired change to the bsr-priority parameter.
 - c Save your changes and close the File Editor.

3 Start the XORP server:


- a Select **Policy > Rules**.
- b Select the rule that uses the pimd service and click **Modify**.
- c Select the **Enable** box.
- d Click **OK** and save your changes.

Viewing PIM-SM configurations

Use the PIM-SM window to view and configure PIM-SM routing parameters.

Select **Network > Routing > Dynamic Routing > PIMSM** or **Policy > Rule Elements > Services > pimd**. The PIMSM window appears.

Figure 288 The Dynamic Routing > PIMSM window



You can view the following information:

- **MRIB information inside PIM** – Displays the unicast routes used to reach other PIM rendezvous points, bootstrap routers, and multicast source.
- **Bootstrap routers** – Displays current bootstrap router and the set of rendezvous points.
- **Groups** – Displays information regarding joined groups: source (sender), rendezvous point, and flags.
- **Interfaces** – Displays each network interface configured for PIM.
- **Neighbors** – Displays information on neighboring PIM routers: interface, priority, and address.
- **RPs** – Displays the current rendezvous point routers: IP address, type, priority, and group range.
- **Multicast Forwarding Cache** – Displays the kernel's multicast forwarding table.

On this window, you can do the following:

- **Edit a configuration file.** Click **Edit** to open the configuration file using the Admin Console File Editor. Edit the file as needed and then save your changes. The firewall automatically restarts the pimd server. See [Configuring PIM-SM \(pimd\)](#) for more information.
- **View and compare files.** Select an option from the list and then click **Retrieve**. A pop-up window appears displaying the requested information. Close this pop-up to return to the main PIMSM window.

Dynamic routing in HA clusters

If you use dynamic routing in HA clusters, note the following considerations:

- [rip, ospf, and bgp] Add a router-id entry to the configuration file. You must specify an address, such as the cluster IP address.

For example, if 10.1.1.15 is your cluster IP address, configure the router-id like the following:

rip

```
router rip
network 10.1.1.15/32
```

ospf

```
router ospf
router-id 10.1.1.15

network 10.1.1.15/32 area 0
```

bgp

```
router bgp
bgp router-id 10.1.1.15
```

In neighbor bgp routers:

```
router bgp
neighbor 10.1.1.15 remote-as 6665
```

- [ospf and rip] If you specify the networks or interfaces by IP address, use the cluster IP address.

Troubleshooting dynamic routing issues

If you need to troubleshoot dynamic routing issues, you can use the following commands to enable debugging, and then either display or save the log files.

- 1 Using a command line session, log into the firewall and switch to the Admn domain by entering:

```
srole
```

- 2 Telnet into the appropriate dynamic routing server by entering one of the following command:

- To access ospfd, enter:

```
telnet localhost ospfd
```

- To access bgpd, enter:

```
telnet localhost bgpd
```

- To access ripd, enter:

```
telnet localhost_n ripd
```

where *n* = the burb index of the burb used as the source burb in the enabled ripd rule. Use `cf burb query` to look up a burb's index. It is also listed on the **Network > Burb Configuration** window as the ID.

A password prompt appears.

- 3 Enter **zebra**.

A `ripd>` prompt appears.

Note: The prompt will reflect the server you logged into.

- 4 Enable the full command set by entering:

```
ripd>en
```

The prompt changes to `ripd#` to indicate that the full command set is enabled.

Note: Enabling debugging at this prompt turns debugging on temporarily. To making debugging persistent, enter the **conf t** command before entering the debug commands.

- 5** Set the debug parameters by entering one or more commands similar to these examples:

```
ripd#debug protocol event
ripd#debug protocol packet [recv|send] [detail]
ripd#debug protocol zebra
```

where *protocol* is rip (case sensitive).

See the online help or Quagga documentation for ospf and bgp commands and for additional debugging flags.

- 6** View the log information in the current window by entering:

```
ripd#term monitor
```

To stop writing debug statement to the current window, enter:

```
ripd#term no monitor
```

- 7** [Optional] To save the log information to a log file, you can edit the configuration file directly and add this line:

```
log file filename
```

The default path for the log file is `/var/run/quagga`. To save the file in a different location, specify the entire path as part of the file name.

If you misconfigure your routing tables, you will need to disable ripd and make corrections to the tables and then restart ripd, either by writing the file changes or saving the configuration file using the Admin Console File Editor. Before restarting ripd, enter the following command at a UNIX prompt to flush the routing tables of all gateways.

```
route flush
```


18 DNS (Domain Name System)

Contents

[What is DNS?](#)

[Configuring transparent DNS](#)

[Configuring firewall-hosted DNS servers](#)

[Reconfiguring DNS](#)

[Manually editing DNS configuration files](#)

[DNS message logging](#)

What is DNS?

The domain name system (DNS) is a service that translates host names to IP addresses, and vice versa. DNS is necessary because while computers use a numeric addressing scheme to communicate with each other, most individuals prefer to address computers by name. DNS acts as the translator, matching computer names with their IP addresses.

Much of the traffic that flows into and out of your organization must at some point reference a DNS server. In many organizations this server resides on a separate, unsecured computer. The Forcepoint Sidewinder provides the additional option to host the DNS server directly on the firewall, eliminating the need for an additional computer.

The Sidewinder offers two main DNS configurations:

- **Transparent DNS** – Transparent DNS is designed for simple DNS configurations. The DNS server is on a separate computer, and DNS requests are proxied through the firewall. It is the default DNS configuration for a newly installed Sidewinder. See [About transparent DNS on page 470](#).
- **Firewall-hosted DNS** – Firewall-hosted DNS represents a more complex DNS configuration that uses the integrated Sidewinder DNS server. See [About firewall-hosted DNS on page 471](#).

Transparent DNS is the default configuration, created during initial configuration using the Quick Start Wizard. If you want to make changes to your existing DNS configuration, you can use one of two methods:

- **Admin Console** – Select **Network > DNS** to view and modify DNS settings. You can also click the **Reconfigure DNS** button to completely reconfigure your DNS settings. See the following for details:
 - [Configuring transparent DNS on page 474](#)
 - [Reconfiguring DNS on page 490](#)

Note: Using the Admin Console to modify your DNS configuration will remove any comments you may have manually inserted into the DNS configuration files.

- **Manual** – You can manually edit the DNS configuration files. This should only be attempted by highly skilled DNS administrators. See [Manually editing DNS configuration files on page 495](#) for details.

Note: An excellent source of information on DNS is the Internet Software Consortium web site at www.isc.org. The book DNS and BIND, by Albitz & Liu (O'Reilly & Associates, Inc.) is also a popular reference.

About transparent DNS

Transparent DNS represents a simplified DNS configuration. When transparent DNS is configured for the Sidewinder, DNS traffic passes transparently through the firewall using a proxy. The firewall uses proxy rules that pass all DNS traffic by proxy to its appropriate burb. DNS requests are then handled by the remote name servers. Other machines do not “see” the Sidewinder, which means there is minimal disruption to your current DNS configurations throughout your network.

Configuring transparent DNS requires specifying the IP address of one or more remote DNS servers. (Alternative server addresses may be used for redundancy.) If a customer is using NAT through the firewall, they should also have an additional DNS server on the outside of their network. The external DNS server handles the external zones of your network and its addresses. This configuration allows you to control which addresses are visible to the outside world.

There are two transparent DNS configuration options:

- **Single server** – The DNS traffic is proxied through the firewall. This configuration is generally used when you plan to use your existing DNS server. If you are using a single internal DNS server, external users have proxied access to your DNS server. External hosts are unaware that the firewall is “transparently” passing the DNS traffic.
- **Two servers** – The DNS traffic is proxied through the firewall, with a remote DNS server communicating with each interface. DNS queries are generally handled by both your internal DNS server and your external ISP. This configuration is more secure than using a single name server because your external server can limit access to your internal naming system. External hosts are unaware that the firewall is “transparently” passing the DNS traffic.

Transparent DNS is the default configuration on a newly installed Sidewinder. One server or two server DNS depends on your entries in the Quick Start Wizard. If you want to change your DNS configuration, see [Reconfiguring DNS on page 490](#).

Note: Transparent DNS is designed for simple DNS configurations. Complex DNS configurations may require DNS services to be hosted directly on the firewall.

About firewall-hosted DNS

Firewall-hosted DNS represents a more complex DNS configuration that uses the integrated Sidewinder DNS server. When configured for hosted services, DNS servers run directly on the firewall. This places the DNS server(s) on a hardened operating system, preventing attacks against these servers from penetrating your network.

You can configure firewall-hosted DNS to use a single server or split servers:

- **Hosted single server DNS** – In a firewall-hosted single server configuration, one DNS server is hosted on the firewall. That server handles all DNS queries. The server is protected by the Sidewinder hardened OS, preventing attacks from penetrating your network. A single server configuration is generally used when you have no concerns for keeping your internal network architecture hidden, such as when your firewall is acting as an “intrapwall” between two sets of private addresses. External hosts will need to be reconfigured to point to the Sidewinder server.
- **Hosted split server DNS** – In a firewall-hosted split server configuration, two DNS servers are hosted on the firewall: one server (the external name server) is bound to the external burb and the other server (the “unbound” name server) is available for use by all internal burbs. Both servers are protected by the Sidewinder hardened OS, which is able to prevent attacks against them from penetrating your network.

We recommend splitting the Sidewinder DNS servers when using hosted DNS. This configuration offers a good security benefit because the external burb of the firewall hides the DNS entries on the internal server from those who only have access to the external burb.

Designating an authoritative server

If your site has multiple internal domains, and there are name servers for each of these domains, the Sidewinder must be designated as an authoritative name server for *all* of the internal domains (the internal name servers also may be authoritative for one or more of the internal domains). This must occur regardless of whether the firewall is a master or a slave name server. The firewall must be an authoritative name server for all internal domains so that it can resolve queries for the internal domains. The firewall will otherwise automatically forward these internal name queries to the Internet, and the query will not be resolved.

In split DNS mode, if a DNS name occurs in the database of both servers, the name will resolve differently depending on the server that is queried. This occurs when the firewall is authoritative for the same domain both internally and externally. Because of this issue, if you try to access the Internet side of the firewall from an internal workstation you must use the appropriate machine name. For example, if the name of your firewall is *chloe*, then use the machine name *chloe-Internet*. This entry is automatically created during installation.

For more information on DNS, see *DNS and BIND* by Albitz & Liu, 3rd edition (O'Reilly).

Using master and slave servers in your network

In a hosted DNS configuration, the Sidewinder requires information about your DNS authority. Generally, there should be only one *master* name server for any fully qualified domain (such as *nyc.example.com*), also called a *zone*. There may be many *slave* servers, for redundancy and better performance, but they derive their information from the one master for each domain.

Typically, a company will use two or more DNS servers to provide domain name service to their customers. This provides for load balancing and redundancy. When more than one DNS server is used, the local administrator designates one DNS server to host the *master* zone files. The other DNS servers are slave servers that merely retrieve copies of the zone files from the master server. To outside users there is no indication or need to know about which of the multiple servers is the master. They all provide equally authoritative answers to all queries. The designation of which DNS server will be the master is only significant to the DNS administrator, because changes are made only at the master DNS server and not at the individual slave servers.

Note: When DNS servers are in an HA cluster, We recommend configuring the Sidewinder name servers as DNS slaves for authoritative zones. This allows the Master DNS servers to update both firewalls in the HA cluster. If you do not configure the Sidewinder name servers as DNS slaves for authoritative zones, DNS changes will be made to the secondary firewall with the next policy push.

Configuring clients and hosts to use firewall-hosted DNS

If you use firewall-hosted DNS, computers in your network must be configured to point to the appropriate DNS servers on the Sidewinder.

- Internal computers going through a proxy transparently to access the Internet must be configured to direct DNS queries in either of these ways:

- If you have internal name servers, the client computers must point to one or more of these name servers. The internal name servers should be authoritative for the internal domain, and should be configured to forward DNS queries to the Sidewinder.
- Reference the firewall on the client computers. For example:
 - In a UNIX system, enter the IP address of the Sidewinder's DNS server in the `/etc/resolv.conf` file.
 - In a Windows system, enter the IP address of the Sidewinder's DNS server in the TCP/IP Properties window.
- If you are using hosted split server DNS, external hosts must be configured to point to the external burb of the Sidewinder DNS servers.
- If you are hosting your own domain, your domain records can be configured to use the external Sidewinder DNS server as an authoritative name server for your domain. This is generally done with your domain registrar.

Enabling and disabling hosted DNS servers

When you configure firewall-hosted DNS services, the Sidewinder will use either one or two DNS servers. The DNS server(s) start automatically when you boot the firewall.

You can manually disable a DNS server on the Server Configuration tab of the DNS window by clearing the **Enable [Unbound/Internet] Domain Name Server** check box.

Keep the following points in mind, however, if you decide to disable a firewall-hosted DNS server:

- **If you have one DNS server**

In this situation, the server is known as an *unbound DNS server*. If you disable the DNS server, only connections that use IP addresses will still work; those that use host names will not.

- **If you have two DNS servers**

This situation is also known as *split DNS mode*. Note the following:

- If you disable the Unbound DNS server, connections that use IP addresses will still work; those that use host names will not.
- If you disable the Internet server, external connections that require host names will not work unless the name is already cached (saved) in the unbound name server's database. Connections that use IP addresses will work. E-mail will be placed in a queue since IP addresses cannot be resolved.
- If you disable both name servers, connections will work only if they use IP addresses rather than host names. Also, mail will not work and other errors will happen as other parts of the system attempt to access the network by name.

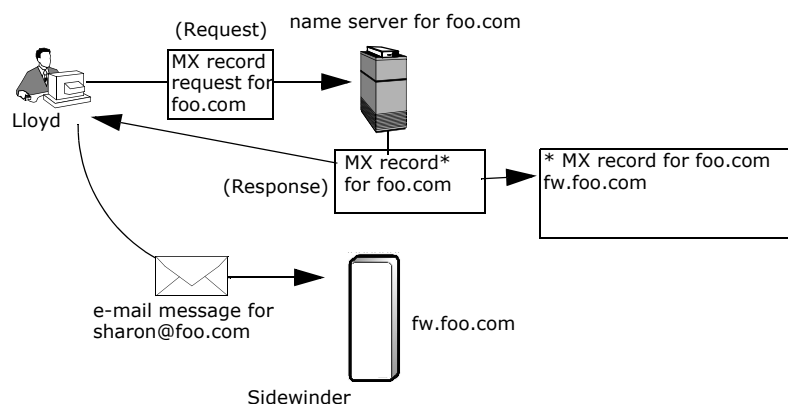
In either case, once you disable a server, the server will remain disabled until you enable it again.

Using hosted sendmail with firewall-hosted DNS

If you use hosted sendmail, you need to create mail exchanger (MX) records when you set up firewall-hosted DNS services for your site. MX records advertise that you are accepting mail for a specific domain(s). If you do not create an MX record for your domain, name servers and users on the Internet will not know how to send e-mail to you. When an e-mail message is sent from a site on the Internet, a DNS query is made in order to find the correct mail exchange (MX) host for the destination domain. The sender's mail process then sends the e-mail to the MX host. The firewall, through the use of mailertables, will forward the mail to the internal mail process, which in turn will forward it to the internal mail host. See [Editing sendmail files on Sidewinder on page 502](#) for more information on mailertables.

Consider the example shown in the figure below. Someone in the Internet, Lloyd, wants to send one of your users, Sharon, an e-mail message, but all Lloyd knows is Sharon's e-mail address: `sharon@foo.com`. The mailer at Lloyd's site uses DNS to find the MX record of `foo.com`. Lloyd's message for Sharon is then sent to the mailhost listed in the MX record for Sharon's site.

Figure 289 Mail exchanger example



A master name server stores and controls your site's MX records. The master name server may be in the external burb of your firewall, or on a host outside of your network (for example, your Internet service provider). If your firewall controls the master name server, then you can make any necessary changes to your MX records; if another host controls your master name server, then changes have to be made on that host. For more information on MX records, see Chapter 5 of *DNS and Bind* by Albitz & Liu.

For information on creating MX records using the Admin Console, see [Configuring the Master Zone Attributes tab on page 482](#).

More points about firewall-hosted DNS

Listed below are some additional points about running DNS on your Sidewinder:

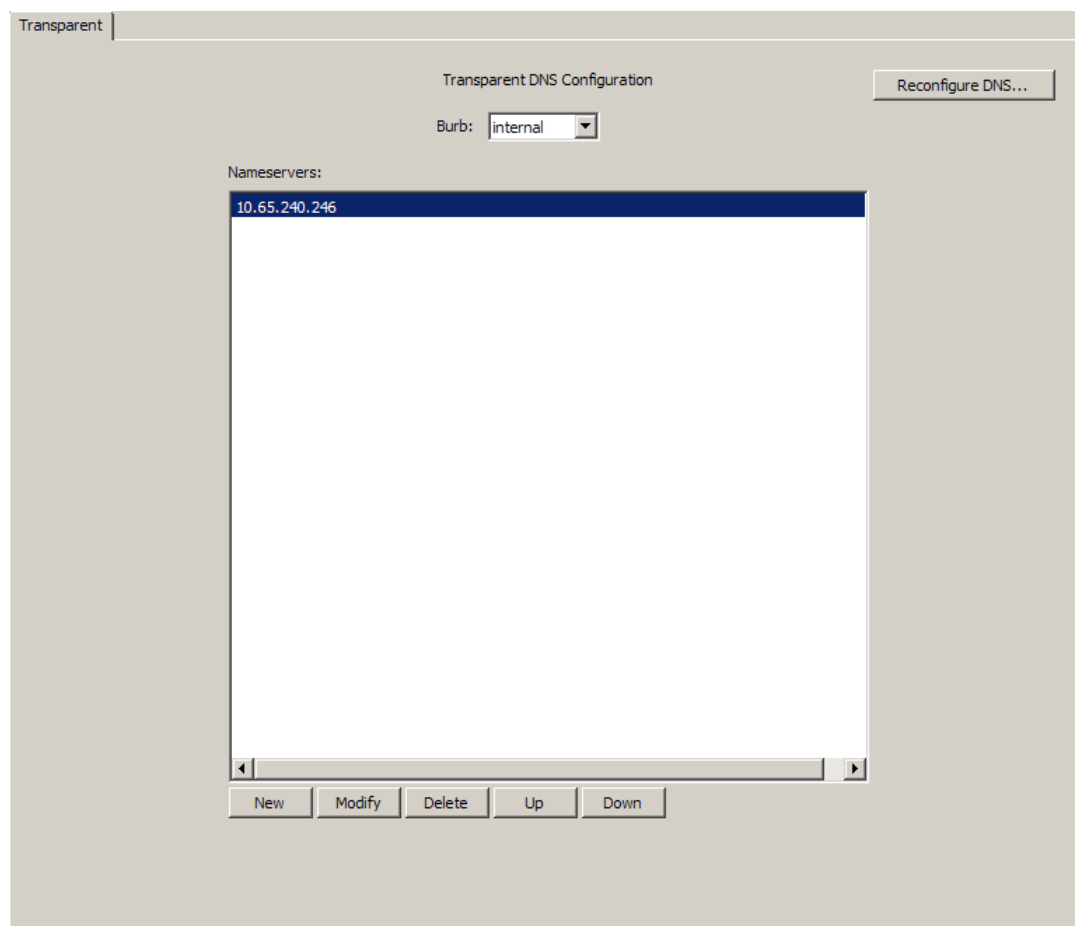
- The Sidewinder uses Berkeley Internet Name Domain (BIND 9).
- The configuration files for the unbound and the Internet name servers are `/etc/named.conf.u` and `/etc/named.conf.i`, respectively. The configuration files specify corresponding directories: `/etc/namedb.u` and `/etc/namedb.i`. When you boot your firewall, the name server daemon (`named`) is started. The `/etc/named.conf.u` and `/etc/named.conf.i` files specify whether the firewall is a master or a slave name server and list the names of the files that contain the DNS database records.
- If you choose to configure the firewall as a master name server on either the unbound (internal) or Internet (external) side, you can modify the `/etc/namedb.u/domain-name.db` and `/etc/namedb.i/domain-name.db` files (where *domain-name* = your site's domain name). You can add the information that is being advertised for these zones.
- The firewall contains a non-blocking DNS resolver to support reverse IP address look-ups in the active proxy rule group, and name-to-address look-ups in the various proxies. The relevant resolver library calls are `gethostbyname()` and `gethostbyaddr()`. The non-blocking DNS resolver provides a small number of DNS resolver daemons (`nbresd`) that are handed queries to resolve on behalf of the client.

Configuring transparent DNS

If you have configured DNS to use transparent services, you can add, modify, or delete transparent name servers. Select **Network > DNS**. The following window appears:

Note: If you want to completely reconfigure your existing DNS configuration (for example, change from transparent DNS to firewall-hosted DNS or vice versa), you must use the Reconfigure DNS window. See [Reconfiguring DNS on page 490](#) for details.

Figure 290 Transparent DNS Configuration window



Use this window to configure name servers for transparent DNS services. You can specify the burb to which the name servers will be assigned from the **Burb** drop-down list. You can assign and order DNS servers for any configured burb. The order in which the servers appear indicates the order in which the Sidewinder queries them.

- To add a new name server to the list, click **New**. To modify a name server, select the name server and click **Modify**.
- To change the name servers' order, select a name server and click the **Up** and **Down** buttons as appropriate.
- To delete a name server, select the name server and click **Delete**.

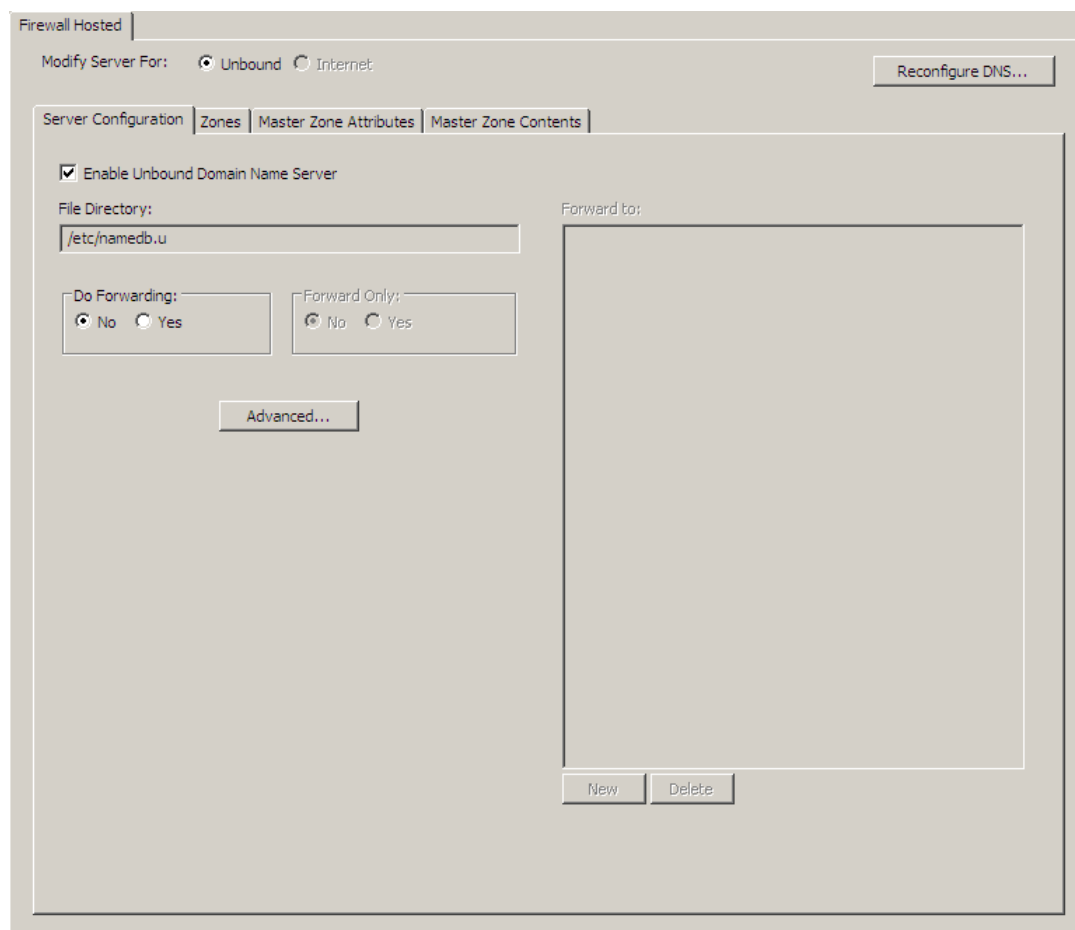
Configuring firewall-hosted DNS servers

If you configure DNS to use firewall-hosted services (single or split), you can define various name server information.

Note: If you want to completely reconfigure your existing DNS configuration (for example, change from transparent DNS to firewall-hosted DNS or vice versa), you must use the Reconfigure DNS window. See [Reconfiguring DNS on page 490](#) for details.

Select **Network > DNS**. The DNS window contains four tabs that allow you to define specific name server information.

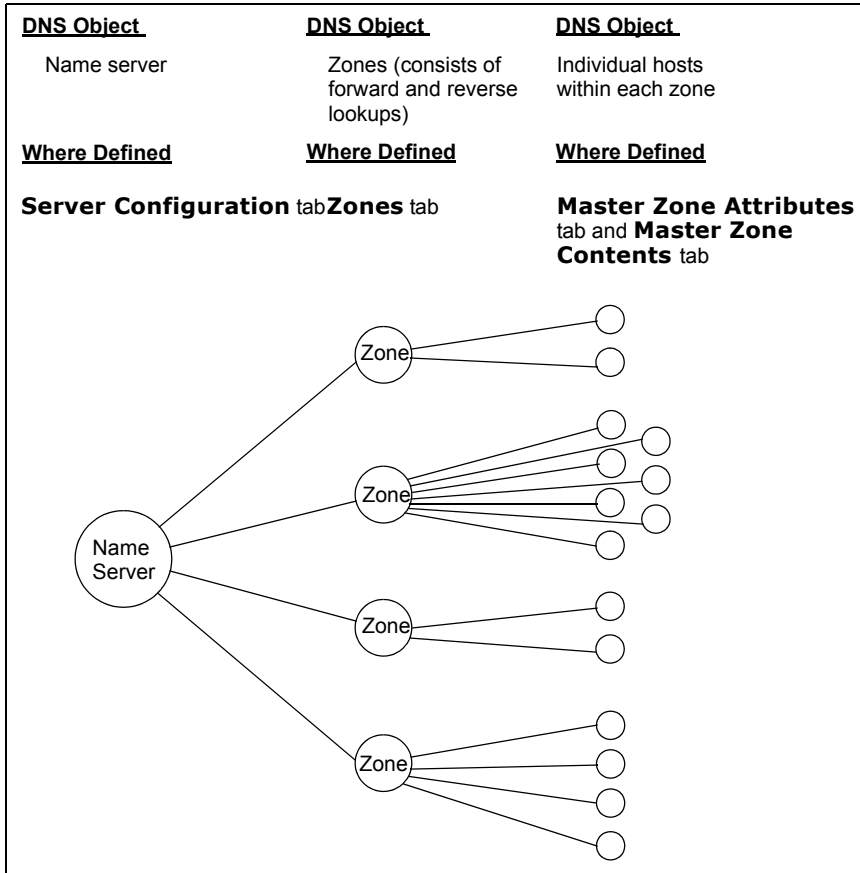
Figure 291 Firewall Hosted DNS window



- The **Server Configuration** tab is used to configure general information about a name server. See [Configuring the Server Configuration tab on page 477](#) for details.
- The **Zones** tab defines each of the master and slave zones associated with the selected name server. See [Configuring the Zones tab on page 480](#) for details.
- The **Master Zone Attributes** tab is used to configure attributes for each *master* zone defined on the **Zones** tab. See [Configuring the Master Zone Attributes tab on page 482](#) for details.
- The **Master Zone Contents** tab defines the hosts associated with each *master* zone defined on the **Zones** tab. See [Configuring the Master Zone Contents tab on page 487](#) for details.

The figure below illustrates the different DNS objects you can configure, how they relate to each other, and which tab is used to configure each object.

Figure 292 DNS objects and the tab used to configure each object



Configuring the Server Configuration tab

The **Server Configuration** tab is used to define configuration settings for the selected name server.

Figure 293 Firewall Hosted: Server Configuration tab

The screenshot shows the 'Server Configuration' tab of a configuration window. At the top are tabs: 'Server Configuration', 'Zones', 'Master Zone Attributes', and 'Master Zone Contents'. The 'Server Configuration' tab is active. It contains a checkbox 'Enable Unbound Domain Name Server' which is checked. Below it is a 'File Directory:' field with the value '/etc/namedb.u'. To the right is a 'Forward to:' list box containing four entries of '127.1.0.1', with the top entry highlighted. Below the 'File Directory' field are two groups of radio buttons: 'Do Forwarding:' with 'No' and 'Yes' (selected), and 'Forward Only:' with 'No' and 'Yes' (selected). Below these is an 'Advanced...' button. At the bottom right of the list box are 'New' and 'Delete' buttons.

Use this tab to define alternate name servers that will be contacted if a query cannot be resolved by the selected name server. The alternate name servers are called *forwarders*. This window is also used to define advanced configuration settings for the name server. To modify the Server Configuration tab:

Note: To completely reconfigure your DNS settings (for example, change from firewall-hosted single server to split server), click Reconfigure DNS. See [Reconfiguring DNS on page 490](#) for details.

- 1 In the **Modify Server For** field, select the name server that you want to modify. (The Internet server is available only if you are using two servers.)
- 2 [Conditional] If you want to disable the selected name server, clear the **Enable Unbound/Internet Domain Name Server** check box. (The Internet Domain Name Server is available only if two servers are defined.)

See [Enabling and disabling hosted DNS servers on page 472](#) for information about the effects of enabling or disabling the servers.

Note: The **File Directory** field displays the location of the files used to store information about the selected server. This field cannot be modified.

- 3 In the **Do Forwarding** field, specify whether the name server will forward queries it cannot answer to another name server. In a split DNS configuration, when modifying the unbound name server this field will default to **Yes** and will forward these unresolved queries to the Internet server (127.x.0.1, where x = the external [or Internet] burb number).

Forwarding occurs only on those queries for which the server is not authoritative and does not have the answer in its cache.


- 4 [Conditional] If you selected **Yes** in the previous step, configure the **Forward Only** field. Specify the following:
 - If you select **Yes**, the name server will forward queries it cannot answer to the name servers listed in the **Forward To** list only. This is the default.

- If you select **No**, the name server forwards the query to the name servers listed in the **Forward To** list. If they cannot answer the query, the name server attempts to contact the root server.
- 5 In the **Forward To** field, specify the alternate name servers that will be used when attempting to resolve a query. This list is consulted only if **Yes** is selected in the **Do Forwarding** field. If multiple name servers are defined, the name servers are consulted in the order listed until the query is resolved. In a split DNS configuration, when modifying the unbound name server this list will by default contain four entries for the Internet name server (127.x.0.1, where x = the external [or Internet] burb number).
- Note:** If you are using a split DNS configuration, We strongly recommend against defining additional alternate name servers for the unbound name server. The Internet (or external) name server should be the only alternate name server defined in this situation.
- 6 To add another entry to the list of authorized name servers, click **New** under the **Forward To** list, then type the IP address of the alternate name server. The alternate name servers are consulted if the primary name server cannot resolve a query.
 - 7 To delete a name server from the **Forward To** list, highlight the name server you want to delete and click **Delete**.
 - 8 [Conditional] To modify an advanced configuration setting for the name server, click **Advanced**. For more information on modifying the Advanced Server Options window, see [About the Advanced Server Options window on page 479](#).
- Note:** Only experienced DNS administrators should modify an advanced configuration setting.
- 9 Save your changes. To configure additional name server information, see [Configuring the Zones tab on page 480](#).

About the Advanced Server Options window

Use this window to define some of the more advanced DNS name server options.

Figure 294 Server Configuration: Advanced Server Options window



- Do not change these options unless you are an experienced DNS system administrator.
- By default, the options on this window are disabled, meaning *there are no restrictions*. If your organization considers this to be a security risk, you should use these options to limit the amount of interaction this name server has with other devices. Use your organization's security policy as a guide.

To modify advanced server options:

- 1 To enable the **notify** option, select the corresponding check box. Enabling this option allows you to specify whether the master server will notify all slave servers when a zone file changes. The notification indicates to the slaves that the contents of the master have changed and a zone transfer is necessary.

If this field is not selected, the field defaults to **Yes**.

- 2 To enable the **allow-query** option, select the corresponding check box. Selecting this option affects who is able to query this name server. The options are the following:
 - If not selected, all requesters are authorized to query the name server. This is the default.
 - If selected and contains IP addresses, only the requesters defined in the allow-query list will be authorized to query this name server. Use the **New** and **Delete** buttons to modify this list.

Note: If you select this option, be sure to include all IP addresses that might need to query the server, such as the heartbeat burbs' IP addresses, loopback addresses, etc.

- 3** To enable the **allow-transfer** option, select the corresponding check box. Selecting this option allows you to limit who is authorized to request zone transfers from this name server.
- If not selected, all requesters are authorized to transfer zones from the name server. This is the default.
 - If selected and no IP addresses are added, no requesters will be authorized to transfer zones from this name server.
 - If selected and contains IP addresses, only the requesters defined in the allow-transfer list will be authorized to transfer zones from this name server. Use the **New** and **Delete** buttons to modify this list.
- 4** Click **OK** to save your changes.

Configuring the Zones tab

A DNS server is responsible for serving one or more *zones*. A zone is a distinct portion of the domain name space. A zone consists of a domain or a subdomain (for example, example.com or sales.example.com). Each zone can be configured as either a master, slave or forward zone on this name server.

Figure 295 Firewall Hosted: Zones tab

The screenshot shows the 'Zones' tab in a configuration window. The window has four tabs: 'Server Configuration', 'Zones', 'Master Zone Attributes', and 'Master Zone Contents'. The 'Zones' tab is active. On the left, there is a list of zones: '248.65.10.in-addr.arpa' (selected), 'example.net', '0.2.127.in-addr.arpa', '0.3.127.in-addr.arpa', '0.0.127.in-addr.arpa', and '0.255.239.in-addr.arpa'. Below this list are 'New' and 'Delete' buttons. On the right, there are two main sections. The top section has 'Zone Type:' set to 'Master' and 'Zone File Name:' set to '248.65.10.db'. Below this is a 'Master Servers:' section with an empty list and 'New' and 'Delete' buttons. To the right of that is a 'Related Zones:' section with a list containing 'example.net'. At the bottom right is an 'Advanced...' button.

Use this tab to define zone information about the name server. Follow the steps below.

Note: To completely reconfigure your DNS settings (for example, change from firewall-hosted single server to split server), click **Reconfigure DNS**. See [Reconfiguring DNS on page 490](#) for details.

- 1 In the **Modify Server For** field, select the name server that you want to modify.
- 2 The **Zones** list defines the zones for which the name server is authoritative. This list initially contains a zone entry for each domain and each network interface defined to the firewall. You can add or delete zone entries as follows:
 - To add a new zone to the list, click **New** and type the name of the forward or reverse zone you want to add to the list.
 - To delete a zone, highlight a zone and click **Delete**.

We strongly recommend against deleting or modifying the following entries:

- Any 127 reverse zones (for example, *0.1.127.in-addr.arpa*). These zones represent local loopback addresses and are required.
- The zone with *0.255.239.in-addr.arpa* in its name. This zone provides multicast support for the Sidewinder failover feature.

There can be two different types of entries in the **Zone** list:

- Reverse zones (for example, *4.3.in-addr.arpa*): This format indicates the entry provides reverse lookup functions for this zone.
- Forward zones (for example, *example.com*): This format indicates the entry provides forward lookup functions for this zone.

The **Related Zones** list displays the zones that are related to the selected zone. For example, if a forward zone is selected, the related reverse lookup zones are displayed. This list cannot be modified.

- 3 In the **Zone Type** field, specify whether the selected zone is a master zone, a slave zone, or a forward zone, as follows:
 - **Master** – A master zone is a zone for which the name server is authoritative. Many organizations define a master zone for each sub-domain within the network. Administrators should only make changes to zones defined as a master.
Tip: You should consider defining a matching reverse zone (an **in-addr.arpa** zone) for each master zone you configure.
 - **Slave** – A slave zone is a zone for which the name server is authoritative. Unlike a master zone, however, the slave zone's data is periodically transferred from another name server that is also authoritative for the zone (usually, the master). If you select *Slave*, the **Master Servers** field becomes active. Be sure to use the **Master Servers** field to define the name server(s) that will provide zone transfer information for this slave zone. Administrators should not make changes to zones defined as a slave.
Caution: When changing a zone from slave to master, the Admin Console changes the slave file into a master file and the file becomes the lookup manager for the zone. The DNS server will have no problems understanding and using the new master file. For large zones (class A or B), however, this file may become too complex to be managed properly using the Admin Console. We recommend either leaving large zones as slaves on the firewall or manually modifying these files.
 - **Forward** – A forward zone allows you to specify that queries for names in the zone are forwarded to another name server.
- 4 In the **Zone File Name** field, specify the name of the file that is used to store information about this zone. The file is located in the directory specified in the **File Directory** field on the Server Configuration tab. We do not recommend changing this name.

- 5 Conditional] When **Zone Type** is **Forward**, the **Forwarders** list defines one or more forwarders for a zone. You can add or delete forwarder entries as follows:
 - To add a new forwarder to the list, click **New** and type the IP address.
 - To delete a forwarder, select that item and click **Delete**.
- 6 [Conditional] When the **Zone Type** is **Slave**, the **Master Servers** list defines one or more master name servers that are authorized to transfer zone files to the slave zone. You can add or delete server entries as follows:
 - To add a new master server to the list, click **New** and type the IP address.
 - To delete a master server, highlight a server and click **Delete**.
- 7 [Conditional] To modify an advanced configuration setting for the selected zone, click **Advanced**. For more information on modifying the Advanced Server Options window, see [About the Advanced Zone Configuration window on page 482](#).

Note: Only experienced DNS administrators should modify an advanced configuration setting.
- 8 Save your changes.

About the Advanced Zone Configuration window

Use the Advanced Zone Configuration window to define some of the more advanced zone configuration options. This window allows you to configure certain options specifically for the selected zone, overriding similar options that may be configured for the global name server (the Unbound or the Internet name server). Follow the steps below.

Note: Only experienced DNS administrators should modify an advanced configuration setting.

- 1 To enable the **notify** option, select the corresponding check box. Enabling this option allows you to specify whether the master server will notify all slave servers when the zone changes. The notification indicates to the slaves that the contents of the master have changed and a zone transfer is necessary. The name servers that are notified are those defined in the **Zone NS Records** field on the **Master Zone Attributes** tab.

If this field is not selected, the field defaults to **Yes**.

- 2 To enable the **allow-query** option, select the corresponding check box. Selecting this option affects who is able to query this zone. The options are the following:
 - If not selected, all requesters are authorized to query the zone. This is the default.
 - If selected and contains IP addresses, only the requesters defined in the allow-query list will be authorized to query this zone. Use the **New** and **Delete** buttons to modify this list.

Note: If you select this option, be sure to include all IP addresses that might need to query the zone, such as the heartbeat burbs' IP addresses, loopback addresses, etc.

- 3 To enable the **allow-update** option, select the corresponding check box. Selecting this option allows you to specify from whom the zone will accept dynamic DNS updates. If this option is selected, only the hosts in the allow-update list are authorized to update this zone. This option is only valid for master zones. Use the **New** and **Delete** buttons to modify this list.


By default the **allow-update** option is not selected, meaning the server will deny updates from all hosts.

- 4 To enable the **allow-transfer** option, select the corresponding check box. Selecting this option allows you to limit who is authorized to request zone transfers from this zone.
 - If not selected, all requesters are authorized to transfer this zone from the name server. This is the default.
 - If selected and no IP addresses are added, no requesters will be authorized to transfer this zone from the name server.
 - If selected and contains IP addresses, only the requesters defined in the allow-transfer list will be authorized to transfer the zone from the name server. Use the **New** and **Delete** buttons to modify this list.

Configuring the Master Zone Attributes tab

Use the **Master Zone Attributes** tab to configure attributes for each master zone defined on the **Zones** tab. Slave zones are not included on this tab because you can only define attributes for those zones for which you are the master.

Figure 296 Firewall Hosted: Master Zone Attributes tab



Use this tab to define the attributes of each master zone defined for the selected name server. In particular, it defines the Name Server record(s) and the Start of Authority (SOA) record for each master zone. The window also enables you to define Mail Exchanger (MX) records for those entries that are forward lookup zones. Follow the steps below.

Note: To completely reconfigure your DNS settings (for example, change from firewall-hosted single server to split server), click **Reconfigure DNS**. See [Reconfiguring DNS on page 490](#) for details.

- 1 In the **Modify Server For** field, select the name server that you want to modify.

The **Master Zones** list defines the zones for which the name server is master. A plus sign (+) will appear in front of any forward lookup zone that contains one or more sub-domains. Click the plus sign to view the sub-domains.

To modify an entry in the list, click the entry name. A menu of options used to characterize the selected entry is presented on the right side of the window.

Note: The **Forward Lookup Zone Name/Reverse Lookup Zone Name** field displays the full zone name associated with the entry selected in the **Master Zones** list.

- 2 To modify the **Zone SOA** tab, click the tab and follow the sub-steps below. The fields on the Zone SOA tab collectively define one Start Of Authority (SOA) record. An SOA record controls how master and slave zones interoperate.

Figure 297 Master Zone Attributes: Zone SOA tab

Zone SOA | Zone Records

Zone Start of Authority (SOA):

DNS Serial #:

DNS Contact:

Refresh:

Retry:

Expiration:


TTL:

(Time in seconds)

The **DNS Serial #** field displays the revision number of this SOA record. This field will increment by one each time you modify this zone. Slave zones use this field to determine if their zone files are out-of-date. You cannot modify this field. (See sub-[Step b](#) for more details.)

- a** In the **DNS Contact** field, specify the name of the technical contact that can answer questions about this zone. The name must be a fully-qualified name, with the @ character replaced by a period (for example, `hostmaster@example.com` becomes `hostmaster.example.com.`).
 - b** In the **Refresh** field, specify in seconds how often a slave will check this zone for new zone files. The slave uses the **DNS Serial #** value to determine if its zone files need to be updated. For example, if the slave's DNS serial number is 4 and the master zone's DNS serial number is 5, the slave knows that its zone files are out-of-date and it will download the updated zone files. Values must be positive integers.
 - c** In the **Retry** field, specify in seconds how long a slave should wait to try another refresh following an unsuccessful refresh attempt. Values must be positive integers.
 - d** In the **Expiration** field, specify in seconds how long a slave can go without updating its data before expiring its data. For example, assume you set this value to 604800 (one week). If the slave is unable to contact this master zone for one week, the slave's resource records will expire. After expiration, queries to that zone will fail. Values must be positive integers.
 - e** In the **TTL** field, specify the time to live (TTL) value. This value defines how long a resource record from this zone can be cached by another name server before it expires the record. The value specified here is used as the default in records that do not specify a TTL value. Values must be positive integers.
 - f** To add a sub-domain to the selected zone, click **Add Sub...**. This button is only available if a forward lookup zone is selected in the **Zones** list. For information on adding a sub-domain, see [Adding a forward lookup sub-domain on page 485](#).
 - g** To delete a sub-domain from the selected zone, click **Delete Sub...**. This button is only available if a forward lookup zone is selected in the **Zones** list.
- 3** To modify the **Zone Records** tab, click the tab. This tab contains NS (Name Server) and MX (Mail Exchange) records for forward zones. This tab contains only NS Records for reverse zones.

Figure 298 Master Zone Attributes: Zone Records tab



- The **Zone NS Records** table contains DNS NS records that indicate what machines will act as name servers for this zone. By default the table contains an entry for the machine you are currently using.
 - To add a Zone NS Records entry, click **New**. In the **NS Record** field, type the domain name associated with this NS record. The name must be a fully qualified name and must end with a period. The name you specify should be a pre-existing domain name that maps to a valid IP address.
 - To delete a Zone NS Records entry, select the entry and click **Delete**.

If this zone is configured to notify all slave servers when a zone file changes, the notify commands are sent to all NS hosts specified here. (See [About the Advanced Zone Configuration window on page 482](#) for a description of the **notify** field.)

- The **Zone MX Records** list is available only if the selected zone entry is a forward lookup entry. It is used to specify entries in the Mail Exchangers table for the selected zone. The Mail Exchangers table contains DNS MX records that indicate what machines will act as mail routers (mail exchangers) for the selected domain.
 - To add a Zone MX Records entry, click **New**. Type a fully qualified host name, and a priority level for this record. Valid values are 1–65535. The lower the value, the higher the priority.
 - To delete a Zone MX Records record entry, select the entry and click **Delete**.
- The **Zone A Record** field is available only if the selected zone entry is a forward lookup entry. It defines a DNS A record (an Address record) for the zone itself. A DNS A record is used to map host names to IP addresses. The address you specify must be entered using standard dotted quad notation (for example 172.14.207.27).
- The **TXT Record** field is available if the selected zone entry is a forward lookup entry. This optional field allows you to enter comments or additional information about this zone, such as sender id information.

4 Save your changes.

Adding a forward lookup sub-domain

Use this window to add a forward lookup sub-domain to the selected forward lookup zone. By adding a sub-domain you are delegating authority for a portion of the parent domain to the new sub-domain. Follow the steps below.

- 1** In the **Forward Sub-Domain Name** field, type the name of the sub-domain. Do not type a fully qualified name. For example, assume you have a domain named *example.com* that contains a sub-domain named *west*. You would type *west* in this field rather than *west.example.com*.
- 2** In the **Sub-Domain NS Records** field, specify entries in the Name Servers table for this sub-domain. The Name Servers table contains DNS NS records that indicate what machines will act as name servers for this sub-domain.
 - To add an NS Records entry, click **New**. In the **NS Record** field, type the domain name associated with this NS record. The name must be a fully qualified name and must end with a period. The name you specify should be a pre-existing domain name that maps to a valid IP address.
 - To delete an NS Records entry, select the entry and click **Delete**.
- 3** [Optional] In the **Sub-Domain MX Records** field, specify entries in the Mail Exchangers table for this sub-domain. The Mail Exchangers table contains DNS MX records that indicate what machines will act as mail routers (mail exchangers) for the sub-domain.
 - To add an MX Records entry, click **New**. Type a fully qualified host name, and a priority level for this record. Valid values are 1–65535. The lower the value, the higher the priority.
 - To delete an MX Records record entry, select the entry and click **Delete**.
- 4** Click **Add** to add the specified sub-domain. Click **Close** to exit the window.

Configuring the Master Zone Contents tab

The **Master Zone Contents** tab is used to define the hosts that are associated with each master zone.

When you select the **Master Zone Contents** tab, a window similar to the following appears.

Note: If you are adding a large number of hosts (hundreds or thousands) to a master zone, you may want to consider manually adding the required host information directly to the appropriate DNS files using one of the available editors on the firewall to save time. However, only experienced Sidewinder administrators should attempt this. (Using the manual method will still require you to manually define each host.)

Figure 299 Firewall Hosted: Master Zone Contents tab

The screenshot shows the 'Master Zone Contents' tab in a web-based configuration interface. The interface has a tabbed menu at the top with 'Server Configuration', 'Zones', 'Master Zone Attributes', and 'Master Zone Contents'. The 'Master Zone Contents' tab is active. On the left, under 'Master Zones:', there is a tree view showing a hierarchy: '/etc/namedb.u' expanded to 'example.net', which is further expanded to 'localhost'. Other zones listed include '10.65.248', '127.2.0', '127.3.0', '127.0.0', and '239.255.0'. On the right, under 'Address Record', there are input fields for 'Entry Name:' (localhost), 'A Record IP:' (127.2.0.1), 'CNAME Rec:', and 'TXT Record:'. Below these is a section for 'Entry MX Records:' with a large empty box and 'New' and 'Delete' buttons. To the right of the MX records box are two more input fields: 'HINFO - Type:' and 'HINFO - OS:'.

For each host you define in a forward lookup zone you should also create a matching entry in the associated reverse lookup zone. Follow the steps below.

Note: To completely reconfigure your DNS settings (for example, change from firewall-hosted single server to split server), click **Reconfigure DNS**. See [Reconfiguring DNS on page 490](#) for details.

- 1 In the **Modify Server For** field, select the name server that you want to modify.

The fields that are available on this tab will vary depending on whether a zone, a host in a forward lookup zone, or a host in a reverse lookup zone is selected.

- 2 [Conditional] If you are modifying a zone, do the following:

- a In the Master Zones area, select the zone you want to modify.
- b To add a host to the selected zone, click **Add Entry**.
 - If you are adding a host to a forward lookup zone, see [Adding a new forward lookup entry on page 489](#) for details.
 - If you are adding a host to a reverse lookup zone, see [Adding a new reverse lookup entry on page 489](#).
- c To delete a host from the selected zone, click **Delete Entry**. The **Hosts in Zone** field lists all the hosts currently defined within the selected zone. Select the host you want to delete and click **Delete Host**. You can only delete one host at a time.

3 [Conditional] If you are modifying a host in a reverse lookup zone, the following two fields appear:

- **Name (Host portion of IP)** – The field displays the host portion of either the IP address or of the fully-qualified domain name of this entry. You cannot modify this field. If you need to change the name, you must delete the entry from the list, then add the entry back using the new name.
- **Fully-Qualified Domain Name** – This field displays the domain name of the host. You can modify this field by typing in a new value. Be sure to type the fully-qualified domain name of the host.

Note: The **Name** field and the **Fully-Qualified Domain Name Entry** field collectively define a PTR Record for the selected reverse lookup zone. The PTR record is used in a Reverse Addresses table and maps an IP address to a host name.

4 [Conditional] If a host in a forward lookup zone is selected, the following fields appear:

- **Entry Name** – This field defines the host portion of the fully-qualified domain name of this entry.
- **A Record IP** – This field defines a DNS A record (an Address record), which is used to map host names to IP addresses. In this case the field displays the IP address of the selected host. You can modify this field by typing in a new value. The address you specify must be entered using standard dotted quad notation (for example 172.14.207.27).
- **CNAME Rec** – This field defines a DNS CNAME record, which is used to map an alias to its canonical name. The field, if populated, displays the name of the Canonical Record of the selected host. You can modify this field by typing in a new name. The name you specify must be entered using the fully qualified primary name of the domain.

Note: A host in a forward lookup zone requires either an A Record or a CNAME Record.

- **TXT Record** – This field allows you to enter comments or additional information about this zone, such as sender id information.
- **Entry MX Records** – This field is used to specify entries in the Mail Exchangers table for the selected host. The Mail Exchangers table contains DNS MX records that indicate what machines will act as mail routers (mail exchangers) for the selected host.
 - To add an MX Records entry, click **New**. Type a fully qualified host name, and a priority level for this record. Valid values are 1–65535. The lower the value, the higher the priority.
 - To delete an MX Records entry, select the entry and click **Delete**.
- **HINFO-Type** – This field provides information about a host's hardware type.
- **HINFO-OS** – This field provides information about a host's operating system.

Note: For security reasons, many organizations elect not to use the HINFO fields.

5 Save your changes.

Adding a new forward lookup entry

Use this window to define a new host for a forward lookup zone.

To add a forward lookup entry:

- 1** In the **Entry Name** field, specify the host portion of the fully-qualified domain name of this entry.
- 2** In the **A Record IP** field, specify a DNS A record (an Address record), which is used to map host names to IP addresses. The address you specify must be entered using standard dotted quad notation (for example 172.14.207.27). This field and the **CNAME Rec** field are mutually exclusive.
- 3** In the **CNAME Rec** field, specify a DNS CNAME record, which is used to map an alias to its canonical name. The name you specify must be entered using the fully-qualified primary name of the domain. This field and the **A Record IP** field are mutually exclusive.
- 4** [Optional] In the **TXT Record** field, enter comments or additional information about this zone, such as sender ID information.
- 5** [Optional] The **Entry MX Records** field lists entries in the Mail Exchangers table for this host. The Mail Exchangers table contains DNS MX records that indicate what machines will act as mail exchangers for the host.
 - To add an MX Records entry, click **New**. Type a fully qualified host name, and a priority level for this record. Valid values are 1–65535. The lower the value, the higher the priority.
 - To delete an MX Records record entry, select the entry and click **Delete**.
- 6** [Conditional] The **HINFO-Type:** field provides information about a host's hardware type.

Note: For security reasons, many organizations elect not to use the HINFO fields.
- 7** [Conditional] The **HINFO-OS** field provides information about a host's operating system.
- 8** Click **Add** to save the new entry. Click **Close** to exit this window.

Adding a new reverse lookup entry

Use this window to define a new host for a reverse lookup zone. Follow the steps below.

- 1** In the **Entry Name** field, specify the host portion of the IP address of this entry.
- 2** In the **Fully-Qualified Name Entry** field, specify the domain name of the host. Be sure to type the fully-qualified domain name of the host.

Note: The **Entry Name** field and the **Fully-Qualified Name Entry** field collectively define a PTR Record for the selected reverse lookup zone. The PTR record is used in a Reverse Addresses table and maps an IP address to a host name.
- 3** Click **Add** to save the new entry. Click **Close** to exit this window.

Reconfiguring DNS

The Reconfigure DNS window allows you to completely reconfigure DNS on your Sidewinder.

- Make sure you create a configuration backup before reconfiguring DNS.
- After using the DNS configuration utility, reboot the firewall.
- Any active DNS servers on the firewall will be disabled during the reconfiguration process.
- Any prior modifications you have made to your DNS configuration will be lost when you save your changes. You will need to re-apply the modifications.

Reconfiguring transparent DNS

To reconfigure DNS to use transparent services, select **Maintenance > Reconfigure DNS**. (You can also click the **Reconfigure DNS...** button on the DNS window.) The following window appears:

Figure 300 Reconfigure transparent DNS window

Firewall Hosted: Reconfiguring DNS

Reconfiguring DNS

Current DNS Configuration: One firewall-hosted server

New DNS Configuration: transparent

☐ Internal Name Server:

IP Address:

Alternate IP Address:

Burb: internal

☐ Internet Name Server:

IP Address:

Alternate IP Address:

Burb: external

OK Cancel Help

To reconfigure your DNS settings to use transparent DNS services:

- 1 In the **New DNS Configuration** drop-down list, select **Transparent**.
- 2 To configure the Sidewinder to use the internal name servers:
 - a Select the **Internal Name Server** check box.
 - b In the corresponding **IP Address** field, type the IP address of the name server located in the internal burb.
 - c [Optional] In the **Alternate IP Address** field, type the IP address of an alternate name server.
 - d In the **Burb** drop-down list, select your internal burb.

3 To configure the Sidewinder to use the external (Internet) name servers:

- a** Select the **Internet Name Server** check box.
- b** In the corresponding **IP Address** field, type the IP address of the name server located in the external (Internet) burb (that is, your ISP's name server).
- c** [Optional] In the **Alternate IP Address** field, type the IP address of an alternate name server.
- d** Save your DNS settings. A pop-up message appears informing you whether the reconfiguration was successful.


Note: The pop-up message that appears may contain additional information or warnings about your Sidewinder configuration. Please read this message carefully before you click **OK**.

4 Reboot the firewall: Select **Maintenance > System Shutdown**.

Reconfiguring single server hosted DNS

To reconfigure DNS to use single server hosted services, select **Maintenance > Reconfigure DNS**. (You can also click the **Reconfigure DNS...** button on the DNS window.) The following window appears:

Figure 301 Reconfiguring firewall Hosted (single server) DNS window



To reconfigure your DNS settings to use hosted single server DNS services:

- 1 In the **New DNS Configuration** drop-down list, select **Firewall Hosted**.
- 2 Select the **1 Server** radio button.
- 3 In the **Domain** field, verify that the correct domain name appears.
- 4 In the **Authority** field, select one of the following options:
 - **Master** – Select this option if the server you are defining will be a master name server. A master name server contains name and address information for every computer within its zone.
 - **Slave** – Select this option if the server you are defining will be a slave name server. A slave name server is similar to a master name server, except that it does not maintain its own original data. Instead, it transfers data from another name server.
- 5 [Conditional] If you selected **Slave** in the previous step, type the IP address of the master authority server in the **Master IP** field.
- 6 Save your DNS settings. A pop-up message appears informing you whether the reconfiguration was successful.

Note: The pop-up message that appears may contain additional information or warnings about your Sidewinder configuration. Please read this message carefully before you click **OK**.
- 7 Reboot the firewall: Select **Maintenance > System Shutdown**.

Reconfiguring split server hosted DNS

To reconfigure DNS to use split server hosted services, select **Maintenance > Reconfigure DNS**. (You can also click the **Reconfigure DNS...** button on the DNS window.) The following window appears:

Figure 302 Reconfiguring DNS: Firewall Hosted (split server) window

Firewall Hosted: Reconfiguring DNS

Reconfiguring DNS

Current DNS Configuration: One firewall-hosted server

New DNS Configuration: Firewall Hosted

☐ 1 Server
☒ 2 Servers

Server: Unbound

Domain: example.net

Authority:

☐ Master
☒ Slave

Master IP:

Server: Internet

Domain: example.net

Authority:

☐ Master
☒ Slave

Master IP:

OK Cancel Help

To reconfigure your DNS settings to use hosted split server DNS services:

- 1** In the **New DNS Configuration** drop-down list, select **Firewall Hosted**.
- 2** Select the **2 Servers** radio button.
- 3** To configure the **Unbound** server, do the following:
 - a** In the **Domain** field, verify that the correct domain name appears.
 - b** In the **Authority** field, select one of the following options:
 - **Master** – Select this option if the server you are defining will be a master name server. A master name server contains name and address information for every computer within its zone.
 - **Slave** – Select this option if the server you are defining will be a slave name server. A slave name server is similar to a master name server, except that it does not maintain its own original data. Instead, it transfers data from another name server.
 - c** [Conditional] If you selected **Slave** in the previous step, type the IP address of the master authority server in the **Master IP** field.

- 4** To configure the **Internet** server, do the following:
 - a** In the **Domain** field, verify that the correct domain name appears.
 - b** In the **Authority** field, select one of the following options:
 - **Master** – Select this option if the server you are defining will be a master name server. A master name server contains name and address information for every computer within its zone.
 - **Slave** – Select this option if the server you are defining will be a slave name server. A slave name server is similar to a master name server, except that it does not maintain its own original data. Instead, it transfers data from another name server.
 - c** [Conditional] If you selected **Slave** in the previous step, type the IP address of the master authority server in the **Master IP** field.
- 5** Save your changes to reconfigure your DNS settings. A pop-up message appears informing you whether the reconfiguration was successful.

Note: The pop-up message that appears may contain additional information or warnings about your Sidewinder configuration. Please read this window carefully before you click **OK**.

Manually editing DNS configuration files

You can manually edit the DNS configuration files.

- Files with a *u* extension are for the unbound nameserver, and files with an *i* extension are for the Internet nameserver.
- You should only edit zone files for a master name server. Never edit the slave name server files. The file names shown below are for a master name server.

To manually edit DNS configuration files:

- 1 From a console attached to the firewall, log in and enter **srole** to switch to the Admin domain.

The following two steps assume you have zone files named **domain.db** and **reverse.db** in your system. Substitute your file names as required.

- 2 Open the `/etc/namedb.u/domain.db` and `/etc/namedb.i/domain.db` files in a UNIX text editor and make the necessary changes.
- 3 Open the `/etc/namedb.u/reverse.db` and `/etc/namedb.i/reverse.db` files in a UNIX text editor and make the necessary changes.
- 4 Open the `/etc/named.conf.u` and `/etc/named.conf.i` files in a UNIX text editor and make the necessary changes.

Note: If you edit the `/etc/named.conf.*` files to change an existing master zone into a slave zone, you must also manually remove the old zone file in your `/etc/namedb.*` directories.

- 5 If you have added new files, you must change the files to the correct Type Enforcement types.

To do this, type the following command and insert the names of the file(s) you edited in steps 2, 3 and 4:

```
chtype DNSx:conf filename
```

- For non-Internet (unbound) burbs, in place of **x** type the identifier **u**.
- For the Internet burb, in place of **x** type the index number of the Internet burb. (Use the **region show** command to determine the index number.)

- 6 Increment the serial number after every change to the master files.
- 7 Enter the following command to restart DNS.

```
cf daemond restart agent=named_unbound
cf daemond restart agent=named_intenet
```

Note: Any files created by **named** daemons, such as zone backup files or query log files, have types of DNSu:file or DNSx:file.

- 8 Check `/var/log/daemon.log` for any errors.

DNS message logging

DNS messages, Type Enforcement errors and process limit errors are logged in the following locations on the Sidewinder.

- `/var/log/audit.raw`: Contains information in the Sidewinder audit format.
- `/var/log/daemon.log`: Contains traditional syslog format messages.

You can view the `audit.raw` file using the Audit windows in the Admin Console (See [Chapter 11, Auditing](#) for more information). The `daemon.log` file can be viewed using any text editor.

19 E-mail

Contents

[Overview of mail on Sidewinder](#)

[Setting up and reconfiguring mail](#)

[Understanding sendmail on Sidewinder](#)

[Editing sendmail files on Sidewinder](#)

[Configuring advanced sendmail features](#)

[Managing mail queues](#)

[Receiving mail sent by Sidewinder](#)

Overview of mail on Sidewinder

When you run mail on a network protected by the Forcepoint Sidewinder, you have two options for getting messages through the firewall: use the SMTP proxy or use the Sendmail® server hosted on the firewall. If you use the firewall-hosted Sendmail server, you can also use firewall's mail anti-virus service.

The two mail configuration options are described in the following sections.

About transparent mail (SMTP proxy)

This configuration option allows you to use transparent SMTP services (without sendmail processes running directly on the firewall). Transparent SMTP service indicates that all inbound and outbound mail passes by proxy through the firewall, just as other proxy traffic does.

When using transparent mail services, the following mail filtering features are available:

- Internal mail infrastructure masking
- Message size filtering
- Destination address filtering
- Server reply length checks
- Command filtering
- Header filtering
- Extension filtering

The Sidewinder SMTP proxy is best used as a frontline defense together with additional mail filtering devices. The SMTP proxy can stop a large portion of inbound spam while using a comparatively small amount of computational resources. As a result, the workload on dedicated mail filtering devices behind the firewall is substantially reduced.

Set up transparent mail by doing the following:

- 1 Make a configuration backup (**Maintenance > Configuration Backup**).
- 2 Select the transparent mail option. See [Setting up and reconfiguring mail](#).
- 3 Create a rule using a mail proxy service and a Mail (SMTP Proxy) Application Defense.
 - To configure a rule, see [Creating, modifying, and duplicating rules](#).
 - To configure a mail proxy service, see [Create and modify services](#).
 - To configure a Mail (SMTP Proxy) application defense, see [Creating Mail \(SMTP proxy\) Defenses](#).

About Sidewinder-hosted mail (sendmail)

This configuration option allows you to have two sendmail servers running directly on the firewall, each supported on its own burb: the external burb and one non-Internet burb that you choose. The firewall's sendmail servers will route mail through the firewall only for these two burbs.

This configuration protects your internal mailhost from malicious attacks, and offers a variety of additional mail-handling options. When using secure split mail services, the firewall's external sendmail server is the mail host to which all external SMTP hosts will connect. The firewall's internal sendmail server will connect with internal mail hosts in its same burb.

Your internal mail host must run mail software that can accept incoming messages from, and send outgoing messages to, your Sidewinder. This system might be running sendmail or some other mail package such as Microsoft Exchange or cc:Mail with a Simple Mail Transport Protocol (SMTP) gateway.


When using hosted sendmail, you also have the option of purchasing the anti-virus/anti-spyware service that runs directly on Sidewinder.

For instructions on setting up the hosted mail option, see [Setting up and reconfiguring mail](#). From instructions on configuring anti-virus services, see [Chapter 6, Content Inspection](#). Using hosted sendmail services is a very secure option, but does require more mail administration on the Sidewinder. Read the rest of this chapter for important management instructions.

Setting up and reconfiguring mail

The Reconfigure Mail window is used to configure both firewall mail options. In the Admin Console, select **Maintenance > Reconfigure Mail**. (You can also access this window within the sendmail server's Properties window.) The Reconfigure Mail window appears.

Figure 303 Reconfigure Mail window



The screenshot shows the 'Reconfigure Mail' window with the following fields:

- Current SMTP Mode: Secure Split SMTP Servers
- New SMTP Mode: [Dropdown menu]
- Internal SMTP Burb: [internal] [Dropdown menu]
- Internal SMTP Mail Server: [mailhost.example.net] [Text field]

Use this window to set your initial mail configuration or reconfigure your existing mail configuration. However, before you make any changes you should be aware that if you manually edited any sendmail configuration files, changing your mail configuration in the Reconfigure Mail window will overwrite the changes you made.

To establish or change your mail configuration:

- 1 Verify that DNS is configured correctly.
- 2 Make a configuration backup before you change your mail configuration: Select **Maintenance > Configuration Backup**.
- 3 On the Reconfigure Mail window, expand the **New SMTP Mode** drop-down list and select the mail configuration mode you want to configure. The current mode is listed in the **Current SMTP Mode** field. The following options are available:
 - **Transparent** – Use this option when you want to pass mail by proxy through the firewall. If you select this option, only the files necessary to send administrative messages (including firewall-generated alerts, messages, and logs) will be configured.
 - **Secure Split SMTP Servers (firewall-hosted)** – Use this option to use the firewall-hosted sendmail server(s). This configuration allows you to take advantage of additional sendmail features, including header stripping, mail routing, aliases, and masquerading.
- 4 In the **Internal SMTP Burb** field, select the burb in which your site's internal SMTP server resides.
- 5 In the **Internal SMTP Mail Server** field, type the fully qualified name of your site's internal SMTP server.
Do not use simple host names or IP addresses.
- 6 Click **Save** in the toolbar (or click **Apply** if you are accessing this window from the Services window) to reconfigure your mail mode. A confirmation window appears when the reconfiguration process is complete.
- 7 [Conditional] If you accessed Reconfigure Mail from the Services window, click **Close** to return to the sendmail server Properties window.
- 8 Select **Policy > Rules** and create or modify the necessary rules:
 - If you selected Transparent:
 - Create two rules: one for inbound mail and one for outbound mail.
 - Use the **smtp (Mail Proxy)** service.
 - Use two **Mail (SMTP proxy)** application defenses: one for inbound traffic and one for outbound traffic, each with direction-appropriate settings.
 - If you selected Secure Split SMTP Servers:
 - Create two rules: one for inbound mail and one for outbound mail.
 - Use the **sendmail (Sendmail Server)** service.
 - The rules' destination burb must be **<Any>**. The endpoint must also be **<Any>**.
 - Use two **Mail (Sendmail)** application defenses: one for inbound traffic and one for outbound traffic, each with direction-appropriate settings.

See [Chapter 9, Rules](#), for more information on how to manage rules.

Note: If you are changing your mail configuration, you will need to update or replace your existing mail rules to reflect the new configuration.

- 9 Save your changes.

The firewall now has a new mail configuration.

- If you selected Transparent, mail management is primarily handled off-box and does not require changes to any Sidewinder mail files.
- If you selected Secure Split SMTP Servers and are an experienced mail administrator, you may want to edit the configuration files. See [Editing sendmail files on Sidewinder](#) and [Configuring advanced sendmail features](#) for more information.

Understanding sendmail on Sidewinder

Hosting sendmail on Sidewinder requires knowledge of both the sendmail application and how the firewall interacts with it. Read the following sections to learn about domain and file names, Interoperability considerations, and filtering services.

Using sendmail on Sidewinder

When you use the Reconfigure Mail window (**Maintenance > Reconfigure Mail**) to select secure split services, you specify a mail host on your internal network and the burb where that server is located. This is the internal mail host that will send mail to, and receive mail from, the firewall-hosted mail server. The necessary configuration files and everything you need to run the firewall-hosted mail server are automatically set up for you, such as:

- The three mail domains: *mtac*, *mtaX*, and *mtaY* (where *X* = the number of the external burb, and *Y* = the number of an internal burb), are in place. Sendmail is already configured to route mail among the three sendmail servers.
- Mail addressed to users on your internal network will be forwarded to the mail host you specified in the Reconfigure Mail window.
- Messages that are sent to the person administering a mail system are generally addressed to “postmaster.” During the Quick Start Wizard (initial configuration), you set up an administrator’s account. Postmaster messages are automatically routed to that administrator’s firewall-hosted mail account. (We recommend that all administrators redirect their local mail to a non-firewall-hosted e-mail account.)

Note: You will need to configure your internal mail server to forward non-local mail to the firewall. This procedure differs depending on the type of mail program your network runs. Refer to your mail software’s documentation for details.

When you configure secure split SMTP services, there are three separate sendmail servers that each have a different purpose:

- Local

The local server handles mail that is sent directly from the firewall. For example, if an administrator sends a mail message from the firewall, it is sent through the local server. This sendmail process runs in the *mtac* domain and forwards all mail to the firewall’s internal network.

- Internal

The internal server runs in the *mta#* domain, where *#* is the burb index of an internal burb that you specify when running Reconfigure Mail.

This internal sendmail server *receives* mail from one of three sources:

- a host on the internal network
- a sendmail process transferring mail from the local sendmail server
- a sendmail process transferring mail from the external sendmail server

This internal sendmail server *delivers* mail to one of three places:

- If the message is for a user local to the firewall, such as an administrator with a mailbox on the firewall, it delivers the message to the user’s mailbox using the `mail.local` program.
- If the message is for a user on the internal network, it connects to the mail host on the internal network and delivers the mail there.
- If the message is not for either of the above, it assumes the message is for an external user and transfers the message to the external burb for that user.

- external

The external server runs in the *mta#* domain, where # is the burb index of the Internet burb. This sendmail server *receives* mail from one of two sources:

- a host on the external network
- a sendmail process transferring mail from the internal sendmail server

The external server *delivers* mail to one of two places:

- If the message is for an external user, it connects to an external host and delivers the mail there.
- If the message is for a user local to the firewall (such as an administrator) *or* for a user on the internal network, it transfers the mail to the internal burb for delivery to that user.

When using firewall-hosted SMTP services, all mail for a user local to the firewall goes to the internal *mta* domain for delivery. Local delivery does not take place in the external *mta* domain or the *mtac* domain. Running sendmail on the firewall works as it does in any other UNIX environment, with the following exceptions:

- The firewall runs three separate sendmail servers (as described in the previous section).
- Type Enforcement restricts sendmail so that its security flaws cannot be exploited. For example, firewall administrators cannot execute shell scripts or other executables through sendmail, as they could do on a standard UNIX system.
- Aliases allow users to send their mail to another mailbox that may be at a different location. For example, firewall administrators might choose to redirect their mail to a mailbox located on the internal network so they receive all of their mail in one place. Administrators can use the */etc/mail/aliases* file, but this file cannot contain commands to run other programs, such as program mailers (for example, *procmail*). For more information on aliasing mail, see [Setting up e-mail aliases for administrator accounts](#).
- If a server is too busy to send a message, or if the machine it is sending mail to is not responding, the messages are sent to a mail queue. The firewall has a separate queue for each sendmail server: */var/spool/mqueue.#*, */var/spool/mqueue.#*, and */var/spool/mqueue.c* (# = the burb number).

Note: If mail cannot be delivered on the first attempt, it is placed in a queue. By default, the system checks the queues every 30 minutes and attempts redelivery.

You can check if there are messages in the mail queues by following the steps described in [Managing mail queues](#).

Mail is an extremely complex subject and can require a great deal of effort to configure. If you want additional information on managing mail, the best resource is the book *sendmail* by Bryan Costales (O'Reilly & Associates, Inc.).

Mail filtering services on Sidewinder

An advantage to using hosted sendmail is the ability to do on-box mail filtering. To filter messages, you must create rules using the sendmail server and a Mail Application Defense with the filter options configured. The following mail filtering services are available:

Note: You must have Secure Split SMTP mail servers configured to use the mail filtering listed here.

MIME/Virus/Spyware filtering

The MIME/Virus/Spyware application defense options allow you to:

- Allow, deny, or scan specific types of MIME elements and specific file extensions
- Configure how to handle infected files
- Specify file attachment size restrictions (per message, not per attachment)
- Determine whether mail messages will be scanned as a whole (entire message is allowed or denied) or in segments (attachments may be dropped if they do not meet filtering criteria, but the acceptable portions of the mail message will still reach the recipient)
- Reject all mail if scanning services become unavailable

The virus scanner must be configured at **Policy > Application Defenses > Virus Scanning**. Virus scanning can then be applied per-rule using the rule's application defense. See [Configuring virus scanning services](#) and [Configuring the Mail \(Sendmail\): MIME/Virus/Spyware tab](#) for more information.

Note: You must purchase and activate the anti-virus add-on module before the MIME/Virus/Spyware filter rules you create will scan mail messages for viruses. MIME type filtering does not require the Anti-Virus license.

TrustedSource

TrustedSource is a reputation service that filters incoming mail connections and then provides precise information about an e-mail sender's reputation based on its IP address. The TrustedSource reputation service is a tool for reducing the amount of spam that reaches your organization's inboxes.

You can enable TrustedSource at **Policy > Application Defenses > TrustedSource**. Select **Perform TrustedSource filtering on inbound mail**, then set the threshold to a value from 0 to 120.

See [About TrustedSource](#) for more information.

Key word search filtering

The Keyword Search filter allows you to filter mail messages based on the presence of defined key words (character strings).

See [Configuring the Mail \(Sendmail\): Keyword Search tab](#).

Configure size limitations for mail messages

The size filter performs a check on e-mail messages for the number of bytes the message contains, including the message header. Messages that equal or exceed the specified size you specify will be rejected.

See [Configuring the Mail \(Sendmail\): Size tab](#).

Anti-relay controls

Anti-relay control uses access control to prevent your mailhost from being used by a hacker as a relay point for spam to other sites. This option is automatically enabled for all Mail defenses and cannot be disabled.

See [Configuring the Mail \(Sendmail\): Control tab](#).

Editing sendmail files on Sidewinder

If sendmail is hosted on your Sidewinder, then the sendmail configuration information is stored in *sendmail.cf* files. These files contain information such as which delivery agents to use and how to format message headers. You should change your configuration options only if you are directed to do so by Forcepoint support, or if you are an experienced sendmail user and want to customize the files for your site.

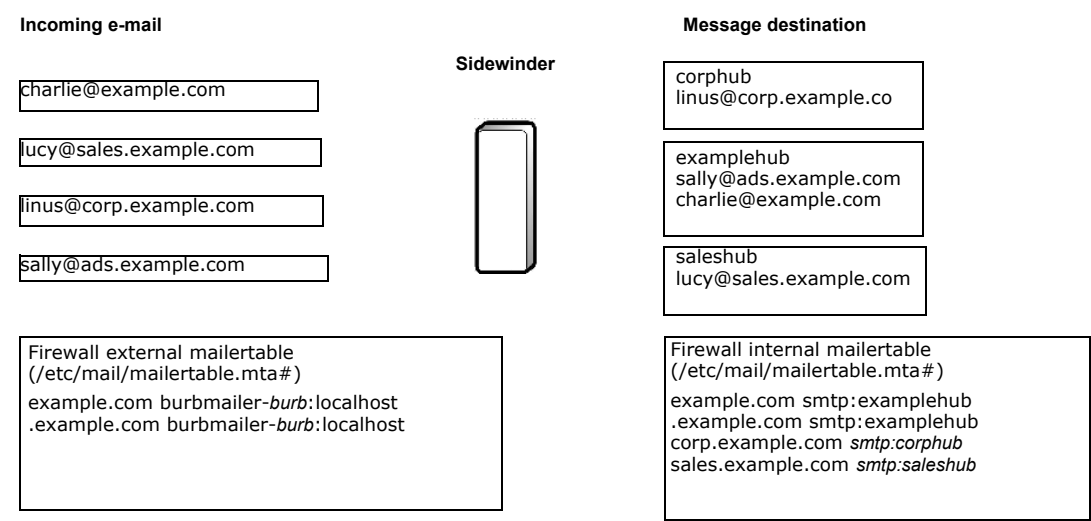
Sendmail allows you to create configuration files using macros written for the `m4` preprocessor. Sections 19.5 and 19.6 in the *UNIX System Administration Handbook* describe these macros. You can also refer to the book *sendmail* by Bryan Costales (O'Reilly & Associates, Inc.).

The firewall sets up two mailertables for you: one internal and one external.

- The external mailtable, `/etc/mail/mailertable.mta#` (`#` = the number of the external burb), processes the mail and directs it to the internal mailtable.
- The internal mailtable, `/etc/mail/mailertable.mta#` (`#` = the number of an internal burb), sorts the mail by host name, and sends the mail to the correct internal mail host.

The following figure shows an example of the route along which incoming mail messages travel.

Figure 304 Sidewinder mailertables




The recommended method of editing the mail files is to change the sendmail server's properties using the Admin Console. This opens a file editor that knows to automatically rebuild and restart the sendmail server when you save a file.

To edit the mail configuration files using this method:

Caution: Only experienced administrators should modify sendmail configuration files.

- 1 In the Admin Console, select **Policy > Rule Elements > Services** and then select **sendmail**.
- 2 Click **Properties**. The following window appears:

Figure 305 Services: Sendmail Configuration window



- 3 Select the configuration file you want to modify in the appropriate burb configuration file list. There are separate files for each sendmail server running on the firewall. You may edit the following files:

Tip: Before making any changes, select **File > Backup** and create a backup file. Also, for best results, do not edit these files with any other file editor, as those editors will not automatically rebuild and restart the sendmail server.

- **Access Table** – This file defines anti-relaying and anti-spamming policies for the SMTP server.
- **Aliases file** – (Available only in the internal burb.) This file defines the mail aliases that are used to redirect e-mail to another person or location.
- **Alternate Host Names file** – This file identifies alternate host names by which the firewall is known. E-mail addressed to any of the alternate names is treated as local mail by the firewall.
- **Domain Table** – This file provides a mapping from an old domain name to a new domain name. For example, you might modify this file if your organization's external domain name changes.
- **M4 Config file** – This file defines the initial sendmail configuration. Modify this file as needed to account for your site-specific requirements.
- **Mailer Table** – This file maps a domain to a mail relay that is responsible for mail delivery in that domain.

Note: Only edit mail configuration files if it is necessary for your site's e-mail functionality.

- 4 Save and then close the file.

- 5 Open the appropriate mailtable file and edit as necessary.

Note: Only edit mailtable files if it is necessary for your site's e-mail functionality.

The mailtable files are named `/etc/mail/mailtable.mta#` (# = the appropriate burb number).

- 6 Enter the correct domain, mailer, and host in the following format:

```
domain<tab>mailer:host
```

On the internal side of the network, the mailtable appears as:

```
.example.net<tab>smtp:examplehub  
example.net<tab>smtp:examplehub  
corp.example.net<tab>smtp:examplehub  
sales.example.net<tab>smtp:examplehub
```

On the external side of the network, the mailtable should appear as:

```
example.net<tab>burbmailer-burb:localhost  
.example.net<tab>burbmailer-burb:localhost
```

where *burb* = the external burb number and *y* = the internal (trusted) burb number.

The entries that begin with a dot (.) act as a wildcard, matching anything with that domain name. The entries that do not begin with a dot match the full domain name. See the `/usr/share/sendmail/README` file for more information on creating mailtables.

- 7 Save and then close the file.

- 8 Click **Save** to save the configuration changes and rebuild the configuration and database files. This will also automatically restart the sendmail servers.

Note: If at a command prompt, use `cf sendmail rebuild` and `cf daemon restart agent=sendmail`.

The firewall has updated sendmail with your changes and is now ready to process mail.

This window also has a shortcut to the **Reconfigure Mail** area. See [Setting up and reconfiguring mail](#) for details on changing your mail configuration.

Configuring advanced sendmail features

Once you run Reconfigure Mail to set up hosted mail and create the appropriate rules, the basic mail services are enabled. However, sendmail provides several additional features that you may choose to configure. Of those listed here, mail routing, header stripping, and the RealTime Blackhole list are the most popular additional sendmail features. The details for implementing these features are described in the sections that follow.

- **Mail routing** – Enables you to reroute e-mail from one domain name to another domain name by editing the mailertable files.

See [Editing sendmail files on Sidewinder](#).

- **Header stripping** – Enables you to remove header information from an outbound message to conceal internal host information from the outside world.

See [Configuring sendmail to strip message headers](#).

- **Blackhole list** – Enables you to eliminate unwanted and unsolicited e-mail. The types of spam control you might implement include use of a Realtime Blackhole list, Promiscuous Relaying, and so on.

See [Configuring sendmail to use the RealTime Blackhole list](#).

- **Filter mail based on the user** – Enables you to allow or deny mail based on a specific user or users.

See [Allowing or denying mail on a user basis](#).

- **Masquerading** – Enables you to transform a local host address in the header of an e-mail message into the address of a different host.

See [Configuring sendmail to hide internal e-mail addresses](#).

Tip: You can also configure aliases for e-mail accounts. Creating aliases for the firewall administrator accounts is particularly useful because system messages sent to these accounts can, if left unattended, fill up the firewall's hard drive.

Configuring sendmail to strip message headers

During the normal operation of sendmail, the path a message traces is appended to the message by each host through which the mail passes. This enables internal host names and IP addresses to be allowed beyond the firewall.

Note: Header information can only be removed for outbound mail (that is, mail leaving the firewall). Therefore, you should only enable header stripping in the destination (or external) burb for a message. If you configure header stripping in the source burb of a message, header stripping will not happen for that message.

You can configure sendmail to strip (remove) or scrub (change to a different value) the following headers from messages leaving the firewall:

- Received (stripped)
- X400-received (stripped)
- Via (stripped)
- Mail-from (stripped)
- Return-path (stripped)
- Message-id (scrubbed)
- Resent-message-id (scrubbed)

Note: Stripping the headers will **not** alter the To and From hosts. The To and From hosts can be eliminated using rules in the **sendmail** configuration file. You can also modify the To and From hosts using masquerading or by editing the domain tables.

To configure sendmail to strip or scrub headers:

- 1 Select **Policy > Rules Elements > Services**.
 - 2 Select **sendmail** and click **Modify**.
 - 3 Click **Properties**. A list of mail files appears. Note that separate configuration files are maintained for each burb.
 - 4 Select the **M4 Config File** in the external burb list and click **Edit File**.
- Tip:** Before making any changes, select File > Backup and create a backup of this file.
- 5 Locate the `C{STRIP_DOMAINS}` line in the file and append the domain name on which to perform header stripping. For example:

```
C{STRIP_DOMAINS} domainx
```

where *domainx* = the domain name on which to perform header stripping.

You can define multiple domains by entering multiple domain names on one line (for example, `C{STRIP_DOMAINS} abc.com xyz.com`).

Note: STRIP_DOMAINS contains the list of domains that will trigger header stripping. Each message processed by **sendmail** in the external burb will be subjected to header stripping if it is received from a domain in this list.

- 6 Save and then close the file.
- 7 Click **OK** to return the main Services window.
- 8 Click **Save** to save the configuration changes and rebuild the configuration and database files. This will also automatically restart the sendmail servers.

Note: If at a command prompt, use `cf sendmail rebuild` and `cf daemon restart agent=sendmail`.

Configuring sendmail to use the RealTime Blackhole list

Sendmail is able to use the services of the RealTime Blackhole List (RBL). The Blackhole List, a list of known spam domain names, is maintained by an organization called MAPS (Mail Abuse Prevention System). The mail server checks each mail message against the Blackhole list. Any e-mail message originating from a domain in the list will be rejected.

Note: You must subscribe to the Trend Micro Blackhole List in order to use it. Go to www.trendmicro.com/ for details. Other blackhole lists may work with the Sidewinder, but are not supported.

To configure the firewall to use a Realtime Blackhole List:

- 1 Select **Policy > Rules Elements > Services**.
- 2 Select **sendmail** and click **Modify**.
- 3 Click **Properties**. A list of mail files appears. Note that separate configuration files are maintained for each burb.
- 4 Select the **M4 Config File** in the external burb list and click **Edit File**.

Tip: Before making any changes, select **File > Backup** and create a backup of this file.

- 5 Add the following line to the file.

```
FEATURE(`dnsbl', `domain') dn1
```

The *domain* that you enter in the above line will depend on the type of service for which you have subscribed. If you subscribe to MAPS, they will provide you with the correct domain (for example, *blackholes.mail-abuse.org*) to use.

- 6 Save and then close the file.
- 7 Click **OK** to return the main Services window.
- 8 Click **Save** to save the configuration changes and rebuild the configuration and database files. This will also automatically restart the sendmail servers.

Note: If at a command prompt, use **cf sendmail rebuild** and **cf daemon restart agent=sendmail**.

Now when the firewall receives mail, it will query the RBL to see if the sender's domain is on the list. If the domain is a match, sendmail rejects the message.

Allowing or denying mail on a user basis

Sendmail will allow or deny mail on a domain basis. However, you can also instruct sendmail to allow or deny mail to/from specific users, IP addresses, and subnets within a domain. To do this, follow the steps below:

- 1 Select **Policy > Rules Elements > Services**.
- 2 Select **sendmail** and click **Modify**.
- 3 Click **Properties**. A list of mail files appears. Note that separate configuration files are maintained for each burb.
- 4 Select the **Access Table** file for the appropriate burb and click **Edit File**.

Tip: Before making any changes, select File > Backup and create a backup of this file.

- 5 Add the specific allow (**RELAY**), deny and notify the sender (**REJECT**), and/or deny without notifying the sender (**DISCARD**) information to the access table.

For example, if you want to allow mail addressed to Lloyd and Sharon but deny mail addressed to everyone else, you would add the following lines:

```
# Allow mail addressed to these users
To:Lloyd@example.com RELAY
To:Sharon@example.com RELAY
# Deny mail for everyone else
To:example.com REJECT
```

Note: For additional information, see the `README` file in the `/usr/share/sendmail` directory on the firewall.

- 6 Save and then close the file.
- 7 Click **OK** to return the main Services window.
- 8 Click **Save** to save the configuration changes and rebuild the configuration and database files. This will also automatically restart the sendmail servers..

Note: If at a command prompt, use `cf sendmail rebuild` and `cf daemon restart agent=sendmail`.

Mail from those specified users, IP addresses, and subnets will now be handled as indicated in the file.

Configuring sendmail to hide internal e-mail addresses

Occasionally, you may use domain names on your internal network that you do not want the rest of the Internet to know about. You can instruct sendmail to change the header information so that it hides internal domains before relaying the mail on to the final destination. This is called *masquerading*. Masquerading also involves modifying the “From “ or “From:” field before the mail is relayed. To do this, follow the steps below:

Tip: Using `masquerade_entire_domain` field is effective, but it leaves open the possibility of showing internal addresses that are included on the message (such as the CC or To fields). Use `masquerade_envelope` field to masquerade all addresses in the envelope containing the domain using the specified domain.

- 1 Select **Policy > Rules Elements > Services**.
- 2 Select **sendmail** and click **Modify**.
- 3 Click **Properties**. A list of mail files appears. Note that separate configuration files are maintained for each burb.
- 4 Select the **M4 Config File** for the Internet burb (generally the external burb).

Note: Before making any changes, select File > Backup and create a backup.

- 5 Identify the domains you want hidden:

- a Locate the `MASQUERADE_DOMAIN` section. The default looks like this:

```
dn1 # MASQUERADE_DOMAIN(`hide_me.acme.com hide_me_too.acme.com')dn1
```

- b Uncomment the line by deleting “dn1 #”.

- c Changed the listed domains to the domain or domains that you want to hide. For example, ``hide_me.acme.com hide_me_too.acme.com'` becomes ``sales.example.net'`.

- 6 Enter the domain you want to show:

- a Locate the `MASQUERADE_AS` section. The default looks like this:

```
dn1 # MASQUERADE_AS(`newdomain.com')dn1
dn1 # FEATURE(`masquerade_entire_domain')dn1
dn1 # FEATURE(`masquerade_envelope')dn1
```

- b Uncomment the section by deleting each “dn1 #”.

- c In the `MASQUERADE_AS` line, change the listed domain to the domain that should replace all internal domains. For example, ``newdomain.com'` becomes ``example.net'`.

- 7 Save and then close the file.
- 8 Click **OK** to return the main Services window.
- 9 Click **Save** to save the configuration changes and rebuild the configuration and database files. This will also automatically restart the sendmail servers.

Note: If at a command prompt, use `cf sendmail rebuild` and `cf daemond restart agent=sendmail`.

Enabling Sendmail TLS

The Sendmail implementation of RFC 2487, SMTP over TLS, is supported on Sidewinder.

Sendmail can act as either a client or server in a TLS session:

- When acting as the server, it advertises the STARTTLS feature in the response to the EHLO command, then responds positively to the subsequent STARTTLS command.
- When acting as the client, it issues the STARTTLS command if the remote server advertises STARTTLS on the EHLO response.

In both cases, after the STARTTLS command and positive response, the client and server negotiate a TLS session.

Note: As part of the implementation, Sendmail TLS also enforces FIPS mode.

For more information on enabling Sendmail TLS, see the Knowledge Base article [9003](#).

Managing mail queues

If a sendmail message cannot be delivered (for example, if the destination system is down), messages are temporarily placed in queues until they can be delivered. There are separate queues for each server: */var/spool/mqueue.c* (local) and */var/spool/mqueue.#* for the Internet and the trusted burbs. The following sections explain how to view mail, how to change some of the basic queue settings, and how to manually force sendmail to attempt to deliver queued mail.

Tip: You should check the queues periodically. If there are a lot of messages that are several days old, you may have a problem with your system or its configuration.

Viewing the mail queue

To view the mail queue output, type the following command:

```
mailq
```

The output of this command lists the messages currently in the queue you chose, along with information about each message. Each message is assigned a unique identification number, which is shown in the first column. In the following example, the external queue shows a message still in queue due to some temporary error. The internal queue shows a valid message ready to be delivered or possibly currently being delivered. The common_sendmail queue shows no mail queued up which should normally, but not always, be the case.

Listing the **external** Queue

```
                /var/spool/mqueue.1 (1 request)
-----Q-ID----- --Size-- -----Q-Time-----
-----Sender/Recipient-----
kA6M17qb008045      4 Mon Nov  6 16:01 me@mydomain.com
                    (Deferred: Connection refused by yourdomain.com.)
                    you@yourdomain.com

Total requests: 1
```

Listing the **internal** Queue

```
                /var/spool/mqueue.2 (1 request)
-----Q-ID----- --Size-- -----Q-Time-----
-----Sender/Recipient-----
kA6M4gd8008175      4 Mon Nov  6 16:04 admin@fwdomain.com
```

Changing how long a message waits between delivery attempts

By default, undelivered e-mail messages remain in the mail queues 30 minutes before another delivery attempt is made. If you want to change the length of time e-mail messages remain in the mail queues before another delivery attempt is made, do the following:

- 1 Select **Policy > Rules Elements > Services**.
- 2 Select **sendmail** and click **Modify**.
- 3 Click **Properties**. A list of mail files appears. Note that separate configuration files are maintained for each burb.
- 4 Select the **M4 Config File** for the burb that is running sendmail, and click **Edit File**.

Tip: Before making any changes, select **File > Backup** and create a backup of this file.

- 5 Scroll to the **Set the Queue Interval** area and edit the following line:

```
define(`confQUEUE_INTERVAL', `Xm')dnl
```

where:

X is the amount of time that the message remains in the queue before an attempt is made to resend the message.

m indicates that the time will be measured in minutes. You can also use other time measurements, such as seconds (s), hours (h), days (d), etc.

Note: The default value is 30 minutes.

- 6 Save and then close the file.
- 7 Click **OK** to return the main Services window.
- 8 Click **Save** to save the configuration changes and rebuild the configuration and database files. This automatically restarts the sendmail servers.

The time a message waits before sendmail attempts to deliver it has now been changed.

Manually attempting to deliver queued messages

Occasionally, you may need to attempt to send all queued messages immediately instead of waiting for them to be pushed automatically. This process is called *flushing* the mail queue. If you want to force sendmail to attempt to deliver its queued messages, do the following:

- 1 At a firewall command prompt, enter the following command to change to the Admn role:

```
srole
```

- 2 Instruct sendmail to manually attempt to deliver mail in one or more mail queues:

- **cf sendmail flush** – Flushes all three queues.
- **cf sendmail flush queue=*burbname*** – Flushes only the queue for that burb.
- **cf sendmail flush queue=*common*** – Flushes the queue containing mail sent by the firewall, such as system updates and alerts.

The firewall immediately attempts to send all mail in the queue.

Changing how long a message waits before it is returned to its sender

By default, undelivered e-mail messages remain in the mail queues 5 days before they are returned to their senders as undeliverable and deleted from the queue. If you want to change the length of time e-mail messages remain in the mail queues before they are considered undeliverable, do the following:

- 1 Select **Policy > Rules Elements > Services**.
- 2 Select **sendmail** and click **Modify**.
- 3 Click **Properties**. A list of mail files appears. Note that separate configuration files are maintained for each burb.
- 4 Select the **M4 Config File** for the burb that is running sendmail, and click **Edit File**.

Tip: Before making any changes, select **File > Backup** and create a backup of this file.

- 5 Locate the **Set the Queue Interval** area and edit the following line:

```
define(`confTO_QUEUERETURN', `Xd')dnl
```

where:

X is the amount of time that the message remains in the queue its sender is notified that it was undeliverable and the message is deleted.

d indicates that the time will be measured in days. You can also use other time measurements, such as seconds (s), minutes (m), hours (h), etc.

The default value is 5 days.

- 6 Save and then close the file.
- 7 Click **OK** to return the main Services window.
- 8 Click **Save** to save the configuration changes and rebuild the configuration and database files. This automatically restarts the sendmail servers.

The time an undelivered message waits before sendmail returns it to its sender has now been changed.

Receiving mail sent by Sidewinder

Sidewinder sends status updates and alerts to root and administrator accounts. By default, these accounts are hosted on the firewall and must be checked using a command line session. If you want to redirect mail from your administrators' firewall mailboxes to a different destination, you can edit the `/etc/mail/aliases` file. The following sections provide information on how to create e-mail aliases or access the mail messages directly on the firewall.

Setting up e-mail aliases for administrator accounts

On the firewall, messages and other files are often e-mailed to system users such as *root* and *postmaster*. To redirect these system messages to an external account, you can set up an alias.

Tip: Remember to update aliases when there are personnel changes.

Aliases are stored in the `/etc/mail` directory, which can be accessed through the sendmail service. Do the following to set up a mail alias for system users:

- 1 At a firewall command prompt, enter the following command to change to the Admn role:

```
srole
```

- 2 Using a file editor, open `/etc/mail/aliases`.
- 3 Locate the root line in the file. The default, root, is automatically aliased to the administrator account created during the Quick Start Wizard and looks like this:

```
#root: username
```

- 4 Uncomment the line by deleting the #, and then replace the existing address with the off-box e-mail address of the person who will receive system messages. If you want to add multiple accounts, separate them with commas and do not include spaces.

The line now looks like this:

```
root: username_a@example.com,username_b@example.com
```

By default, all other system roles are aliased to root, and mail sent to those accounts will also be sent to the e-mail address entered above. To redirect other system roles' mail to other accounts, use the same format.

- 5 Save the changes and then close the file.

- 6 Enter the following commands:

```
cf sendmail rebuild
```

```
cf daemond restart agent=sendmail
```

This rebuilds the configuration and database files, and restarts the sendmail servers.

System mail messages will now be sent to the aliased account instead of accumulating on the firewall's hard drive.

Viewing administrator mail messages on Sidewinder

By default, a root alias is created for the administrator you set up when you configured your system. This alias automatically redirects system messages addressed to root to that first administrator's firewall-hosted account. A mailbox will be created the first time an administrator sends or receives a mail message. Mailboxes for firewall administrators are stored in the `/var/mail` directory.

Note: Do not ignore the e-mail that accumulates on the firewall as it contains important information about your network and the firewall, and also uses disk space. Routinely read and delete mail sent to the firewall, or have it redirected elsewhere. To redirect mail to another destination, see the previous section, [Setting up e-mail aliases for administrator accounts](#).

To view system messages sent to firewall-hosted accounts, follow the steps below.

- 1 At a firewall command prompt, enter the following command to change to the Admn role:

```
srole
```

- 2 View e-mail messages by entering one of the following:

E-mail

Receiving mail sent by Sidewinder

- **mail** – Displays your messages (messages for the logged-in administrator)
- **mail -f root** – Displays messages addressed to root
- **mail -f username** – Displays messages addressed to that administrator

Tip: Refer to the **mail** man page for detailed information on using the **mail** command. If you prefer, you may use an alternate mail program.

Remember to check mail frequently, particularly if you have attack and system event responses sent to *root*.

Note: For more on responses, see [Chapter 13, IPS Attack and System Event Responses](#).

E-mail
Receiving mail sent by Sidewinder

20 Virtual Private Networks

Contents

[About the Sidewinder VPN solution](#)

[Planning your VPN](#)

[Creating VPN policy](#)


[VPN user interface reference](#)

[Example VPN Scenarios](#)

About the Sidewinder VPN solution

The Forcepoint Sidewinder VPN solution provides secure data transmission across unsecured networks through an encryption and decryption process. The firewall uses the IPsec protocol suite to support this process.

Figure 306 Sidewinders, an IPsec or IKE remote site, or a VPN client machine



The Sidewinder VPN solution is embedded in the architecture, making it an operating characteristic of the OS. This integration not only lets you apply access rules to VPNs in the same way you do for physically connected networks, but also means that you use the Sidewinder VPN solution to coordinate corporate-wide network security policies.

As companies expand to new locations and employees spend more time working out of the office, VPN solutions are becoming more important to businesses. Consider the value of encrypting and authenticating data in these situations:

- Passing traffic from firewall to firewall between offices located in different cities (gateway-to-gateway VPN)
- Passing traffic from employees working remotely to your network (client-to-gateway VPN)

Protecting your information

The Internet is a broadcast medium that is used to send information. While information is in transit, anyone can choose to monitor or intercept this information.

Sending information beyond your firewall via the Internet is like sending an unsealed envelope of important information via a courier service: you must trust that the courier will not read or steal the information.

To address this danger, the IETF (Internet Engineering Task Force) developed a standard for protecting data on unprotected (or untrusted) networks such as the Internet. The standard has become known as *IPsec*, meaning Internet-Protocol Security. In brief, IPsec calls for encrypting the data before it leaves the local host, then decrypting it when it is received at the destination or remote host. Once it is decrypted, the data assumes its original form and can be read as intended. No matter how long or circuitous its route through the Internet, the data remains private by virtue of this encryption.

What are encryption and authentication?

The two main components of IPsec security are *encryption* and *authentication*.

- **Encryption** – Encryption is the means by which plain text is translated into code. It ensures that the transmitted data remains private and unreadable until properly decrypted. Sidewinder uses an encryption key to encipher and decipher each unit of data sent between your site and the “partner” or remote VPN site. (See [About IPsec keys](#).)
- **Authentication** – VPN authentication prevents unauthorized individuals from tampering with the contents of the data being transmitted. It also prevents them from creating messages that claim to come from a particular place but are actually sent from somewhere else (such as the hacker’s home computer). Authentication is accomplished through two methods:
 - Data-integrity checking, which allows the receiver to verify whether the data was modified or corrupted during transmission.
 - Sender identification, which allows the receiver to verify whether the data transmission originated from the source that claims to have sent it.

When used together, encryption and authentication are very much like writing an encoded message, sealing it in an envelope, and then signing your name across the flap. The receiver can first verify that the signature is yours as a means of determining the origin of the message. Next, the receiver can determine if the contents have been viewed or altered by checking that the envelope seal has not been compromised. Once the receiver is assured of the authenticity of the message, they can decode the contents and know that the contents are as intended.

About IPsec keys

A key is a number that is used to electronically sign, encrypt, and authenticate data when you send it, and to decrypt and authenticate your data when it is received. When a VPN is established between two sites, two keys are generated for each remote end: an encryption key and an authentication key.

To prevent these keys from being guessed or calculated by a third party, a key is a large number. Encryption and authentication (or session) keys are unique to each VPN definition you create.

Once generated, these keys are exchanged (either automatically or manually) between the sites, so that each end of the VPN knows the other end’s keys.

To generate key pairs, the firewall gives you two options:

- **Manual key generation** – If the remote site is not Internet Key Exchange (IKE)-compliant, you may want to choose the manual method of key generation. With this method, the firewall provides randomly generated encryption and authentication keys (or you can create your own), which you must copy and pass to the remote end of the VPN via secure e-mail, diskette, or telephone. Repeat this process each time you generate keys. Manual keys are more labor intensive than automatic keys and rarely used.
- **Automatic key generation using IKE** – If the remote end of your VPN uses the IKE protocol, the firewall can manage the generation of session keys between sites automatically. This process also regularly changes the keys to avoid key-guessing attacks. Automatic keys are very common in today’s network environments.

Planning your VPN

Before you create new VPN policy, plan the characteristics of the VPN based on your network environment, security requirements, authentication requirements, and the type of IPsec-compliant remote device to which the VPN will be established.

Use the following sections to plan your VPN:

- [Choosing the appropriate VPN attributes](#)
- [Choosing the appropriate authentication type](#)
- [Ordering VPN definitions](#)
- [Restricting VPN access with a virtual burb](#)

Understanding the options associated with each concept will assist you in creating your VPN definition. Study the following information to help you determine which VPN configuration best suits your network environment.

Choosing the appropriate VPN attributes

Before you create a new VPN definition, select the attributes of your VPN:

- [Determine the VPN mode](#)
- [Choose an IKE version](#)
- [Determine the encapsulation method](#)
- [Consider using a client address pool](#)

Determine the VPN mode

Creating a VPN involves establishing an association (or a trust relationship) between your Sidewinder and an IPsec-compliant remote firewall, host, or client. These devices are referred to as “VPN peers.”


There are two types of VPN peers:

- **VPN gateways** provide VPN access to other devices. Your Sidewinder functions as a VPN gateway.
- **VPN clients** use VPN access provided by VPN gateways.

The type of VPN peer you are working with determines the type of VPN you will create. There are two VPN types:

- **gateway-to-gateway** – A VPN tunnel is established between both gateways to allow hosts behind them to communicate with each other.


Figure 307 Example gateway-to-gateway VPN



In the example above, the VPN tunnel is established between the firewalls, allowing the hosts in the 50.1.0.0/16 and the 250.1.1.0/24 networks to communicate securely.

- **client-to-gateway** – The client(s) establish a VPN tunnel to the gateway to communicate with hosts behind the gateway.

Figure 308 Example client-to-gateway VPN



In this example, a VPN tunnel is established between each VPN client and the firewall, allowing the clients to communicate with the 250.1.1.0/24 network securely.

The type of VPN that meets your requirements affects the VPN modes that are available when you create a new VPN definition on your firewall. Use the table below to determine what type of VPN is appropriate for your needs and which VPN mode provides the features you require.

Table 62 Choosing VPN mode

If...	then the VPN type is	and the available modes are
The VPN peer has a dynamic IP address and will provide VPN access to clients that are behind it	gateway-to-gateway	Dynamic IP Restricted Client
The VPN peer has a static IP address and will provide VPN access to clients that are behind it	gateway-to-gateway	Fixed IP
Multiple roving clients need VPN access	client-to-gateway	<ul style="list-style-type: none"> • Dynamic IP Client • Dynamic IP Restricted Client
The VPN peer does not support IKE	gateway-to-gateway	Manually Keyed VPN

More about VPN modes:

- **Fixed IP** – Select this option if the IP address of the remote end is always the same. You must also provide the IP address of the remote end in the **Remote IP** field.
- **Dynamic IP Client** – Select this option if the remote end is a device whose IP address is not fixed. *Example: A traveling executive that gains Internet access from a laptop.*
- **Dynamic IP Restricted Client** – Select this option if the remote end is a device whose IP address is not fixed. This differs from the *Dynamic IP Client* option in that you restrict the range of IP addresses available to the remote end by using either the **Client Address Pool** field or the **Dynamic Virtual Address Range** field. *Example: A member of a sales team that gains Internet access from a laptop.*
- **Manually Keyed VPN** – Select this option if you want to exchange session keys manually (for example, over the phone). You configure specific properties of the manual key exchange on the Crypto tab.

Note: This mode is most often used to provide compatibility with remote peers that do not support the Internet Key Exchange (IKE) protocol.

Choose an IKE version

IKEv1 and IKEv2 are both available options on Sidewinder. Some differences to note about the two versions:

- IKEv2 is simpler, more robust, and more reliable. However, not many products currently support the newer IKEv2. Check your product documentation.
- IKEv1 is not compatible with IKEv2. Both sides of a VPN connection must use the same version of IKE.
- When using IKEv2, each side of a VPN connection can use a different authentication method. With IKEv1, both sides must agree on an authentication method.
- In IKEv2, extended authentication (XAUTH) can be used as a standalone authentication method. In IKEv1, extended authentication must be used in conjunction with password/certificate authentication.

Determine the encapsulation method

There are two methods for encapsulating packets in a VPN connection: transport mode and tunnel mode.

- **Transport mode** – In transport mode, only the data portion of the packet gets encrypted. This means that if a packet is intercepted, a hacker will not be able to read your information, but will be able to determine where it is going and where it has originated. This mode existed before firewalls and was designed for host-to-host communications.
- **Tunnel mode** – In tunnel mode, both the header information and the data is encrypted and a new packet header is attached. The encryption and new packet header act as a secure cloak or “tunnel” for the data inside. If the packet is intercepted, a hacker will not be able to determine any information about the true origin, final destination, or data contained within the packet because the packet header is encrypted. This mode is designed to address the needs of hosts that exist behind the firewall.

Consider using a client address pool

Note: Client address pools can only be used with VPN definitions using Dynamic IP Restricted Client mode.

Client address pools are used to simplify the management of VPN clients by allowing the firewall to assign the following information to the client:

- An IP address for the client's virtual network adapter
- DNS servers that are available to the client
- WINS servers that are available to the client

All the client needs is:

- Client software that supports ISAKMP Mode Configuration (IKEv1) or Configuration Payload (IKEv2) exchange
- Authorization information (a client certificate, a password, etc.)
- The address of the firewall

For a client address pool, an administrator creates a range or a “pool” of IP addresses that will be used by remote peers when they attempt to make a VPN connection. When a client attempts a connection, the firewall assigns it one of the IP addresses available in the address pool. The firewall also negotiates with the client to determine other VPN requirements, such as which DNS and/or WINS servers will be made available to the client. If the negotiation is successful, the client is connected and the VPN connection is established.

Note: Not all VPN client software supports the negotiation of every client address pool parameter. Be sure to verify that your client(s) support the necessary features.

You define the number of IP addresses available in the client address pool. Even though the client may have a fixed IP address, the address used within the VPN definition is the address assigned to it from the address pool. The address pool works for both fixed and dynamic clients. See the following VPN scenarios where address pools could be used:

- [Scenario 2: Simple deployment of remote users](#)
- [Scenario 3: Large scale deployment of clients](#)

You can also create multiple client address pools. Grouping VPN clients into distinct pools allows you to limit the resources the clients in each group can access.

For information on configuring client address pools, see [Managing client address pools](#).

Choosing the appropriate authentication type

Use the following sections to determine which authentication method(s) are appropriate for your VPN:

- [About IKE VPN authentication](#)
- [What type of VPN authentication should I use?](#)
- [Guidelines for selecting a VPN authentication type](#)
- [Extended authentication for VPN](#)

About IKE VPN authentication

When using automatic key generation, after you gather the initial information for the remote end of the VPN, there is no further direct contact between you and the remote end of the VPN. Session keys are automatically and continually generated and updated based on this initial identifying information. As a result, the firewall requires a way to assure that the machine with which you are negotiating session keys is actually who they claim to be—a way to authenticate the other end of the VPN. To allow automatic key generation, the firewall offers the following authentication techniques:

Note: If you are using manual key generation, each time you generate session keys you must communicate directly with the other end of the VPN via telephone, diskette, or e-mail. By contacting the remote end of the VPN each time you change session keys, you manually verify that the remote end is actually who they claim to be.

- **Pre-shared password** – When you generate keys, the firewall and the remote end must both use the agreed-upon password, defined during the initial configuration of the VPN, to authenticate each peer.
- **Single certificate** – Single certificate authentication requires that the firewall generate a certificate and private key to be kept on the firewall and a certificate and private key to be exported and installed the peer. Each certificate, once installed on its end of a VPN connection, acts as a trust point. A single certificate (also referred to as a “self-signed certificate”) differs from certificate authority (CA) based certificates in that no root certificate is necessary.
- **Certificate Authority policy** – The firewall can be configured to trust certificates from a particular certificate authority (CA). The firewall will trust any certificate that is signed by the chosen CA and meets certain administrator-configured requirements on the identity contained within the certificate. We recommend that only locally administered certificate authorities be used in this type of policy. Certificate authorities are described further in [About Certificate/Key Management](#).

What type of VPN authentication should I use?

The Sidewinder supports four VPN authentication methods. The characteristics of a remote VPN peer determine which type of authentication best fits your VPN configuration. Extended authentication may be added to any automated authentication method for increased security.

Note: Extended authentication is **not** available for firewall-to-firewall configurations or any configuration that uses a manual key exchange.

Table 63 VPN authentication options

Authentication	Summary
Manual key VPN	<ul style="list-style-type: none">• Suitable for gateways• Authenticates using a manual key exchanged over a telephone or other secure connection—keying information is cumbersome to enter and not changed often, which reduces security• Uncommon in today's networks, but used for resolving interoperability problems with other vendors' IPsec products• <i>Cannot</i> be used for dynamic IP-assigned clients or gateways• Each VPN peer requires its own firewall VPN configuration
Automatic key shared password VPN	<ul style="list-style-type: none">• Suitable for gateways• Suited for clients when paired with a strong extended authentication, such as SafeWord PremierAccess• Primary authentication is password sharing with the VPN peer• May be used with dynamic IP-assigned clients, but the clients must be configured to use Aggressive mode or IKEv2
Automatic key single certificate VPN	<ul style="list-style-type: none">• Suitable for gateways and clients• Authenticates using a self-signed public certificate—each VPN peer must first import the corresponding peer's certificate• Ideally used for a small number of remote peers• Used with dynamic IP-assigned clients and gateways• Each peer certificate requires its own firewall VPN definition
Automatic key certificate authority-based VPN	<ul style="list-style-type: none">• Suitable for gateways and clients• Authenticates each VPN peer by using a certificate signed by a certificate authority trusted by the other peer• Ideally suited for roving client VPN peers (such as those using laptop computers)• Used with dynamic IP-assigned clients and gateways• Single firewall VPN definition can be used to administer many VPN clients
Extended authentication (XAUTH)	<ul style="list-style-type: none">• Suitable for clients• If using IKEv1, both sides of a VPN connection must use the same version of IKE and the same authentication method• If using IKEv2, each side of a VPN connection can use a different authentication method• If using IKEv1, XAUTH can be used in conjunction with Password, Certificate and Certificate Authority, and Single Certificate• If using IKEv2, XAUTH can be used as the sole authentication method

Guidelines for selecting a VPN authentication type

Follow these general guidelines when deciding which type of VPN to use:

- If the VPN peer is a third-party product, and all other types of VPN methods do not work, try the manual key VPN.
- For a small number of VPN peer clients with dynamically assigned IP addresses, the single certificate VPN is a cost-effective solution. A shared password VPN in conjunction with extended authentication is also an option.
- If the VPN peer has a static IP address, the pre-shared password VPN is the easiest to configure. Extended Authentication would not be used in a gateway to gateway configuration as there is no one to provide the challenge/response.
- If there is a large number of VPN peer clients with dynamically assigned IP addresses (such as a traveling sales force), the CA-based VPN is often the easiest to configure and maintain. Another popular option is to use a pre-shared password VPN in conjunction with extended authentication.

Extended authentication for VPN

The extended authentication (XAUTH) option requires the *person* requesting the remote access VPN connection to validate their identity. The extended authentication option is most useful if you have traveling employees that connect remotely to your network using laptop computers. If a laptop computer is stolen, without extended authentication it might be possible for an outsider to illegally access your network, since the information needed to establish the VPN connection (the self-signed certificate, etc.) is saved within the VPN client software. When extended authentication is used, however, a connection will not be established until the user enters an additional piece of authentication information that is not saved on the computer—either a one-time password, passcode, or PIN. This additional level of authentication renders the VPN capabilities of the laptop useless when in the hands of a thief.

Implementing extended authentication on your Sidewinder is a simple two-step process.

1 Specify the authentication method(s) that are available on your firewall:

- a** Select **Policy > Rule Elements > Services**.
- b** Select **isakmp** from the list and then click **Modify**.
- c** Click **Properties**.
- d** In the **Allowed XAUTH Methods** field, enable the desired methods.

For information on supported methods, see [Types of authentication methods](#).

2 Select XAUTH as an authentication method in the VPN Definitions: Remote Authentication tab. See [About the VPN Properties: Remote Authentication tab](#) for more details.

Note: Extended authentication must also be enabled on the remote client. See your client software documentation for information on configuring and enabling extended authentication.

Ordering VPN definitions

The order of definitions in the VPN Definitions window affects how packets are matched to definitions. The first definition that matches a connection request is used to allow or deny that connection.

For example, the table below shows the first two positions in a list of VPN definitions. If a packet has a source IP address of 10.69.106.5 and a destination IP address of 10.69.104.20, it is matched to the first definition in the list. The search is stopped before the packet is compared to the more precise match in the second definition.

Table 64 VPN definition ordering

Position	Local Network	Remote Network
1	10.69.106.0/24	10.69.104.0/24
2	10.69.106.5/32	10.69.104.20/32

You can also select certain traffic to bypass IPsec policy evaluation and be sent outside the encrypted tunnel. Other security policy rules will apply to this traffic. You select this option on the VPN Definitions: General tab.

Example: Traffic between two networks at two different sites is encrypted, but you want traffic to and from the web server to be sent outside the encrypted tunnel. You would configure a Bypass definition and place it in front of a more general definition in the VPN Definitions list.

The table below shows a VPN Definitions list with a Bypass VPN definition in the first position.

Table 65 Bypass definition in the VPN Definitions list

Position	Action	Local Network	Remote Network
1	Bypass	10.69.106.0/24	10.69.104.20/32
2	IPsec	10.69.106.0/24	10.69.104.0/24

Restricting VPN access with a virtual burb

Use the sections below to familiarize yourself with virtual burb concepts and to determine if you should use a virtual burb to restrict the access allowed by your VPN:

- [About virtual burbs](#)
- [Using rules to direct VPN traffic](#)
- [VPNs with rules using NAT and redirection](#)

About virtual burbs


A termination burb is the burb in which VPN traffic transitions between plain-text and encrypted data. You can increase security and control of that transition by using a virtual burb as your termination burb.

A virtual burb is a burb that does not contain a network interface card (NIC). VPNs terminated in a virtual burb require policy rules to take traffic from the virtual burb to and from the internal burb. Using a virtual burb separates VPN traffic from non-VPN traffic, and it allows you to enforce a security policy that applies strictly to your VPN users.


Consider a VPN policy that is implemented without the use of a virtual burb. Not only will VPN traffic mix with non-VPN traffic, but there is no way to enforce a different set of rules for the VPN traffic. This is because proxies and rules are applied on a burb basis, not to specific traffic within a burb. By terminating the VPN in a virtual burb, you effectively isolate the VPN traffic from non-VPN traffic. In addition, you are able to configure a unique set of rules for the virtual burb that allows you to control precisely what your VPN users can or cannot do. [Figure 309](#) illustrates this concept.

Figure 309 Virtual burb vs. a non-virtual burb VPN implementation

VPN without a virtual burb



VPN with a virtual burb



- Traffic originating from the remote peer must now traverse a proxy from the virtual burb to the internal burb.
- A separate rule is used for traffic originating from the internal network destined for the remote peer.

Once the traffic is decrypted on the firewall, it must also traverse a proxy configured for traffic from the virtual burb to the internal burb. This allows administrators to have a finer grain of control of the services allowed inbound and outbound from their controlled networks.

You can define up to 63 physical and virtual burbs. For example, if you have two distinct types of VPN definitions and you want to apply a different set of rules to each type, create two virtual burbs, then configure the required proxies and rules for each virtual burb.

One question that might come to mind when using a virtual burb is: "How does VPN traffic get to the virtual burb if it doesn't have a network card?" All VPN traffic originating from the Internet initially arrives via the network interface card in the Internet burb. A VPN connection, however, can internally route and logically terminate VPN traffic in any burb on the firewall. By defining a VPN connection to terminate the VPN in a virtual burb, the VPN traffic is automatically routed to that virtual burb within the firewall. Thus, the trusted network now recognizes the virtual burb as the source burb for your VPN traffic. From the virtual burb, a proxy and rule are needed to move the traffic to a trusted burb with network access.

See [Creating and using a virtual burb with a VPN](#) for more information.

Using rules to direct VPN traffic

You can use VPN definitions in conjunction with rules to gain more control over your network security policy.

There are several advantages to using rules to restrict traffic bound for a VPN, including:

- You can control which services are allowed through the VPN by filtering of traffic based on protocol and port.
- Rule-level auditing of inbound/outbound connections is included. Since the VPN audit only indicates when the encrypted connection is established or torn down, rules can provide the added benefit of auditing for each connection between a VPN's protected hosts.
- Traffic can be inspected for malicious content by applying application defenses to the rules. All features of the firewall rules can be used to protect internal resources and monitor what external resources are being used.
- You use rules to restrict access at the user-level, such as per-user authentication.

You can use virtual burbs and NAT/redirection to further control VPN traffic.

VPNs with rules using NAT and redirection

How does the firewall know that a packet between two networks is supposed to go from the internal burb to the VPN burb? Understanding how the firewall makes this determination leads to understanding how NAT and redirection can change how the firewall views a packet. When the firewall first receives a packet, it has three pieces of information:

- The source address/port of the packet
- The destination address/port of the packet
- The burb on which the packet was received

Using this information, the firewall determines the burb to which the packet is destined. The destination burb, as well as the other pieces of packet information, are then used to make a determination on what policy rule applies to this packet.

The firewall determines the packet's destination burb by consulting the IPsec security policy database (SPD). Each SPD entry contains a local or source network and remote or destination network along with the termination burb associated with that VPN. These parameters are configured in the Admin Console's VPN Definitions area. The firewall compares the inbound packet's source and destination against the networks in the SPD and, upon a successful match, returns the termination burb. This termination burb is used as the packet's destination burb.

If a packet fails to match an SPD entry, the firewall consults the routing table to find a route associated with the destination address of the packet. The firewall then uses the burb associated with the matching route as the destination burb.

Figure 310 illustrates the information the packet provides the proxy. The firewall uses the IPsec policy engine, policy rules, and the existing packet information to route the packet to its intended destination.

Figure 310 Information initially provided by the packet

source burb: internal to packet	
header	
IP	source IP: 192.168.100.1 dest IP: 172.17.111.1
TCP	source port: 34567 dest port: 23
additional packet information	

As described above, the first action that the firewall takes is to determine the intended destination burb for the packet. This is done by querying the IPsec policy table and routing table. For this example, a matching IPsec policy is found, as shown in Figure 311.

Figure 311 Matching packet information to determine the destination burb

VPN Definitions: VPN Properties

General Remote Authentication Local Authentication Crypto Advanced

Name: my_vpn Enabled: ☒ Yes ☐ No

Mode: Fixed IP Client Address Pool: <disabled>

IKE Version: ☒ V1 ☐ V2

Encapsulation: ☒ Tunnel ☐ Transport

Burb: VPN

Local

Local IP: Use Localhost IP localhost

Local Network / IP: 192.168.100.0 / 24

New Modify Delete

Remote

Remote IP:

Remote Network / IP: 172.17.0.0 / 16

New Modify Delete

Comments:

Add Close Help

Based on the local network and remote network, the packet's destination burb is set to VPN.

When the destination burb is determined, it is used with the source burb and the source and destination addresses to query the policy rules for a match. For this example, a matching Telnet rule is found (see [Figure 312](#)).

Figure 312 Matching policy rule

Source and destination burbs and source and destination addresses match the packet to this policy rule.

If the matching policy rule specifies that it must apply either NAT or redirection, further action is taken by the firewall prior to sending the packet to the intended destination.

First, we examine NAT. For example, suppose the rule has returned a NAT address of 192.168.200.1. The firewall rewrites the original packet so that it looks similar to [Figure 313](#).

Figure 313 Using NAT to change the source burb and address

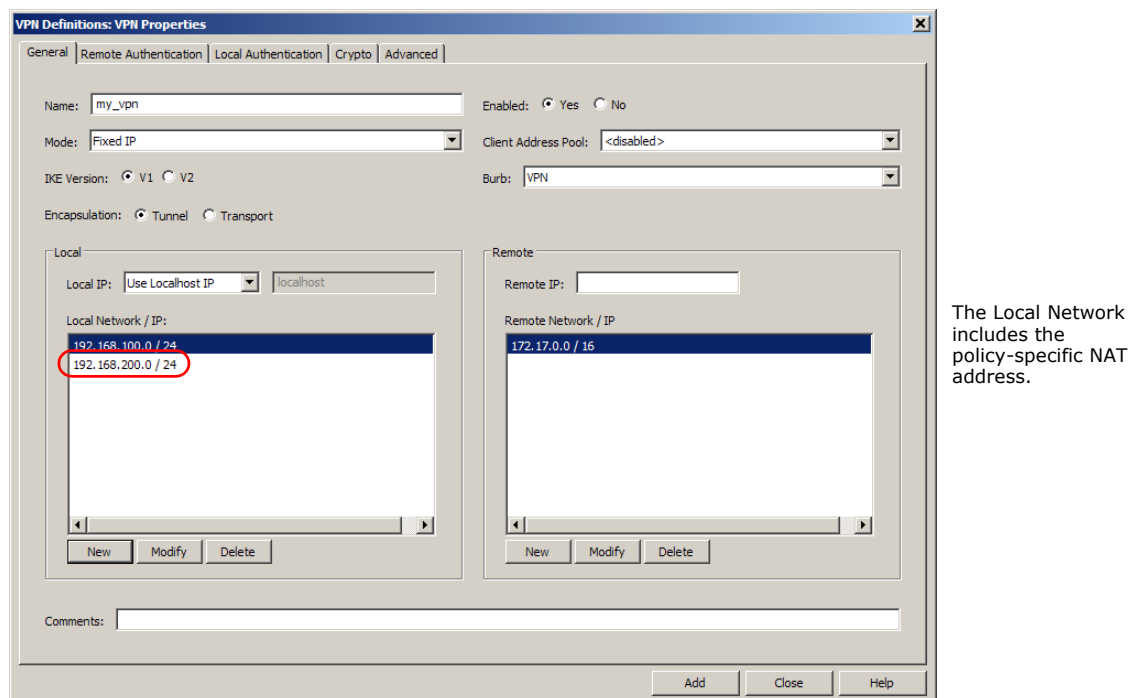
source burb: vpn packet	
header	
IP	source IP: 192.168.200.1 dest IP: 172.17.111.1
TCP	source port: 45678 dest port: 23
additional packet information	

Notice the change in the source address and the source burb. The firewall rewrites the source address to match the NAT address returned by the rule query. The source burb has changed due to the firewall creating a new connection in the destination burb used in the original rule query. The result is that the firewall changed the source address and sent the packet into the vpn burb, which is a virtual burb. The kernel then examines the packet to determine which action to take next (for example, encrypt via VPN tunnel, route out via interface, send to proxies, etc.).

While the original packet may have matched the VPN policy in [Figure 311](#), the transformed packet no longer does. Therefore, the packet will not be encrypted via the VPN definition. There is no interface associated with a virtual burb, so the packet cannot be sent directly. It must then be sent up the stack to any potential proxies, where it is likely either NSS will issue a netprobe (if no proxy is listening on the given port in the virtual burb) or a proxy will pick up the connection and go through the process of the policy evaluation, most likely ending up with a deny rule.

How can this situation be avoided for policies that require NAT to mask protected network addresses? The fundamental resolution is that the VPN policy must match the packet both before and after any address translation. This is done by either selecting a NAT address that falls within the range of one of the local networks of the VPN definition (for example, a NAT address of 192.168.100.254 would work in the above example) or by redefining the VPN policy to take into account any NAT translation. In our example, the VPN definition could be changed as shown in [Figure 314](#).

Figure 314 Matching your security policy with your NAT policy



The *Local Network* has been redefined to include both the protected subnet for which packets will be received and the specific NAT address associated with the rule.

Note: When NAT is being performed on a rule, the **localhost (Host)** option cannot be used for connections where the destination burb is a virtual burb. This option translates the source address of a connection to the interface address associated with the destination burb of that connection. In the case where the destination burb is a virtual burb, there can be no translation using **localhost (Host)** because no interface exists in a virtual burb. The result of attempting to perform this type of NAT in a virtual burb will cause the connection to fail. For best results, select a specific IP address when configuring NAT.

Next, consider how redirection interacts with proxy connections and VPN definitions. Again, consider the same original packet in [Figure 315](#) and the original policy in [Figure 316](#) (shown previously in [Figure 310](#) and [Figure 311](#)).

Figure 315 Original packet

source burb: internal to <i>packet</i>	
header	
IP	source IP: 192.168.100.1 dest IP: 172.17.111.1
TCP	source port: 34567 dest port: 23
additional packet information	

Figure 316 Original VPN definition

The screenshot shows the 'VPN Definitions: VPN Properties' window with the 'General' tab active. The 'Name' field contains 'my_vpn'. The 'Enabled' radio buttons have 'Yes' selected. The 'Mode' dropdown is set to 'Fixed IP'. The 'Client Address Pool' dropdown is set to '<disabled>'. The 'IKE Version' radio buttons have 'V1' selected. The 'Encapsulation' radio buttons have 'Tunnel' selected. The 'Burb' dropdown is set to 'VPN'. Under the 'Local' section, 'Local IP' is set to 'Use Localhost IP' with 'localhost' in the adjacent field. The 'Local Network / IP' list contains '192.168.100.0 / 24'. Under the 'Remote' section, the 'Remote IP' field is empty. The 'Remote Network / IP' list contains '172.17.0.0 / 16'. At the bottom, there are 'New', 'Modify', and 'Delete' buttons for both the Local and Remote sections, and a 'Comments' text area. The 'Add', 'Close', and 'Help' buttons are at the bottom right.

Assume that the matching policy rule has NAT disabled, but specifies a redirection address of 172.16.222.1. To perform the redirection, the firewall needs to perform two tasks. First, the firewall rewrites the packet's destination address using the redirection address (172.16.222.1 in our example). The transformed packet now appears as shown in [Figure 317](#).

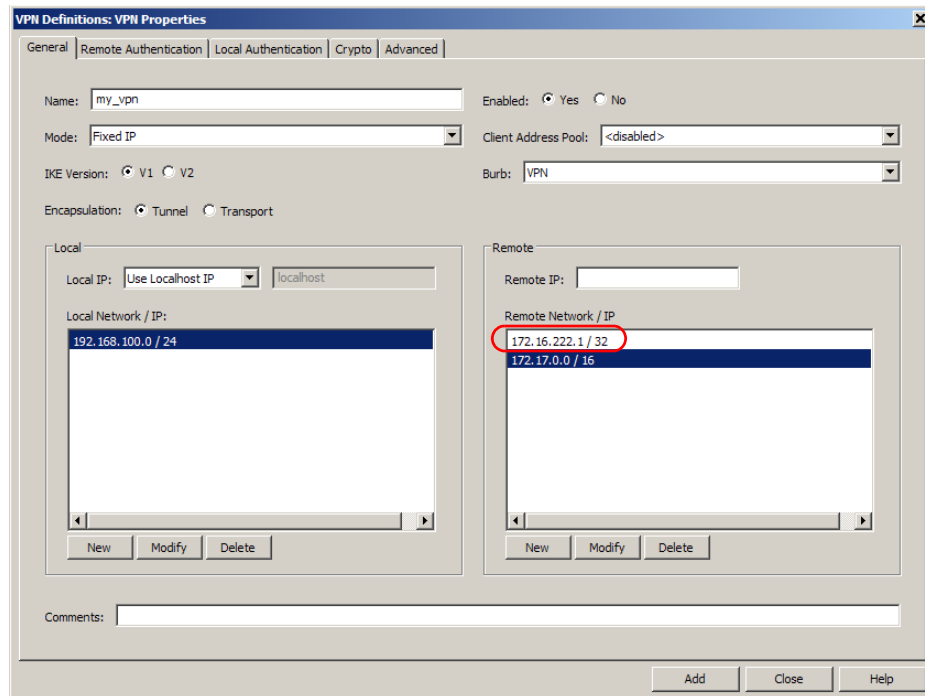
Figure 317 Using redirection to change the destination address

source burb: internal to <i>packet</i>	
header	
IP	source IP: 192.168.100.1 dest IP: 172.17.222.1
TCP	source port: 34567 dest port: 23
additional packet information	

Second, because we changed the destination address of the packet, the firewall needs to recompute the destination burb for the outbound connection. In this case, as with NAT, the packet will no longer match the intended VPN, so the destination burb will not be the virtual VPN burb. Failing to find a matching IPsec policy means that routing is then used to determine the destination burb, based on the destination address. Assuming that a matching route (or default route) exists, the destination burb will be the burb associated with that route. For our example, we assume that the new destination address falls through to the default route, and the default route is out the external burb. It is this burb that the firewall will use to create the outbound connection.

Once again, to remedy this situation and force the transformed packet to match the intended VPN, the VPN definitions and rule-defined address translation must be aligned so that the packet matches the intended VPN both before and after NAT and redirection have been applied. In our example, this means either choosing a redirection address within the existing VPN policy (for example, 172.17.1.1) or changing the VPN policy to include the new redirection address. As with NAT, the specific redirection address has been added to the policy definition in [Figure 318](#) (this time to the remote or destination networks' definition).

Figure 318 Matching your security policy with your redirection policy



The Remote Network includes the policy-specific redirect address

If a rule specified both NAT and redirection, the VPN policy would have to reflect changes to both local and remote network definitions.

Creating VPN policy

This section provides a high level overview of the tasks required to create a Sidewinder VPN. Because the configuration possibilities are substantial, this section does not cover specifics. Before you use this section to create a new VPN, make sure you have planned your VPN configuration. See [Planning your VPN](#).

Note: For specific information about the user interface, see [VPN user interface reference](#).

Perform the following high level tasks to set up a VPN:

- 1 Plan the VPN implementation from end to end. See [Planning your VPN](#).
- 2 Configure the ISAKMP server: Select **Policy > Rule Elements > Services** and double-click **isakmp**. Click **Properties**, then configure the server options and the XAUTH method. See [Setting up the ISAKMP service](#).
- 3 [Conditional] If you are creating a Dynamic IP Restricted Client VPN and you want the firewall to assign resources to the remote peer, create a client address pool: Select **VPN Configuration > Client Address Pools**. See [Configuring client address pools](#).
- 4 Define the VPN: Select **VPN Configuration > VPN Definitions**. See [Managing VPN definitions](#).
Tip: If you want to restrict VPN traffic using a virtual burb, We recommend that you simplify the configuration process by first configuring a new VPN without a virtual burb. After you have verified that the VPN is working, you can then add a virtual burb to restrict access.
- 5 [Conditional] If you want to restrict VPN access using a virtual burb, create a virtual burb and terminate the VPN in that burb. See [Creating and using a virtual burb with a VPN](#).
- 6 After your VPNs are complete, create a configuration backup: Select **Maintenance > Configuration Backup**.

Setting up the ISAKMP service

If you are using automatic key exchange, you must do the following before using any automatic key VPNs:

- **Configure the ISAKMP (Internet Security Association and Key Management Protocol) server** – See [Managing the ISAKMP server](#).
- **Create a rule to allow access to the ISAKMP service** – Select **Policy > Rules**, then click **New** to create a rule for the ISAKMP service. The ISAKMP rule must contain the following values:
 - **Service** – isakmp (ISAKMP Server)
 - **Source Burb** – The burb receiving traffic from the VPN peer(s)
 - **Source endpoint** – <Any> (or addresses of remote VPN peer(s))
 - **Destination Burb** – Match the source burb setting

Configuring client address pools

To create a new client address pool:

1 Select **Network > VPN Configuration > Client Address Pools**.

2 Create a new address pool, configuring the following tabs:

- **Subnets** tab:
 - Add a virtual subnet for each range of IP address you want the firewall to assign to remote VPN clients.
 - Add a local subnet for each network you want to allow remote VPN clients to communicate with.
- **Servers** tab:
 - Add any DNS servers that the firewall should assign to the remote VPN clients.
 - Add any NBNS or WINS servers that the firewall should assign to the remote VPN clients.

• [Optional] **Fixed IP Map** tab:

If you want the firewall to assign fixed IP addresses to selected remote VPN clients, add an entry for each client to the Fixed IP Map table.

3 Save your changes.

The new client address pool can now be selected on any VPN definition that uses Dynamic IP Restricted Client mode.

Creating and using a virtual burb with a VPN

To create a virtual burb and use it with a VPN:

1 Create the virtual burb:

a Select **Network > Burb Configuration**.

b Click **New**.

- In the **Burb Name** field, type the name for your virtual burb.
- Select the appropriate connection options.
- Click **OK**.

2 Save your changes.

3 Configure the rules: Select **Policy > Rules** and define the rules that allow access to and from the virtual burb.

The virtual burb should be specified as either the source or destination burb, depending on the type of rule being defined. See [Using rules to direct VPN traffic](#) for more information.

4 Terminate the desired VPN connection(s) in the virtual burb:

a Select **VPN Configuration > VPN Definitions** and open the VPN definition you want to terminate in the virtual burb.

b From the **Burb** drop-down list, select the virtual burb you created in [Step 1](#).

c Click **OK** and then save your changes.

See [Managing VPN definitions](#) for information on creating or modifying a VPN definition.

For more information on virtual burbs, see [Restricting VPN access with a virtual burb](#).

VPN user interface reference

The following sections describe the VPN user interface:

- [Managing VPN definitions](#)
- [Managing client address pools](#)
- [Managing the ISAKMP server](#)

Managing VPN definitions

Use the VPN Definitions window to manage, view, and order the current VPN definitions on Sidewinder. You can also check the status of VPNs, and create, modify, or delete VPN definitions.

To view and order VPN definitions, select **Network > VPN Configuration > VPN Definitions**. The VPN Definitions window appears.

Use the toolbar to perform the following functions:

Figure 319 VPN Definitions toolbar




Table 66 below describes the toolbar buttons in detail.

Table 66 VPN toolbar

Icon	Action
New	Create a VPN definition by clicking New . The VPN Properties window appears. Enter distinct information about a VPN definition on the five tabs of the window. See Creating VPN policy for details.
Modify	Make changes to a VPN definition by selecting it and clicking Modify . The VPN Properties window appears. Make your changes in the VPN Properties window. Note: Read-only administrators can click View to view a VPN definition.
Delete	Delete an existing VPN definition by selecting it and clicking Delete .
Duplicate	Duplicate an existing VPN definition by selecting the original definition and clicking Duplicate . The default name of the new item is Copy_of_x, where x is the original definition's name.
Rename	Rename an existing VPN definition by selecting it and clicking Rename . Type the new name in the pop-up window and click OK .
Enable/Disable	Enable or disable a VPN definition by selecting it and clicking the appropriate icon. A disabled VPN definition is grayed out in the list. To select multiple consecutive definitions to enable or disable at one time, select the first definition, then press the Shift key while selecting the last definition. To select multiple non-consecutive definitions at one time, press the Ctrl key while selecting each desired definition.
Move Up/Move Down	Set the matching order of VPN definitions by selecting a definition in the list and clicking the Move Up and Move Down arrows to place it in the desired position for matching purposes. The first definition that matches a connection request is used to allow or deny that connection.
Find	Search for specific elements in the list by typing your search criteria in the Find field. Objects with matching elements appear in the list.
VPN Status	Click VPN Status to view the status of all configured VPNs in a pop-up window. Click Refresh Now to update the information in the pop-up window.

To create or modify VPN definitions:

Note: You cannot create more than 16,383 VPN definitions.

- 1 Select **Network > VPN Configuration > VPN Definitions**. The VPN Definitions window appears.
- 2 Click **New** or select an existing definition from the list and click **Modify**. The VPN Properties window appears.
The VPN Properties window contains five tabs that are used to enter distinct information about a VPN definition.
- 3 Enter the necessary information to create a VPN definition.
- 4 When you are done entering information, save your changes.
- 5 Click **Add** to close the VPN Properties window and add the VPN definition to the VPN Definitions list.
- 6 Use the arrows to place the new definition in the desired position. See [Managing VPN definitions](#) for more information.

The VPN Properties window contains the following tabs:

- **General** – Enter basic information about the VPN definition. See [About the VPN Properties: General tab](#) for details.
- **Remote Authentication** – Define the authentication method that will be used by the remote peer in this VPN definition. You also define the characteristics of the selected authentication method.
See [About the VPN Properties: Remote Authentication tab](#) for details.
- **Local Authentication** – Define the authentication method that will be used by the firewall in this VPN definition. You also define the characteristics of the selected authentication method.
See [About the Local Authentication: Password view](#) for details.
- **Crypto** – Define the IPsec cryptographic properties according to the type of key exchange:
 - Manual key exchange – Define manual authentication for this VPN definition. See [About the VPN Properties: Crypto tab \(Manual key exchange\)](#) for details.
 - Automatic key exchange – Define the IPsec cryptographic and hashing algorithms used to in this VPN definition. See [About the VPN Properties: Crypto tab \(Automatic key exchange\)](#) for details.
- **Advanced** – Define some of the more advanced points of a VPN definition on this tab, including NAT Traversal. Only administrators that are highly schooled in VPN should modify the information on this tab. This information is used only with automatic key exchange.

About the VPN Properties: General tab

Use the General tab to enter basic information about the VPN definition.

Figure 320 VPN Properties: General tab

The screenshot shows the 'VPN Definitions: VPN Properties' dialog box with the 'General' tab selected. The dialog has five tabs: General, Remote Authentication, Local Authentication, Crypto, and Advanced. The General tab contains the following fields and controls:

- Name:** A text input field.
- Enabled:** Radio buttons for 'Yes' (selected) and 'No'.
- Mode:** A drop-down menu currently set to 'Fixed IP'.
- Client Address Pool:** A drop-down menu currently set to '<disabled>'.
- IKE Version:** Radio buttons for 'V1' (selected) and 'V2'.
- Burb:** A drop-down menu currently set to 'internal'.
- Encapsulation:** Radio buttons for 'Tunnel' (selected) and 'Transport'.
- Local section:** Includes a 'Local IP' field with a 'Use Localhost IP' button and a text field containing 'localhost'. Below it is a 'Local Network / IP' list box with 'New', 'Modify', and 'Delete' buttons.
- Remote section:** Includes a 'Remote IP' field and a 'Remote Network / IP' list box with 'New', 'Modify', and 'Delete' buttons.
- Comments:** A text area at the bottom.
- Buttons:** 'Add', 'Close', and 'Help' buttons at the bottom right.

To configure the General tab, follow the steps below.

- 1 In the **Name** field, type the name of this VPN.
 - 2 In the **Enabled** field, select **Yes** to enable this VPN definition, or select **No** to disable it.
 - 3 From the **Mode** drop-down list, select how the remote end is operating:
 - **Fixed IP** – Select this option if the IP address of the remote end is always the same. You must also provide the IP address of the remote end in the **Remote IP** field.
 - **Dynamic IP Client** – Select this option if the remote end is a device whose IP address is not fixed. Example: A salesperson that gains Internet access from a laptop.
 - **Dynamic IP Restricted Client** – Select this option if the remote end is a device whose IP address is not fixed. The difference from the *Dynamic IP Client* option is that the you restrict the range of IP addresses available to the remote end by using either the **Client Address Pool** field or the **Dynamic Virtual Address Range** field. An example is a salesperson that gains Internet access from a laptop.
- Note:** You can use Dynamic IP Client or Dynamic IP Restricted Client only if automatic key management is used.
- **Manually keyed VPN** – Select this option if you want to exchange session keys manually (for example, over the phone). You configure specific properties of the manual key exchange on the Crypto tab.
- See [About the VPN Properties: Crypto tab \(Manual key exchange\)](#) for detailed information.
- **Bypass** – Select this option if you want certain traffic to bypass IPsec policy evaluation and be sent outside the encrypted tunnel. Other security policy rules will apply to this traffic. See [Managing VPN definitions](#) for more information.

- 4 [Conditional] If you want remote peers to make connections using only the IP addresses contained within one of the available client address pools, select a client address pool from the **Client Address Pool** drop-down list. The firewall then selects an IP address from the available pool and assigns it to the client.
 - This field is available only if the Mode is **Fixed IP** or **Dynamic IP Restricted Client**.
 - See [Example VPN Scenarios](#) for information on creating a client address pool.
- 5 If you are using automatic key exchange, select which **IKE Version** to use. See [Choose an IKE version](#).
- 6 From the **Burb** drop-down list, select the burb you want to assign this VPN to. The firewall terminates each VPN in a burb so that access rules can be applied to the VPN.

See [About virtual burbs](#) for information about virtual burbs and VPNs.
- 7 In the **Encapsulation** field, select one of the following:
 - **Tunnel** – The more common form of VPN encapsulation. Both the data and the source and destination IP addresses are encrypted within the encapsulated payload.
 - **Transport** – Transport mode encrypts the data but the source and destination IP addresses are not concealed.See [Determine the encapsulation method](#) for a more detailed explanation of these terms.
- 8 In the **Local IP** field, indicate which IP address to use as the local gateway by selecting one of the following:
 - **Use Localhost IP** – Select this option to have the firewall assign the IP address. The firewall uses its routing table to automatically determine which interface or alias address is associated with a route to reach the remote gateway.
 - **Specify IP** – Select this option to configure a specific IP address. This IP address should be one of the firewall's interface or alias addresses, and that interface must have a route to reach the remote gateway.

Note: If configuring a VPN for an HA cluster, be sure to use the localhost option or specify an alias shared by the cluster.
- 9 To add or modify a local network address to the **Local Network/IP** list, click **New** or select an address from the list and click **Modify**. The Local Network List window appears:
 - a In the **IP Address** field, type the IP address used in this VPN definition.
 - b In the **Number of bits in Netmask** field, use the up/down arrows to select the number of bits that are significant in the network mask. The value specified is used to identify the network portion of the IP address.The **Local Network/IP** list shows the network names or IP addresses the firewall can use in a VPN definition. The addresses in this list and the addresses in the **Remote Network/IP** list together identify allowed and reachable addresses for this VPN tunnel.
- 10 [Conditional] In the **Remote IP** field, type the IP address of the remote peer. This field is available for **Fixed IP** and **Manually Keyed VPN** mode.
- 11 [Conditional] You can add or modify an entry in the **Remote Network / IP** list if the Mode is **Fixed IP**, **Manually Keyed VPN**, or **Bypass**: Click **New** or select an address in the list and click **Modify**. The Remote Network List window appears:
 - a In the **IP Address** field, type the IP address used in this VPN definition.
 - b In the **Number of bits in Netmask** field, use the up/down arrows to select the number of bits that are significant in the network mask. The value specified is used to identify the network portion of the IP address.The networks configured here represent real networks located behind the remote VPN peer. The addresses in this list and the addresses in the **Local Network/IP** list together identify allowed and reachable addresses for this VPN tunnel.
- 12 [Conditional] If the Mode is **Dynamic IP Restricted Client**, you can add or modify an entry to the **Dynamic Virtual Address Range** list: click **New** or select an address range and click **Modify**. The Dynamic Virtual Address Range List window appears.

- a** In the **IP Address** field, type the IP address range a client can use when initiating a VPN connection.
- b** In the **Number of bits in Netmask** field, use the up/down arrows to select the number of bits that are significant in the network mask. The value specified is used to identify the network portion of the IP address.
- This list defines the range of addresses a client can use when initiating a VPN connection.
- These are the addresses of client machines on their remote networks. They can be internet-routable addresses or virtual addresses.
- With this option the client assigns their own IP address, although the address must be within the approved address range.

13 [Optional] In the **Comments** field, type a short description for this VPN definition.

About the VPN Properties: Remote Authentication tab

Use this tab to define the authentication method that the remote peer will use to authenticate to the firewall in this VPN definition. Authenticating the remote peer prevents access to the VPN from Internet hosts masquerading as the remote peer.

From the **Remote Authentication Method** drop-down list, select one of the following methods. The settings for the selected authentication method will populate the window.

- **Password** – Select this option if you and the remote end want to use a password to verify the key exchange. The same password must be used on both ends of this connection.

See [About the Remote Authentication: Password view](#) for detailed information.

- **Certificate + Certificate Authority** – Select this option if you want to use one or more trusted CAs and Remote Identities to validate the certificate of the remote end. This method is commonly used by organizations that have many remote users who must access resources behind the firewall.

See [About the Remote Authentication: Certificate + Certificate Authority view](#) for detailed information.

- **Single certificate** – Select this option if you want to validate the remote end using a self-signed certificate generated by the firewall, or using a certificate generated by a CA server. This method is commonly used by organizations that have a small number of people that travel but need secure access to your network.

See [About the Remote Authentication: Single Certificate view](#) for detailed information.

- [Conditional] **XAUTH** – Select this option to verify the *person* rather than the *machine*. XAUTH uses the authentication methods configured to authenticate proxy users connecting through the firewall.

Note: XAUTH must be enabled in the ISAKMP service Properties window.

The use of XAUTH is determined by the version of IKE you select:

- **XAUTH + Password/Certificate + Certificate Authority/Single Certificate** – With IKEv1, XAUTH is an added layer of security to be paired with one of the other remote authentication options. These are automatic key exchanges.
- **XAUTH** – With IKEv2, XAUTH is a separate remote authentication method available for selection.

About the Remote Authentication: Password view

Use this view to define the remote peer's identity and the password that the remote peer uses in this VPN definition to authenticate to the firewall. The firewall must use the same password to authenticate to the remote peer.

Note: If you are using IKEv1, Password-based authentication should be used only with fixed IP-configured VPN or with extended authentication.

To define the remote peer's identity and password:

- 1** In the **Enter Remote Password** field, type the password to be used each time automatic key exchange takes place.
- 2** In the **Verify Remote Password** field, confirm the password.
- 3** In the **Remote Identity** section, select an identity the remote peer will be required to use to authenticate to the firewall. The firewall uses this information to determine the password the remote peer should use.

- Select **Gateway IP Address (not specified)** to require the remote peer to use its gateway address as its identity to authenticate to the firewall. This option is strongly recommended.
- Select **Remote One or More Remote IDs From List** to require the remote peer to use the configured remote identity.

Note: The remote identity is optional for Fixed IP VPN definitions because the firewall can use the IP address to determine who the remote peer is and thus what password the remote peer should be using.

- 4 [Optional] To create a new remote identity, click **Remote Identities**. The Remote Identities window appears. Create the new identity, then click **Close** to return to the Remote Authentication: Password view.

About the Remote Authentication: Certificate + Certificate Authority view

Use this view to define the Certificate and Certificate Authority that the remote peer uses in this VPN definition to authenticate to the firewall.

- 1 In the Certificate Authorities section, select a Certificate Authority to use for this VPN definition. The remote VPN peer is required to use a certificate that was signed by one of the configured CAs for this VPN definition.

To add Certificate Authorities to this list:

- a In the Certificate Authorities section, click the **Certificate Authorities...** button. The Certificate Authorities window appears.
- b Create the new Certificate Authority, then click **Close** to return to the Remote Authentication: Certificate + Certificate Authority view.

See [Managing certificate authorities](#) for details. You can add several Certificate Authorities to this list.

- 2 In the Remote Identities section, select a remote identity to use for this VPN definition. This allows the system administrator to further restrict access.

To create a new remote identity:

- a Click **Remote Identities**. The Remote Identities window appears.
- b Create the new identity, then click **Close** to return to the Remote Authentication: Certificate + Certificate Authority view.

See [Configuring and displaying remote identities](#) for details. You can add several remote identities to this list.

About the Remote Authentication: Single Certificate view

Use this view to define the single certificate that the remote peer uses in this VPN definition to authenticate to the firewall. From the **Remote Certificate** drop-down list, select the certificate used on the remote end of the VPN. The remote VPN peer will be required to use this certificate to authenticate.

- The values used as the selected remote certificate appear in the field below the drop-down list. This value is filled in automatically using the information from the selected certificate. The field cannot be modified.
- To create or import a certificate for this definition:
 - a** Click **Remote Certs...**. The Remote Certificates window appears.
 - b** Create or import the new certificate, then click **Close** to return to the Remote Authentication: Single Certificate view.

See [Configuring and displaying remote certificates](#) and [Importing firewall certificates](#) for details. You can add several remote identities to this list.

About the Remote Authentication: XAUTH view

[Applicable to IKEv2 only] Use this view to define extended authentication (XAUTH) that the remote peer uses in this VPN definition to authenticate to the firewall.

In the list of remote identities, select the identities who will use XAUTH to authenticate to the firewall. A check mark appears next to selected remote identities.

- XAUTH must be enabled in the ISAKMP server before it can be used in a VPN definition. See [Managing the ISAKMP server](#) for details.
- To create a new remote identity:
 - a** Click **Remote Identities**. The Remote Identities window appears.
 - b** Create the new identity, then click **Close** to return to the Remote Authentication: Certificate + Certificate Authority view.

See [Configuring and displaying remote identities](#) for details. You can add several remote identities to this list.

About the Local Authentication: Password view

Use this view to define the password that the firewall uses in this VPN definition to authenticate to the remote peer.

The selections available on the Local Authentication tab are dependent on the version of IKE being used:

- **IKEv1** – The firewall must use the same authentication method as the remote peer. For example, if you select Password on the Remote Authentication tab, Password is automatically selected as the Local Authentication Method, and cannot be changed on the Local Authentication tab.
- **IKEv2** – The firewall can use a different authentication method than the remote peer. For example, if you select Password on the Remote Authentication tab, you can select Certificate on the Local Authentication tab.

To define a password:

- 1** [IKEv2 definition only] From the **Local Authentication Method** drop-down list, select **Password**.
 - If you are using IKEv1, this selection matches the remote method by default and cannot be changed in this view.
 - If you are using IKEv2 and you select **Certificate**, see [About the Local Authentication: Certificate view](#) for details.
- 2** In the **Enter Local Password** and **Verify Local Password** fields, type and verify the password the firewall uses to authenticate to the remote peer.

If Password is the remote authentication method, these fields are automatically populated with the remote password and cannot be modified here.
- 3** From the **Local Identity Type** drop-down list, select the type of identity to use when identifying the firewall to the remote peer:
 - E-Mail
 - Fully Qualified Domain Name
 - IP Address

Note: E-mail addresses are not recommended, as they are rarely used in the context of a security gateway.

- 4** In the **Value** field, type the actual value used as the firewall identity.

The value must be of the type selected as the **Local Identity Type**. For example, if you selected **IP Address** in the **Local Identity Type** drop-down list, you must type an IP address in the **Value** field.

About the Local Authentication: Certificate view

Use this view to define the certificate that the firewall uses in this VPN definition to authenticate to the remote peer.

The selections available on the Local Authentication tab are dependent on the version of IKE being used:

- **IKEv1** – The firewall must use the same authentication method as the remote peer. For example, if you select Certificate on the Remote Authentication tab, Certificate is automatically selected as the Local Authentication Method, and cannot be changed on the Local Authentication tab.
- **IKEv2** – The firewall can use a different authentication method than the remote peer. For example, if you select Password on the Remote Authentication tab, you can select Certificate on the Local Authentication tab.

To define a certificate:

- 1 [IKEv2 definition only] From the **Local Authentication Method** drop-down list, select **Certificate**.
 - If you are using IKEv1, this selection matches the remote method by default and cannot be changed in this view.
 - If you are using IKEv2 and you select **Password**, see [About the Local Authentication: Password view](#) for details.
- 2 From the **Local Certificate** drop-down list, select the certificate used on the firewall end of the VPN. To create or import a certificate for this definition:
 - a Click **Local Certs....** The Firewall Certificates window appears.
 - b Create or import the new certificate, then click **Close** to return to the Local Authentication: Certificate viewSee [Configuring and displaying remote certificates](#) and [Importing firewall certificates](#) for details. You can add several remote identities to this list.
- 3 From the **Local Identity Type** drop-down list, select the type of identity to use when identifying the firewall to the remote peer. Only those identities defined within the selected certificate will be available in this field. Valid options are:
 - Distinguished Name
 - E-Mail
 - Fully Qualified Domain Name
 - IP Address

The values used as the selected remote certificate appear in the **Value** field. This value is filled in automatically using the information from the selected certificate. The field cannot be modified.

About the VPN Properties: Crypto tab (Manual key exchange)

Use this tab to configure manual authentication for this VPN definition. The firewall generates key values that are shared with the remote peer. The firewall and remote peer must enter the key values exactly to authenticate to each other.

Follow the steps below.

- 1 In the **IPSEC Transformations** drop-down list, select the appropriate form of IPsec transformation. The valid options are:
 - **Authentication Header (AH)** – Provides authentication only
 - **Encapsulating Security Payload (ESP)** – Provides encryption only
 - **Separate AH + ESP** – Performs separate transformations for authentication and encryption
 - **Combined ESP + AH** – Performs a single transformation that provides authentication and encryption
- 2 In the **Encryption** drop-down list, select the type of encryption you and the remote peer have chosen to use. The choices are:

Table 67 Encryption options

Encryption type	Key length
AES256	256-bit
AES128	128-bit
CAST128	128-bit
3DES	168-bit
DES	56-bit
Null	0

- 3** In the **Authentication Hash** drop-down list, select the type of authentication you and the remote peer have chosen to use. The valid options are:

- **sha1**
- **md5**

- 4** Click **Generate Keys** to create keys and SPI index values. Randomly generated keys appear in the key and SPI fields.

- The key and SPI fields available are dependent on the IPSEC Transformations selection.
- You can type your own unique key and SPI index, but it is not recommended. Since manually generating random keys is difficult, the firewall provides randomly generated authentication and encryption keys and Security Parameters Index (SPI) value for you and the remote peer to use. It is highly recommended that you use the default keys provided.

Note: Once you have chosen the keys, they must be kept a secret. You should only exchange the keys by a secure method, such as diskette, encrypted e-mail (such as PGP), or via the telephone. If attackers learn the key, they can decrypt all of your VPN traffic.

- 5** Send the generated keys and SPI values to the remote peer via a secure method (diskette, encrypted e-mail, or telephone).

The remote peer must enter the inbound and outbound keys and SPIs in the opposite fields:

Table 68 Key relationship between the firewall and the remote peer

If the key on the Sidewinder is in this field:	That key is entered on the remote peer in this field:
AH Inbound Key and SPI	AH Outbound Key and SPI
AH Outbound Key and SPI	AH Inbound Key and SPI
ESP Inbound Key and SPI	ESP Outbound Key and SPI
ESP Outbound Key and SPI	ESP Inbound Key and SPI

About the VPN Properties: Crypto tab (Automatic key exchange)

Use this tab to define the cryptographic and hashing algorithms that the firewall uses in this VPN definition to secure the traffic to and from the remote peer.

Follow the steps below.

- 1 In the **IPSEC Encryption Algorithms** pane, select an algorithm from the list of encryption algorithms. You can select multiple algorithms.
 - The Null option contains an encryption header but does not specify an encryption algorithm. It is generally only used during testing.
 - To authenticate only, without performing encryption, clear all encryption algorithms.
- 2 In the **IPSEC Authentication Algorithms** pane, select an algorithm from the list of authentication algorithms. You can select multiple algorithms.

About the VPN Properties: Advanced tab

The Advanced tab defines more advanced points of a VPN definition.

- As a general rule only administrators that are highly schooled in the nuts and bolts of VPN should modify the information on this tab.
- The information on this tab is used only with automatic key exchange.

The Advanced tab contains the following fields and buttons:

Internet Key Exchange (IKE) data fields

- **IKE V1 Exchange** – [If IKEv1 is the selected mode] Select which mode is used for key exchange:
 - **Main** – Has three exchanges between the initiator and the receiver; slower but secure. May not be used with dynamic IP clients with password authentication.
 - **Aggressive** – Has fewer exchanges between the initiator and the receiver; faster than Main mode, less secure.
- **Hard Limits** – Indicates how often the system must negotiate for new ISAKMP keys and how much ISAKMP traffic this phase can protect. The defaults are 3600 seconds (1 hour) and 0 (no limit to the amount of traffic).
- **Soft Percentage** – Indicates how far in advance of the hard limit to begin negotiating for new keys. This makes sure you have some new keys on hand by the time the hard limit expires.
- **Encryption Algorithms** – Specifies the encryption algorithm to use during Phase 1.
- **Hash Algorithms** – Specifies the hash algorithm to use during Phase 1.
- **PRF Algorithms** – Specifies the PRF algorithm to use during Phase 1 (IKEv2 only).
- **Key Exchange Group** – Indicates the Diffie-Hellman group to use for the derivation of ISAKMP keys.
- **Force XAuth on Rekey** – Select this option to force XAuth to be performed each time the phase 1 session is started or renegotiated.
- **Relax Strict Identity Matching** – Select this option to relax the identity matching restrictions. If you are experiencing issues associated with identity processing with the remote VPN peer, selecting this option can improve interoperability, but decreases security.
- **Enable NAT Traversal** – Allows multiple VPN users behind a NAT device to access a VPN tunnel. A UDP header is added to IPsec traffic and port 500 is changed to port 4500 to allow traffic across a NAT device.
 - Both sides of the VPN tunnel must have NAT Traversal capability.
 - Intended only for Dynamic IP policies.
 - This function works only in Tunnel mode.
- **Enable Initial Contact** – Enable this option to send and receive initial contact notify messages when first connecting with a VPN peer. This causes the peer to reload any previous state and is useful for re-synchronizing state after a device reboot.
- **Encrypt Final Aggressive Mode Packet** – For aggressive mode IKEv1 exchanges, this option will cause the firewall to encrypt the final aggressive mode packet in the exchange. You may need to enable this option if you are experiencing interoperability issues with your VPN peer using aggressive mode.

- **Enable Dead Peer Detection** – [If IKEv2 is the selected mode] Enable this option to send and receive messages to a VPN peer at regular intervals to confirm that the peer is available. If a reply is not received for a period of time, the connection with the peer is ended and no traffic is sent to the peer. The VPN connection must be re-established to send traffic to the peer.

Rekey data fields

- **Hard Lifetimes** – Indicates how often the system must negotiate for new IPsec keys and how much traffic it can encrypt. The defaults are 700 seconds and 0 (meaning no traffic limit).
- **Soft Percentage** – Indicates how far in advance of the hard limit to begin negotiating for new keys. This makes sure you have some new keys on hand by the time the hard limit expires.
- **Forced Rekey** – Forces the connection to rekey when the limits are reached, even if no traffic has passed through the VPN since the last rekey.

Caution: Do not enable the **Forced Rekey** option if you have HA/LS configured and are using static IP addresses for your VPNs. Doing so will cause all firewalls in the cluster to attempt to instantiate the VPN at the same time, resulting in failure.

- **Enable Extended Sequence Numbers** – Selecting this option doubles the IPsec sequence number to a 64-bit number. This option is useful if you expect extremely heavy traffic, ensuring that you can pass traffic over a VPN without running out of sequence numbers.
- **PFS** (Perfect Forward Secrecy) – If this option is enabled, it ensures that the key material associated with each IPsec security association cannot be derived from the key material used to authenticate the remote peer during the ISAKMP negotiation.
- **Oakley Group** – Indicates the Diffie-Hellman group to use for the PFS derivation of IPsec keys. This option is available only if the PFS option is enabled.

Managing client address pools

This section provides information on the client address pool user interface.

Configuring a client address pool

To create or modify a Client Address Pool, select **Network > VPN Configuration > Client Address Pools**. The following window appears.

Figure 321 Client Address Pools

You can perform the following actions:

- **Create a new client address pool** – Click **New** in the **Pools** area and enter information in the New Pool window.
- **Delete a client address pool** – Select the pool in the Pools list and click **Delete**.
- **Configure a client address pool** – To configure the Client Address Pools tabs, see the following:
 - For information on configuring the Subnets tab, see [Configuring the Subnets tab](#).
 - For information on configuring the Servers tab, see [Configuring the DNS and/or WINS servers](#).
 - For information on configuring the Fixed IP Map tab, see [Configuring the fixed IP map](#).

Configuring the Subnets tab

Use the Subnets tab to define the virtual address subnet for this address pool. You can also specify any local networks that you want to be accessible to remote clients using this pool.

To configure the virtual subnet address, select **Network > VPN Configuration > Client Address Pools** and select the client address pool that you want to configure from the Pools list in the left pane. The current configurations of the selected pool appears in the Subnets tab.

Figure 322 Client Address Pools: Subnets tab

The screenshot shows the 'Subnets' tab of the 'Client Address Pools' configuration window. At the top, the 'Pool Name' is set to 'self_signed'. Below this, there are three tabs: 'Subnets', 'Servers', and 'Fixed IP Map'. The 'Subnets' tab is active, displaying two lists: 'Virtual Subnet List' and 'Local Subnet List'. The 'Virtual Subnet List' contains one entry, '10.9.0.0/16', and the 'Local Subnet List' contains one entry, '10.10.0.0/24'. Both lists have scroll bars at the bottom. Below each list are three buttons: 'New', 'Modify', and 'Delete'.

To configure the Subnets tab:

1. Configure the **Virtual Subnet List**. This list defines the virtual subnets that define the IP address ranges that are available within this pool.

Note: The virtual subnets should **not** match the internal network's subnet, as this configuration could cause internal routing and connectivity issues. Virtual addressing works only if the client address pool uses unassigned address space.

You can perform the following actions:

- **Create a new virtual subnet** – Click **New** to define a new entry in the Virtual Subnet List.
 - Type the IP address that defines the network portion of the IP addresses used in the client address pool.
 - In the netmask field, specify the number of bits to use in the network mask. The network mask specifies the significant portion of the IP address.

Note: The virtual subnet cannot already exist in the firewall's routing table.

- **Modify a virtual subnet** – Select an existing subnet entry in the Virtual Subnet List and click **Modify**, then make your changes.
- **Delete a virtual subnet** – Select an existing entry from the Virtual Subnet List and click **Delete**.

2 Configure the **Local Subnet List**. This list defines the local networks available to remote peers that establish a VPN connection using an address from the client address pool. You can perform the following actions:

- **Create a new local subnet** – Click **New** to define a new entry in the Local Subnet List.
 - Type the IP address that defines the network portion of the local network available to remote peers establishing a VPN connection.
 - In the netmask field, specify the number of bits to use in the network mask. The network mask specifies the significant portion of the IP address.
- **Modify a local subnet** – Select an existing subnet entry in the Local Subnet List and click **Modify**, then make your changes
- **Delete a local subnet** – Select an existing entry from the Local Subnet List and click **Delete**.

Configuring the DNS and/or WINS servers

Use the Servers tab to define the DNS server(s) and/or the WINS server(s) that will be made available to remote peers.

- These servers provide name and address resolution services for devices within the local network.
- The DNS servers you specify can reside on the firewall or be located on another machine in a local or remote network.
- WINS servers are never located on the firewall.

To configure the DNS and/or WINS servers, select **Network > VPN Configuration > Client Address Pools** and select the client address pool that you want to configure from the Pools list in the left pane. The current configurations of the selected pool appear in the Servers tab.

Figure 323 Client Address Pools: Servers tab

The screenshot shows a configuration window for a client address pool named 'self_signed'. At the top, there are three tabs: 'Subnets', 'Servers' (which is selected), and 'Fixed IP Map'. Below the tabs, there are two main sections: 'DNS Servers:' and 'NBNS / WINS Servers:'. The 'DNS Servers:' section contains a list box with the IP address '10.10.12.23' selected. The 'NBNS / WINS Servers:' section is currently empty. At the bottom of each section, there are three buttons: 'New', 'Modify', and 'Delete'.

To configure the Servers tab:

- 1 The DNS Servers box lists the DNS servers that will be made available to VPN clients that establish a connection using an address from the client address pool. You can perform the following actions:
 - **Create a new DNS server** – Click **New** and type the IP address that specifies the location of the DNS server.
 - **Modify a DNS server** – Select an existing DNS server and click **Modify**, then make your changes.
 - **Delete a DNS server** – Select an existing DNS server and click **Delete**.
- 2 The NBNS/WINS Servers box lists the NBNS and WINS servers that will be made available to VPN clients that establish a connection using an address from the client address pool. You can perform the following actions:
 - **Create a new NBNS/WINS server** – Click **New** and type the IP address that specifies the location of the NBNS/WINS server.
 - **Modify an NBNS/WINS server** – Select an existing NBNS/WINS server and click **Modify**, then make your changes.
 - **Delete an NBNS/WINS server** – Select an existing NBNS/WINS server and click **Delete**.

Configuring the fixed IP map

Use the Fixed IP Map tab to define fixed addresses for selected clients.

- It enables each of the specified clients to connect to the firewall using its own unique IP address.
- It effectively reserves a specific IP address for a specified client.
- The fixed addresses you specify must be within the range of available IP address as defined by the client address pools.

Caution: Do not use network or broadcast addresses when mapping IP addresses to client IDs. These addresses are reserved and are not considered valid values for client address mappings. For example, if your address range is 192.168.105.0/24, then 192.168.105.0 (the network address) and 192.168.105.255 (the broadcast address) should not be used in a fixed IP client mapping. The network address is that address whose unmasked portion is all 0s, and the broadcast address is that address whose unmasked portion is all 1s.

Figure 324 Client Address Pools: Fixed IP Map tab

Pool Name:

Subnets Servers Fixed IP Map

Fixed IP Client Address Mappings:

IP Address	Client Identification
------------	-----------------------

New Modify Delete

One of the benefits of assigning fixed IP addresses to selected clients is that it allows you to govern what each client can do. For example, you might restrict access to certain clients, and you might grant additional privileges to other clients. You do this by creating a network object for a selected IP address and then using the network object within a rule.

The Fixed IP Map tab contains a **Fixed IP Client Address Mappings** box that lists the current IP address/client mappings. Each unique IP address can appear in the table only once. Multiple identities representing a single client, however, can be mapped to one IP address.

To configure the fixed IP map, select **Network > VPN Configuration > Client Address Pools** and select the client address pool that you want to configure from the Pools list in the left pane. The current configurations of the selected pool appears in the Fixed IP Map tab.

To define a new fixed IP client mapping address:

- 1 Click **New** to define an associate IP address and client identification strings.
- 2 In the **IP Address** field, enter the fixed IP address that will be associated with this mapping. The IP address must be within the virtual subnet for this pool.
- 3 Configure the client identification strings for this entry. All entries listed in the **Client Identification Strings** box will be mapped to the associated IP address. Because a client can use one of several different IDs (a distinguished name, an e-mail address, etc.) when negotiating a session, you can map multiple IDs to one IP address. However, you cannot map two separate clients to the same address.

Defining all the possible IDs for a client means you will be ready regardless of which ID is presented during the negotiation. Note that if a user will be using extended authentication, their user name will override any other ID.

Click **New** to add a client identification string. Select a string and click **Modify** to make changes. Select a string and click **Delete** to delete it.

Note: Each client identification string must be entered separately.

- 4 When you have finished configuring the client identification strings, click **Add** to add the new pool entry to the list.

Managing the ISAKMP server

The ISAKMP server is used by the Sidewinder to generate and exchange keys for VPN sessions. The ISAKMP server properties includes audit, negotiation and connection, and extended authentication parameters.

To configure the ISAKMP server:

- 1 Select **Policy > Rule Elements > Services**.
- 2 Select **isakmp** and click **Modify**.
- 3 Click **Properties**. The following window appears.

Figure 325 The ISAKMP Properties window

To configure the ISAKMP properties:

- 1 In the **Audit Level** field, select the type of audit output for the ISAKMP server.

Use these levels in most situations:

- **Error** – Logs only major errors.
- **Normal** – (Default) This is the most common setting. It outputs major errors and informational messages.
- **Verbose** – Use this level when initially troubleshooting VPN connectivity problems. This audit output is useful for detecting configuration issues.
- **Debug** – Logs all errors and informational messages. Also logs debug information.

Note: Only use **Debug** and **Error** if you are an experienced administrator or under guidance from Forcepoint support. In particular, debug can overflow your audit logs if left on for an extended period of time.

- 2 In the Advanced ISAKMP Server Options area, click **Properties**. The Advanced ISAKMP Options window appears with the following options:

- Certificate negotiation – The default is to allow certificate negotiation. If you clear the **Allow certificate negotiation** check box, all certificates used to authenticate remote peers must either be in the local certificate database or be accessible via LDAP.
- Negotiation properties – Set how long (in seconds) the ISAKMP server will wait for a response to its request to a remote peer before resending the packet, and how many times it will attempt to resend a packet if no response is received.
- Control how many remote peers can establish a connection to the ISAKMP server at one time—the default is **unlimited**. However, if you have a large number of remote users whose sessions will immediately reconnect after reboot, you may experience connection establishment issues and should adjust this limit.

Click **OK** to close the Advanced ISAKMP Options window.

- 3 In the XAUTH (Extended Authentication) Configuration area, configure how the ISAKMP server interacts with extended authentication:

- a In the **Allowed XAUTH Methods** area, select the authentication method(s) you want to be made available for VPN definitions that use extended authentication. A check mark indicates an allowed authentication method.

To add an authentication method to the list, click **New** and select a method from the pop-up menu. Configure the authenticator in the New Authenticator pop-up window.

- See [Extended authentication for VPN](#) for a detailed description of extended authentication.
- See [Configuring an authenticator](#) for details about configuring an authenticator.

- b If two or more authentication methods are selected, specify a default method from the **Default XAUTH Method** drop-down list. If you do not specify a default method, the first method selected in the list is used.

- c** Click **Properties** to open the Advanced XAUTH Options window and configure the following:
- Limit number of active VPNs per user – The default is one active VPN per authenticated user. This limit should work for most security policies. However, if your policy allows multiple users to use the same user name, generally from different VPN clients, you may need to remove this limit. (We strongly advise against allowing more than one user per user name.)
 - Negotiation properties – Set how long (in seconds) the ISAKMP server will wait for a response to its request to an authenticator before resending the packet, and how many times it will attempt to resend a packet if no response is received.

Click **OK** to close the Advanced XAUTH Options window.

4 Click **OK** to return to the main Services window.

5 Save your changes.

The ISAKMP service is now configured.

Example VPN Scenarios

The following sections describe three typical VPN scenarios. Each scenario begins by describing a particular VPN requirement. It then explains how to implement the solution using the Admin Console. These scenarios assume that the proper rule(s) are defined to allow ISAKMP traffic on the proper burb(s). In the scenarios that follow it is assumed a rule has been defined that allows ISAKMP traffic on the Internet burb.


Note: The values used in the following scenarios are for demonstration purposes only.

Scenario 1: Firewall-to-firewall VPN using a shared password

The easiest type of VPN definition to configure is one that uses a shared password for authentication. A shared password is typically used to establish a VPN connection between two corporate offices that have static IP addresses. Such a situation occurs if you have a business partner that requires access to your network, or if you have one or more corporate divisions located in different cities.

The following figure provides the sample configuration information used in this scenario.

Figure 326 VPN between two corporate offices



Requirements

This VPN scenario requires the following:

- A VPN connection between two corporate offices
- Shared password authentication
- Static IP addresses for each peer in the VPN definition

Implementation

The following steps show the fields on the VPN menus that must be defined to create this VPN definition. The configuration steps are performed on the firewall named *fw.east.example.com*.

Select **Network > VPN Configuration > VPN Definitions**, and then click **New**.

1 On the **General** tab:

- **Name** – corporate_west
- **Enabled** – Yes
- **Mode** – Fixed IP
- **Client Address Pool** – <disabled>
- **IKE Version** – V1
- **Burb** – Trusted
- **Encapsulation** – Tunnel
- **Local IP** – localhost
- **Local Network / IP** – 250.1.1.0/24
- **Remote IP** – 100.1.1.1
- **Remote Network / IP** – 50.1.0.0/16

Note: When configuring the firewall named *fw.west.example.com*, the Local Network/IP and the Remote Network/IP values are reversed and the Remote IP value is 200.1.1.1.

2 On the **Remote Authentication** tab:

- **Remote Authentication Method** – password

- **Enter Remote Password** – *samplepassword*
- **Verify Remote Password** – *samplepassword*
- **Remote Identity** – Gateway IP Address (100.1.1.1)

3 On the **Local Authentication** tab:

- **Local Authentication Method** – password (not editable)
- **Enter Local Password** – *samplepassword* (not editable)
- **Verify Local Password** – *samplepassword* (not editable)
- **Local Identity Type** – IP Address
- **Value** – localhost

4 On the **Crypto** tab: Select the algorithms to match the other firewall.

5 On the **Advanced** tab: No changes needed.

6 Click **Add** to save the new VPN definition.

7 Save your changes.

Summary

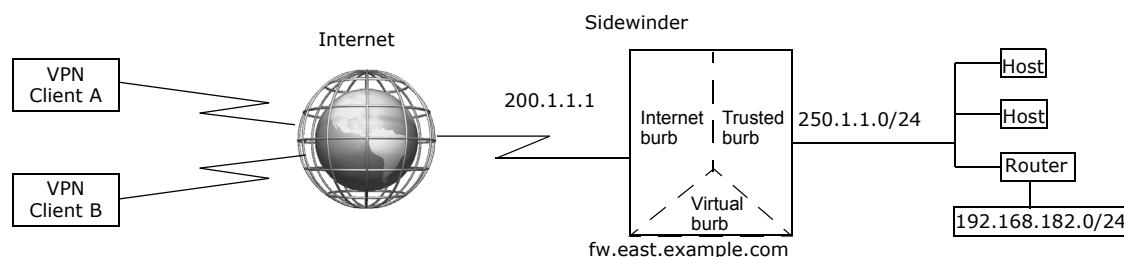
The VPN can be used as soon as you configure the other firewall. Enter the same type of information, changing the IP addresses as appropriate.

Scenario 2: Simple deployment of remote users

A common reason for using a VPN is to allow your travelling employees to connect to your corporate network from a remote site. This connection is typically made between an employee's laptop computer and your corporate Sidewinder. In this type of VPN definition, single (also known as "self-signed") certificates are generated by the firewall and distributed to each remote peer. This type of VPN can be used with dynamic IP-assigned clients and gateways. Create one definition for each client, so this type of VPN is typically used only if you have a small number of remote peers.

The figure below provides the sample configuration information used in this scenario. Note that the remote end of this VPN connection (from the firewall point of view) is a laptop that will be using a dynamic IP address.

Figure 327 One VPN definition per client



Assumptions

This VPN scenario assumes the following:

- A VPN connection between a remote computer and the firewall
- A self-signed firewall certificate that is generated by the firewall
- One or more remote certificates that is generated by the firewall and distributed to the remote peers
- One VPN definition per remote peer
- Each VPN definition is terminated in the Virtual burb
- VPN clients should have access to the 250.1.1.0 network but not the 192.168.182.0 network
- All clients make connections using a virtual IP address assigned from a client address pool
- All clients use VPN client software that supports mode-config

Note: When determining your deployment method, consider what steps you will take to ensure the protection of your private key material. Allowing unauthorized access to your private key material could compromise your entire network.

Implementation

The following steps show the fields on the VPN menus that must be defined in order to create this VPN definition.

1 Create and export a firewall certificate:

- In the Admin Console, select **Maintenance > Certificate/Key Management**.
- On the **Firewall Certificates** tab, click **New** and create a firewall certificate by specifying the following:
 - **Certificate Name** – MyFirewall_cert
 - **Distinguished Name** – CN=MyFirewall,O=bizco,C=US
 - **Submit to CA** – Self Signed
 - **Signature Type** – RSA
 - Click **Add**.
 - Click the **Save** icon.
- [Optional] On the **Firewall Certificates** tab, click **Export** and export the firewall certificate by specifying the following:
 - Select **Export Firewall Certificate to File**

- Click **Browse** and specify where you want to save the firewall certificate. The firewall certificate is often saved to an accessible location (portable storage device or protected network) for distribution to the client.
- Click **OK**.

2 Create a remote certificate for each client:

- a On the **Remote Certificates** tab click **New** and create a self-signed certificate for a client by specify the following:
 - **Certificate Name** – Sales_A
 - **Distinguished Name** – CN=Sales_A,O=bizco,C=US
 - **Submit to CA** – Self Signed
 - **Signature Type** – RSA
 - Click **Add**.
 - Click the **Save** icon.
- b Repeat [Step a](#) for each remote peer.
- c On the **Remote Certificates** tab, click **Export** and export the remote certificate by specifying the following:
 - Select **Export Certificate and Private Key**.
 - Select **Export Certificate/Key As One File**.
 - Export Client Private Key to File: Click **Browse** and specify where you want to save the private key.
 - Click **OK**.
- d Repeat [Step c](#) for each remote peer. When you are finished you should have the public firewall certificate as well as either the PKCS12-formatted object or the certificate/key file pair for that client saved to a location accessible by the remote peer (portable storage device or network).

3 Create a client address pool.

Using a client address pool lets you define which local networks the clients can access. For this example, assume you want to permit access to the 250.1.1.0 network but not the 192.168.182 network.

Note: Your client software must support this capability.

- a In the Admin Console, select **Network > VPN Configuration > Client Address Pools**, and then click **New** to create a new client address pool.
 - b **Enter New Pool Name** – SalesPool
 - c **Virtual Subnet** – 10.1.1.32/27
 - d Click **New**. In the **Local Subnet** field, enter 250.1.1.0/24 and then click **Add**.
 - e Click **Add** to add the new pool.
- Note:** The Subnet and Number of Bits in Netmask fields work in concert to determine the network portion of the addresses in the pool as well as the total number of addresses in the pool. The values shown here provide 30 possible addresses: 10.1.1.33 - 10.1.1.62. Modify these two values as appropriate for your situation. (For example, in this scenario you might alternatively specify IP Address = 10.1.1.16 and Netmask = 28, creating 14 possible addresses: 10.1.1.17 - 10.1.1.30.)
- f On the **Servers** tab: If the client software you are using supports this mode-config capability, specify your internal DNS and WINS servers here.
 - g Click **Add**.

4 Create a VPN definition for each client:

- a Select **Network > VPN Configuration > VPN Definitions**, and then click **New** to configure a new definition.
- b On the **General** tab:
 - **Name** – Sales_A

- **Enabled** – Yes
 - **Mode** – Dynamic IP Restricted Client
 - **Client Address Pool** – SalesPool
 - **Burb** – Virtual
 - **Encapsulation** – Tunnel
 - **Local IP** – localhost
- c** On the **Remote Authentication** tab:
- **Remote authentication method** – Single Certificate
 - **Remote Certificate** – Select the certificate you created in step 1C for this client
- d** On the **Local Authentication** tab:
- **Local authentication method** – Single Certificate (not editable)
 - **Firewall Certificate** – Select the certificate you created in step 1A
- e** On the **Crypto** tab: No changes needed.
- f** On the **Advanced** tab: [Conditional] If the clients are expected to be behind a NAT device, select **Enable NAT Traversal**.
- g** Click **Add** to save the new VPN definition.
- h** Click the **Save** icon to save your changes.
- 5** Repeat [Step 4](#) for each client, changing the name and the remote certificate as appropriate.

Summary

Each individual VPN connection can be used as soon as the remote peers are configured. Each client needs the client-specific certificate and private key information you saved in [Step 1](#) and [Step 2](#) to configure their end of the VPN connection. If you saved this information to removable media you can either hand it to them in person, mail it to them, or perform the imports while the machine is within a trusted network. It is not safe to distribute certificate and private key information via e-mail.


Note: The configuration described above restricts VPN traffic by terminating it in a virtual burb. Proxies and rule entries must be configured to specify what access the VPN clients have to the trusted network.

Scenario 3: Large scale deployment of clients

This scenario is similar to Scenario 2 except that instead of a small number of remote peers it assumes you have hundreds or even thousands of remote peers. Because it is unreasonable to create a unique VPN definition for each client, a Certificate Authority (CA) will be used. The CA, in conjunction with the remote identities you define, allows you to create one VPN that is accessible by all of the clients.

The following figure provides the sample configuration information used in this scenario.

Figure 328 One VPN definition for all clients



Assumptions

This VPN scenario assumes the following:

- A VPN connection between a Sidewinder and many clients
- A Certificate Authority-based VPN
- A single VPN definition for all clients with a like security policy rather than one definition per client
- The VPN connection is terminated in a virtual burb
- The clients can have dynamic or static IP addresses
- VPN clients should have access to the 250.1.1.0 network but not the 192.168.182.0 network
- All clients make connections using a virtual IP address assigned from a client address pool
- All clients are using VPN client software that supports mode-config

Note: It is assumed in this scenario that the clients do not have access to the CA and must rely on the firewall to create and distribute the necessary certificates and private keys.

Implementation

The following steps show the fields on the VPN menus that must be defined in order to create this VPN definition.

1 Define the CA used with this VPN:

- In the Admin Console, select **Maintenance > Certificate/Key Management**.
- Click the **Certificate Authorities** tab.
- Click **New** and create a CA by specifying the following:
 - **CA Name** – BizcoCA
 - **Type** – SCEP (or whatever value is appropriate)
 - **URL** – http://10.18.128.8
- Click **Add**.
- Click the **Save** icon to save your changes.
- Click **Get CA Cert** (Retrieves the CA Cert from the URL address.)
- Click **Get CRL** (Retrieves the Certificate Revocation List for this CA.)

2 Create a firewall certificate that is signed by the CA:

- a Click the **Firewall Certificates** tab.
- b Click **New** and create a firewall certificate by specifying the following:
 - **Certificate Name** – BizcoFW_by_CA
 - **Distinguished Name** – CN=BizcoFW_by_CA,O=Bizco,C=US
 - **Submit to CA** – BizcoCA
 - **Signature Type** – RSA
- c Click **Add**.
- d Click the **Save** icon to save your changes.

At this point the Status field for this certificate will be `PENDING`. This is because the request has been sent to the CA but the certificate has yet to be created. The status will remain `PENDING` until the CA administrator approves your request.

- e Click **Query**. This queries the CA to see if the certificate is approved. If yes, the Status field will change to `SIGNED` and the certificate is imported.

Note: The firewall automatically queries the CA every 15 minutes to see if the request has been accepted. If the request has been accepted, the firewall will retrieve the resulting certificate.

3 Create one or more identities that define who is authorized to use this VPN:

- a Click the **Remote Identities** tab.
- b Click **New** and create one or more identities that define who is authorized to use this VPN.
 - **Identity Name** – Sales_force
 - **Distinguished Name** – CN=*,OU=sales,O=bizco,C=us
- c Click **Add**.
- d Click **Close**.
- e Click the **Save** icon to save your changes.

4 Create a client address pool:

Using a client address pool lets you define which local networks the clients can access. For this example, assume you want to permit access to the 250.1.1.0 network but not the 192.168.182 network.

Note: Your client software must support this capability.

- a Select **Network > VPN Configuration > Client Address Pools**.
- b Click **New** and create a new client address pool by specifying the following:
 - **Enter New Pool Name** – SalesPool
 - **Virtual Subnet** – 10.1.1.0/24
 - Click **New**. In the **Local Subnet** field, enter 250.1.1.0/24 and then click **Add**.
- c Click **Add** to add the new pool.

Note: The **IP Address** and **Number of Bits in Netmask** fields work in concert to determine the network portion of the addresses in the pool as well as the total number of addresses in the pool. The values shown here provide 254 possible addresses: 10.1.1.0–10.1.1.255. Modify these two values as appropriate for your situation.

- d If the client software you are using has mode-config capability, specify your internal DNS and WINS servers on the **Servers** tab.
- e Click the **Save** icon to save your changes.

5 Create the VPN definition:

- a Select **Network > VPN Configuration > VPN Definitions**.
 - b Click **New** to configure a new definition.
 - c On the **General** tab:
 - **Name** – Large_scale_sales
 - **Enabled** – Yes
 - **Mode** – Dynamic IP Restricted Client
 - **Client Address Pool** – VPNPool
 - **Burb** – Virtual
 - **Encapsulation** – Tunnel
 - **Local IP** – localhost
 - d On the **Remote Authentication** tab:
 - **Authentication method** – Certificate + Certificate Authority
 - **Certificate Authorities** – BizcoCA (created in step 1A)
 - **Remote Identities** – Sales_force (created in step 1C)
 - e On the **Local Authentication** tab:
 - **Authentication method** – Certificate (not editable)
 - **Firewall Certificate** – BizcoFW_by_CA (created in step 1B)
 - f On the **Crypto** tab: Order the algorithms to match that of the client.
 - g On the **Advanced** tab: [Conditional] If the clients are expected to be behind a NAT device, select **Enable NAT Traversal**.
 - h Click **Add** to save the new VPN definition.
 - i Click the **Save** icon to save your changes.
- 6 Create the client certificates for each client:
- Note:** You can skip this step and [Step 7](#) for those clients that have online access to the CA. These clients can create and retrieve their own certificates.
- a In the Admin Console, **Maintenance > Certificate/Key Management**.
 - b Click the **Remote Certificates** tab.
 - c Click **New** and create a certificate for a client by specifying the following:
 - **Certificate Name** – Sales_A
 - **Distinguished Name** – CN=Sales_A,OU=sales,O=bizco,C=US
 - **Submit to CA** – BizcoCA
 - **Signature Type** – RSA
 - **Private Key** – Click **Browse** and specify where you want to save the private key associated with this certificate. In this scenario it is common to save the certificate to the same location as the exported firewall certificate.
 - **Certificate** – Click **Browse** and specify where you want to save this certificate. In this scenario it is common to save the certificate to the same location as the private key and the exported firewall certificate.
 - d Click **Add**.
 - e Click the **Save** icon to save your changes.
- 7 Provide certificate information and/or files to clients as necessary

- a Select **Maintenance > Certificate/Key Management**. Export the CA certificate and the public firewall certificate to the same location used in [Step 6](#).
 - b On the **Certificate Authorities** tab, select the CA certificate you created in [Step 1](#), then click **Export** and export the certificate by specifying the following:
 - **Destination** – File
 - **Generated CA Certificate File** – Click **Browse** and specify where you want to save the CA certificate. Add the *.pem* extension to the file name.
 - Click **OK**.
 - c [Optional] On the **Firewall Certificates** tab, select the firewall certificate you created in [Step 2](#), then click **Export** and export the certificate by specifying the following:
 - **Destination** – File
 - **Export Firewall Certificate to File** – Click **Browse** and specify where you want to save the firewall certificate. Add the *.pem* extension to the file name.
 - Click **OK**.
- 8 Repeat [Step 6](#) and [Step 7](#) for each remote peer.

When you are finished your storage location should have four items for each remote peer: the CA certificate, the firewall certificate, the unique private key for the client, and the remote certificate public key for the client.

Summary

The firewall is ready to accept connections across this VPN as soon as the remote peers are configured. To configure their end of the VPN connection, each client needs the client-specific certificate and private key information you saved in [Step 6](#), the CA certificate you created in [Step 1](#), and the firewall certificate you created in [Step 2](#). If you saved this information to removable media you can either distribute the information in person or mail it to them, or perform the imports while the machine is within a trusted network. It is not safe to distribute certificate and private key information via e-mail.

Note: The configuration described above restricts VPN traffic by terminating the VPN connection in a virtual burb. Proxies and rules must be configured to specify what access the VPN clients have to the trusted network.

SECTION 5

Maintenance

Chapter 21, General Maintenance Tasks

Chapter 22, Certificate/Key Management

Chapter 23, High Availability

21 General Maintenance Tasks

Contents

Setting the system date and time

Configuration file backup and restore

Activating the Sidewinder license

Protected host licensing and the Host Enrollment List

Software management

Editing files

Registering with Forcepoint Sidewinder Control Center

Enforcing FIPS

Enabling hardware acceleration

Configuring UPS

Setting the system date and time

You can use the Admin Console to set the date and time for the Forcepoint Sidewinder, and to configure the NTP (Network Time Protocol) service to synchronize clocks.

Setting the date and time

To set the system date and time, select **Maintenance > Date and Time**. The Date and Time window appears.

Note: Applying changes to the date and time will cause the firewall to automatically reboot. Therefore, you should only modify date and/or time settings during off-hours.

Figure 329 Date and Time window

The screenshot shows the 'Date and Time' configuration window. It is organized into three main panels. The top-left panel, titled 'Time Zone', contains a checkbox for 'Use GMT (UTC)' and three dropdown menus for 'Region' (set to 'America -- North and South'), 'Country' (set to 'United States'), and 'Zone' (set to 'Eastern Time'). The top-right panel, titled 'Date and Time', features a 'Date' dropdown menu showing '10/ 8/2008' and a 'Time' spinner control set to '05:39 PM'. The bottom panel, titled 'Time Synchronization', includes a section 'Enable Network Time Protocol (NTP) on:' with checkboxes for 'external' and 'internal'. To the right of this is an 'NTP Servers' table with columns 'Server', 'Burb', and 'Preferred'. Below the table are 'New', 'Modify', and 'Delete' buttons.

To change the date and time:

- 1 Set the time zone: From the **Time Zone** drop-down lists, select the region, country, and time zone where the Sidewinder is located.
To use Greenwich Mean Time (also known as Coordinated Universal Time), select **Use GMT (UTC)**.
- 2 Set the date and time: From the **Date** and **Time** fields, select the current date and time.
- 3 Save your changes.

Configuring NTP on a Sidewinder

For a detailed explanation of NTP and the Sidewinder, see [Understanding Network Time Protocol](#).

Tip: For best results, set the Sidewinder time as close as possible to the NTP server's time. If the times are too far apart, it might take a long time for the NTP server to synchronize with the firewall.

To configure NTP:

1 In the **Enable Network Time Protocol (NTP) on** list, select the firewall burb that receives time updates from the NTP server. This enables the NTP service in the appropriate burb.

2 In the **NTP Servers** area, add one or more NTP servers:

To configure an NTP server that supplies time to the Sidewinder, click **New**, then configure the server in the New NTP Server window.

a In the **Server** field, enter the IP address or the host name of the NTP server.

b From the **Burb** drop-down list, select the Sidewinder burb that communicates with the NTP server. This enables the NTP service in the appropriate burb.

c [Conditional] If you want this to be the preferred NTP server, select **Preferred Server**.

d Click **Add**.

Tip: To modify an existing NTP server, select the server in the list and click **Modify**, then make your changes in the Modify NTP Server window. To delete an existing NTP server, select the server in the list and click **Delete**.

3 Save your changes.

Understanding Network Time Protocol

NTP provides a way to synchronize all clocks on a network, or to synchronize the clocks on one network with those on another network.

When you use NTP with a Sidewinder:

- Time is set accurately.
- Synchronized network systems are useful for audit logs.
- You have a more accurate external time source when synchronizing your network for time-critical services.
- High Availability clusters benefit from synchronized time.


This release of the Sidewinder is compatible with NTP versions 2, 3, and 4. Version 4 is the preferred version and is the default on the firewall.

NTP servers and clients

In NTP, a server is a system that sends a time feed to another system. (The server is also referred to as a host.) The receiving system—the one whose time is being set by the server—is an NTP client. The Sidewinder can be set up as an NTP server or a client.

Figure 330 shows a simple configuration with an NTP time server and two NTP clients (A and B) in the same network. The NTP server supplies the time to NTP clients A and B. Using their own NTP software, each client system must also be set up to receive time from the server.


Figure 330 NTP server-client relationship



The Sidewinder as an NTP client

Figure 331 shows a common NTP setup. It is the recommended configuration, with the Sidewinder configured as a client receiving time from a server labeled "Internal time source." In this configuration, a server in the internal network (shown with an analog clock) is the designated time-setter for the rest of the network. The three other systems in the internal network are also NTP clients.

Figure 331 Sidewinder as an NTP client: internal server provides time to the firewall and to other internal workstations (no time feed to or from Internet)



By means of NTP, the server automatically maintains the correct time on the firewall and also maintains the time on other workstations in the network.

The internal network does not rely on an external time server, so it is not exposed to any security breaches that could result. Since the Sidewinder is not supplying time for other systems but is only receiving it, this setup has minimal effect on firewall performance.

- Do not configure the firewall to receive time from both an internal and external NTP server. It should receive time on only one burb. Input from the external time server cannot be reconciled with that from the internal server.
- For best results, set the Sidewinder time as close as possible to the NTP server's time. If the times are too far apart, it might take a long time for the NTP server to synchronize with the firewall.

To set up a Sidewinder as an NTP client for an internal NTP server:

- 1 Select **Maintenance > Date and Time**.
- 2 In the **Enable Network Time Protocol (NTP) on** list, select the firewall burb that is used to communicate with the NTP server. This enables the NTP service in the appropriate burb.
- 3 In the NTP Servers area, click **New**. The New NTP Server window appears.
- 4 Enter the IP address of the NTP server (1.1.1.1 in the example above).
- 5 From the burb drop-down list, select the firewall burb that is used to communicate with the NTP server.
- 6 [Conditional] Identify this server as the preferred NTP server.
- 7 Click **Add**.

The Sidewinder as an NTP server

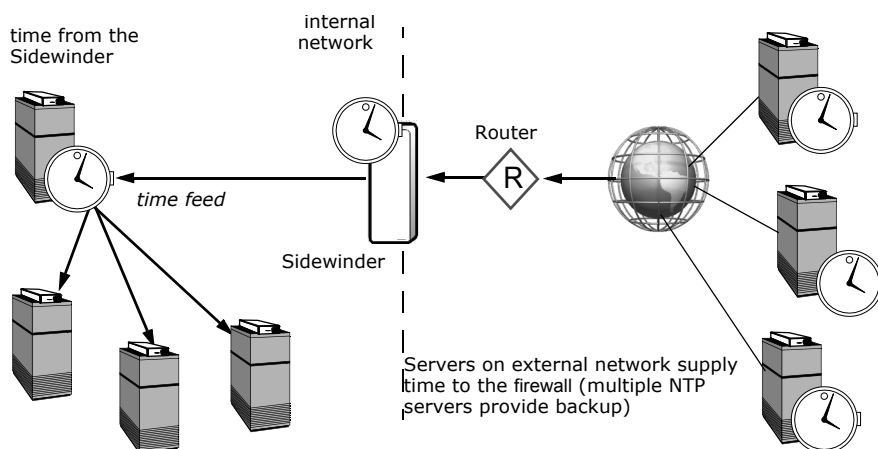
You can also set up the Sidewinder to be a time-setter for the rest of the network. The firewall can receive time from an external NTP server, then feed the time to an internal system, which in turn supplies time to your other workstations.

- Serving time directly from a Sidewinder to several clients can slow the performance of the firewall.
- If the firewall is serving time (host), its clients should not receive time from any other NTP server.

In the figure below, the firewall is receiving time from NTP servers on an external network and passing the time on to the internal network. This would be advantageous if your company required constant and precise time updates to within microseconds of world standard time. In this scenario, the router must be able to handle NTP traffic.

Note: An external-to-internal NTP configuration may introduce security concerns to the firewall and thus to your network. Therefore, this configuration is only recommended for sites that need world standard time.

Figure 332 The Sidewinder as an NTP server: external time servers supply time to the firewall, which passes time on to the internal system



Configure the firewall to receive time from an NTP server and serve time to clients:

- 1** Select **Maintenance > Date and Time**.
- 2** In the **Enable Network Time Protocol (NTP) on** list, enable the NTP service in the appropriate burbs.
 - One burb communicates with external NTP servers.
 - One burb feeds time to the internal time server.
- 3** Add one or more external NTP servers:
 - a** In the NTP Servers area, click **New**. The New NTP Server window appears.
 - b** Enter the IP address of the external NTP server.
 - c** From the burb drop-down list, select the external Sidewinder burb that communicates with the NTP server.
 - d** [Optional] Select **Preferred Server**.
 - e** Click **Add**.
- 4** Add the Sidewinder as a time server:
 - a** In the NTP Servers area, click **New**. The New NTP Server window appears.
 - b** Enter the following IP address: **127.127.1.0**
This tells the firewall to use the local hardware clock as its time source.
 - c** From the burb drop-down list, select the Sidewinder burb that communicates with the internal time server.
 - d** [Optional] Select **Preferred Server**.
 - e** Click **Add**.
- 5** Save your changes.

References

NTP is a complicated protocol with many options. There are numerous places where more information can be obtained. These include RFCs, web sites, and local manual (man) pages. For more information about NTP, see the following sources:

Internet Request For Comments (RFC)

The following RFCs provide information on NTP:

- RFC 1119 Network Time Protocol (Version 2)
- RFC 1305 Network Time Protocol (Version 3)
- RFC 4330 Simple Network Time Protocol (Version 4)

Web sites

Point your browser to the following web site:

<http://www.ntp.org/>

On-line manual (man) pages

Type the following commands:

```
man cf_ntp
man ntpd
man ntpdc
man ntpdate
```


Configuration file backup and restore

Use the Configuration Backup feature to back up and restore Sidewinder configuration files. Backing up the configuration files lets you quickly restore a firewall to a previous operational state.

The figure below shows the three options for a configuration backup.

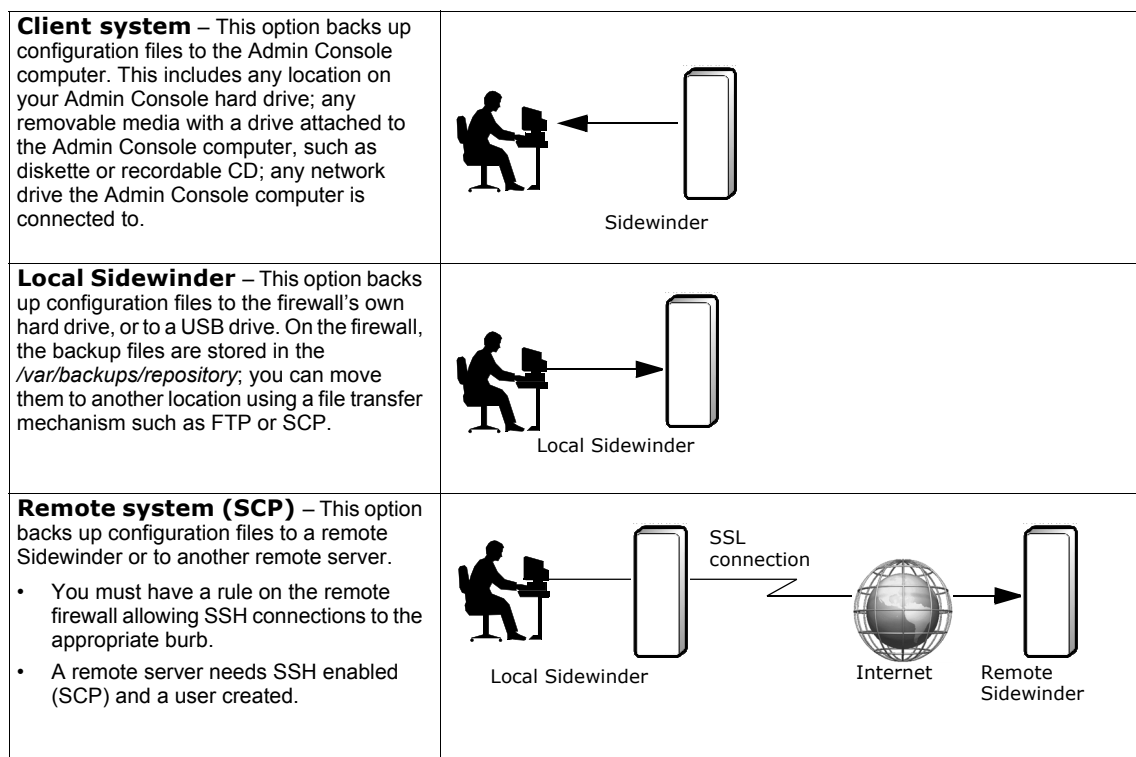


Figure 333 Configuration file backup options

You can also use this feature to manage configuration backups and make a disaster recovery backup.

To back up or restore your configuration files, select **Maintenance > Configuration Backup**. The Configuration Backup window appears.

- To create and manage configuration backups, click the **Configuration Backup** tab. See [About the Configuration Backup: Configuration Backup tab](#) for details.
- To restore configuration backups, click the **Configuration Restore** tab. See [About the Configuration Backup: Configuration Restore tab](#) for details.
- To schedule automatic configuration backups, click the **Schedule** tab. See [About the Configuration Backup: Schedule tab](#) for details.

About the Configuration Backup: Configuration Backup tab

Use this tab to back up Sidewinder configuration files. You can also create a disaster recovery backup and manage local configuration backups.

Figure 334 Configuration Backup tab

Configuration Backup | Configuration Restore | Schedule

Backup McAfee Firewall Enterprise Configuration

☒ Client system

☐ Local McAfee Firewall Enterprise

☐ Remote system (SCP)

Username:

Password:

Hostname: Port:

Directory:

Disaster Recovery Backup

You must connect a USB flash drive to the appliance before initiating this backup. A disaster recovery backup includes the current configuration and installed patches. This backup can only be restored during the re-imaging process.

Manage McAfee Firewall Enterprise Configuration Backups

Current local configuration backups:

Name	Version	Date	Location	Type	Ticket	Sun
initial_configuration	7.0.1.02	Mar 10 2009 03:1...	Local Disk	Complete		Host: fw11v191.ex

You can perform the following actions:

- [Back up configuration files](#)
- [Create a disaster recovery backup](#)
- [Manage configuration backup files](#)

Back up configuration files

Use this feature to back up Sidewinder configuration files.

- You can back up configuration files to the Admin Console computer, the Sidewinder, or a remote system.
- Only configuration files are backed up with this process. For example, the mail queues, the audit trail, the log files, or executable files will not be backed up in a configuration backup.

To back up more information, see [Create a disaster recovery backup](#).

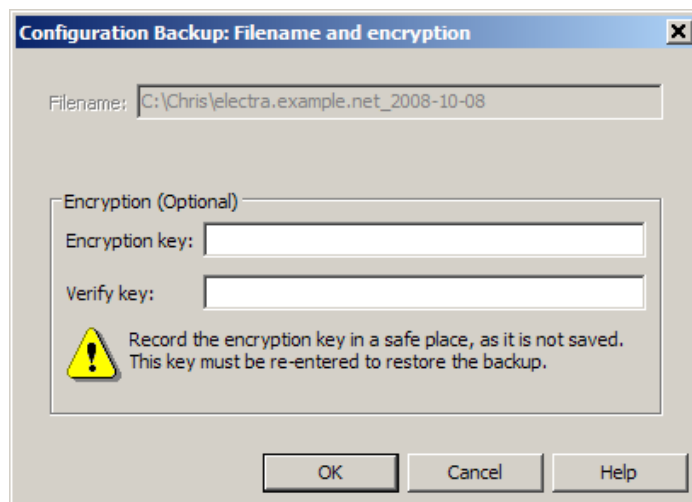
- The backup file must be at the same version as the system it is being restored to.

Backing up configuration files to the Admin Console computer

To back up configuration files to the Admin Console computer:

- 1 In the Backup Sidewinder Configuration area, select **Client System**.
- 2 Click **Backup now**. The Save Configuration Backup window appears.
- 3 Navigate to the location on the Admin Console computer where you want to save the configuration files. You can select any directory, media drive, or network available to the Admin Console computer.
- 4 [Optional] In the **File name** field, enter a name that can easily identify this configuration backup.
A default name consisting of the firewall name plus the current date automatically populates this field.
- 5 Click **Save**. The Filename and encryption window appears.

Figure 335 Filename and encryption widow for Client System backup



- 6 [Optional] Enter a key to encrypt the configuration backup file. Valid values include alphanumeric characters, periods (.), dashes(-), underscores (_), and spaces ().
 - This key will not be saved. You must remember it. You will not be able to restore the configuration file without this key.
 - You do not have to enter an encryption key. If you click **OK** without entering an encryption key, the backup continues.Enter the key again to verify.

- 7 Click **OK**.

A "Configuration backup successful" message appears.

- 8 Click **OK**.

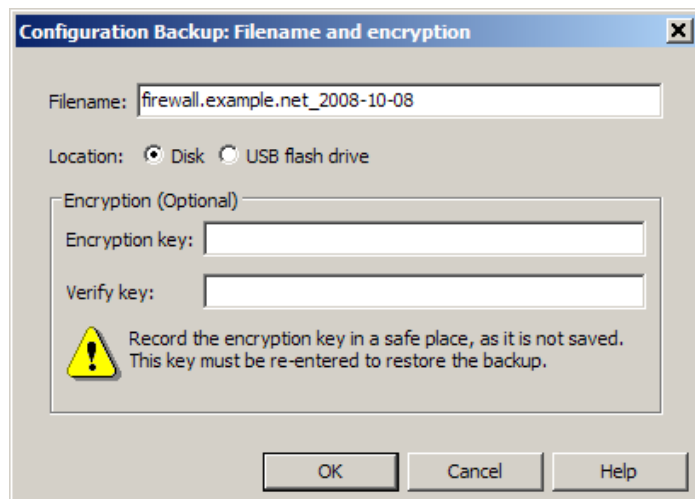
You have finished backing up a configuration file to the Admin Console computer.

Backing up configuration files to the Sidewinder

To back up configuration files to the firewall:

- 1 In the Backup Sidewinder Configuration area, select **Local Sidewinder**.
- 2 Click **Backup now**. The Filename and encryption window appears.

Figure 336 Filename and encryption window for Local Sidewinder backup



- 3 [Optional] In the **File name** field, enter a name that can easily identify this configuration backup.
A default name consisting of the firewall name plus the current date automatically populates this field.
- 4 Select a location for the backup file:
 - To save the backup file on the firewall, select **Disk**.
 - To save the backup file on a flash drive inserted in the USB port on the firewall, select **USB Flash Drive**.
 - Insert the flash drive before performing the backup.
 - Do not remove the flash drive from the firewall until the "Configuration backup successful" message appears
- 5 [Optional] Enter a key to encrypt the configuration backup file. Valid values include alphanumeric characters, periods (.), dashes(-), underscores (_), and spaces ().
 - This key will not be saved. You must remember it. You will not be able to restore the configuration file without this key.
 - You do not have to enter an encryption key. If you click **OK** without entering an encryption key, the backup continues.Enter the key again to verify.
- 6 Click **OK**. A "Configuration backup successful" message appears.
- 7 Click **OK**. The backup appears in the list of current local configuration backups.

You have finished backing up a configuration file to the firewall.

Backing up configuration files to a remote system

Before performing this procedure:

- If backing up to a remote Sidewinder, you must have a rule on the remote firewall allowing SSH connection to the appropriate burb.
- If backing up to another remote system, the remote system needs SSH enabled (SCP) and a user created.

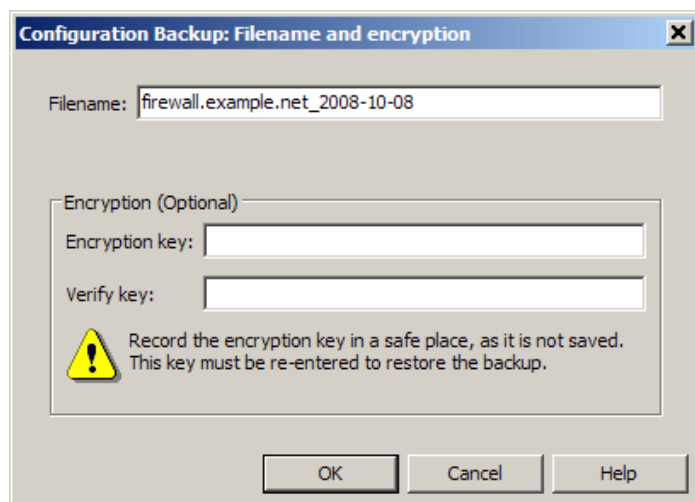
To back up configuration files to a remote system:

- 1 In the Backup Sidewinder Configuration area, select **Remote System (SCP)**.
- 2 Define the remote system that is receiving the configuration backup files:
 - In the **Username** field, enter the user name of a user on the remote system. If the remote system is a Sidewinder, this is a firewall administrator.
 - [Optional] In the **Password** field, enter the password used to authenticate the user to the remote system.
Note: The firewall does not save the password.
 - In the **Hostname** field, enter the host name or the IP address of the remote system.
 - The **Port** field default is 22.
 - In the **Directory** field, enter the directory on the remote system where the configuration files are stored. If the remote system is a Sidewinder, the administrator's home directory is the default.

This information is retained. You can change it at any time.

- 3 Click **Backup now**. The Filename and encryption window appears.

Figure 337 Filename and encryption window for Remote System backup



- 4 [Optional] In the **File name** field, enter a name that can easily identify this configuration backup.
A default name consisting of the firewall name plus the current date automatically populates this field.
- 5 [Optional] Enter a key to encrypt the configuration backup file. Valid values include alphanumeric characters, periods (.), dashes(-), underscores (_), and spaces ().
 - This key will not be saved. You must remember it. You will not be able to restore the configuration file without this key.
 - You do not have to enter an encryption key. If you click **OK** without entering an encryption key, the backup continues. Enter the key again to verify.
- 6 Click **OK**. A "Configuration backup successful" message appears.
- 7 Click **OK**.

You have finished backing up a configuration file to a remote system.

Create a disaster recovery backup

Use this feature to create a disaster recovery backup. If you need to re-install your firewall, you can use this backup to restore the configuration.

- The disaster recovery backup saves configuration files, installed packages, and the Quick Start Wizard data file (qsw_datafile) to a USB drive plugged into the appliance.
- The backup can take several minutes to complete.

To create a disaster recovery backup file:

- 1** Insert a flash drive into the USB port of the Sidewinder appliance.
- 2** On the Admin Console, select **Maintenance > Configuration Backup**.
- 3** Click **Create Disaster Recovery Backup**. The Configuration Backup and Restore: Disaster Recovery window appears.
- 4** [Optional] Enter a key to encrypt the disaster recovery backup. Valid values include alphanumeric characters, periods (.), dashes (-), and underscores (_).
 - This key will not be saved. You must remember it. You will not be able to restore the disaster recovery backup without this key.
 - You do not have to enter an encryption key. If you click **OK** without entering an encryption key, the backup continues.

Enter the key again to verify.

- 5** Click **OK**. A warning message appears.
- 6** Click **Yes** to confirm the backup.

A progress bar appears while the files are backed up to the flash drive.

Note: Do not remove the USB flash drive from the firewall until the “Disaster recovery successful” message appears.

When the backup is complete, a “successful” message appears.

- 7** Click **OK**.
 - 8** Remove the flash drive from the appliance and store it in a safe place.
- You can restore the disaster recovery backup only during re-imaging. For more information, see [Disaster recovery](#).

Manage configuration backup files

Use the Manage Sidewinder Configuration Backups area to move configuration backups between the Sidewinder and the Admin Console computer, and to delete configuration backups. You can also compare changes made between configuration backups and view audits associated with a change ticket.

The table lists the following information about each configuration backup:

- Name, Sidewinder software version, and date the configuration file was backed up
- Where the file is stored
- Whether the configuration backup is Lite or Complete: a Lite configuration backup does not contain debugging information that Technical Support uses for troubleshooting, and it does not contain home directories
- Whether there is a change ticket associated with the configuration backup

Note: The Current local configuration backups list is also on the Configuration Restore tab. Any changes to this list on the Configuration Backup tab also appear on the Configuration Restore tab.

Delete a configuration backup from the firewall

- 1 Select a configuration backup from the list.
- 2 Click **Delete**. Click **Yes** to confirm the deletion.

Move a configuration backup file from the Admin Console to the firewall

- 1 Click **Upload**. The Upload Configuration Backup window appears.
- 2 Navigate to the location on the Admin Console computer where the configuration backup file is stored.
- 3 Select the configuration backup file and click **Open**. An “Upload successful” message appears when the upload is complete.
- 4 Click **OK**. The backup file appears in the list.

Move a configuration backup file from the firewall to the Admin Console

- 1 Select a configuration backup file from the list.
- 2 Click **Download**. The Save Configuration Backup window appears.
- 3 Navigate to the location on the Admin Console computer you want to copy the backup file to. This can include a directory on the hard drive, removable media with a drive attached to the Admin Console computer, or a network drive the Admin Console computer is connected to.
- 4 Click **Save**. A “Download successful” message appears when the download is complete.
- 5 Click **OK**.

Compare configuration backups

- 1 Select configuration backups to compare:
 - To compare the changes that have been made between a configuration backup and the current running configuration, select the backup file from the list.
 - To compare the changes that have been made between two configuration backups, select the two backup files from the list by pressing and holding the **Ctrl** key and selecting the files.
- 2 Click **Compare**. The Compare Backups window appears.
- 3 Review the changes.
- 4 Click **Close**.

View change ticket audits

- 1 Select a Lite configuration backup file from the list.
- 2 Click **Audit**. The View Audit Data window appears.
- 3 When you are done viewing the data, click **Close**.

About the Configuration Backup: Configuration Restore tab

Use this tab to restore a Sidewinder to a previous operational state. You can also manage local configuration backups.

Figure 338 Configuration Restore tab

Configuration Backup

Configuration Restore

Schedule

☒ Client system

Filename:

Browse

☐ Local McAfee Firewall Enterprise

Current local configuration backups:

Name	Version	Date	Location	Type	Ticket	Summary
initial_configuration	7.0.1.02	Mar 10 2009 03:1...	Local Disk	Complete		Host: fw11v191.examp

Delete

Upload

Download

Compare

Audit

Refresh

☐ Remote system (SCP)

Filename:

Username:

Password:

Hostname: Port:

Directory:

Restore Now

You can perform the following actions:

- [Restore configuration backups](#)
- [Manage configuration backup files](#)

Restore configuration backups

Use this feature to restore a Sidewinder to a previous operational state.

- You can restore configuration files from the Admin Console computer, the Sidewinder, or a remote system.
- Only configuration files are restored with this process. For example, the mail queues, the audit trail, the log files, or executable files will not be restored from a configuration backup.
- The backup file must be at the same version as the system it is being restored to.

Restoring configuration files from the Admin Console

Note: The firewall will reboot after the configuration files have been restored.

- 1 Select **Client System**.
- 2 Click **Browse...** and navigate to the location where the backup file is stored. Select the backup file and click **Open**. The backup file appears in the **Filename** field.

You can also type the path and file name in the **Filename** field.
- 3 Click **Restore now**. A message appears stating that a system restore will cause a reboot of the firewall.
- 4 Click **Yes**. The Filename and encryption window appears.
- 5 [Optional] Enter an encryption key to restore the configuration backup file. Valid values include alphanumeric characters, periods (.), dashes(-), underscores (_), and spaces ().

Enter the key again to verify.

Note: If you did not enter an encryption key during the configuration backup, click **OK** to continue with the restore.

- 6 Click **OK**. When the restore is complete, a "Configuration restore successful" message appears.
- 7 Click **OK**. Your Admin Console is disconnected while the firewall reboots.

You have finished restoring a configuration backup from the Admin Console.

Restoring configuration files from the Sidewinder

Note: The firewall will reboot after the configuration files have been restored.

- 1 Select **Local Sidewinder**.
- 2 From the list of configuration backups, select the configuration you want to restore to the firewall.

If the configuration backup is on a flash drive, insert the flash drive in the firewall's USB port, then click the **Refresh** button to see the configuration backup in the list.
- 3 Click **Restore now**. A message appears stating that a system restore will cause a reboot of the firewall.
- 4 Click **Yes**. The Filename and encryption window appears.
- 5 [Optional] Enter an encryption key to restore the configuration backup file. Valid values include alphanumeric characters, periods (.), dashes(-), underscores (_), and spaces ().

Enter the key again to verify.

Note: If you did not enter an encryption key during the configuration backup, click **OK** to continue with the restore.

- 6 Click **OK**. When the restore is complete, a "Configuration restore successful" message appears.
- 7 Click **OK**. Your Admin Console is disconnected while the firewall reboots.

You have finished restoring a configuration backup from the Sidewinder.

Restoring configuration files from a remote system

Note: The firewall will reboot after the configuration files have been restored.

- 1 Select **Remote System (SCP)**.
- 2 In the **Filename** field, enter the name of the configuration backup file you are restoring.
- 3 Define the remote system where the configuration backup files are stored.
 - In the **Username** field, enter the user name of a user on the remote system. If the remote system is a Sidewinder, this is a firewall administrator.
 - [Optional] In the **Password** field, enter the password used to authenticate the user to the remote system.
Note: The firewall does not save the password.
 - In the **Hostname** field, enter the host name or the IP address of the remote system.
 - The **Port** field default is 22.
 - In the **Directory** field, enter the directory on the remote system where the configuration files are stored. If the remote system is a Sidewinder, the administrator's home directory is the default.

This information is retained. You can change it at any time.

- 4 Click **Restore now**. A message appears stating that a system restore will cause a reboot of the firewall.
- 5 Click **Yes**. The Filename and encryption window appears.
- 6 [Optional] Enter the key you used to encrypt the configuration backup file. Valid values include alphanumeric characters, periods (.), dashes(-), underscores (_), and spaces ().

Enter the key again to verify.

Note: If you did not enter an encryption key during the configuration backup, click **OK** to continue with the restore.

- 7 Click **OK**. A warning message appears asking if you want to disconnect after starting the restore.
- 8 Click **Yes**. When the restore is complete, a "Configuration restore successful" message appears.
- 9 Click **OK**. Your Admin Console is disconnected while the firewall reboots.

You have finished restoring a configuration backup from a remote system.

Manage configuration backup files

Use the Current local configuration backups list to move configuration backups between the Sidewinder and the Admin Console computer, and to delete configuration backups. You can also compare changes made between configuration backups and view audits associated with a change ticket.

The table lists the following information about each configuration backup:

- Name, Sidewinder software version, and date the configuration file was backed up
- Where the file is stored
- Whether the configuration backup is Lite or Complete: a Lite configuration backup does not contain debugging information that Technical Support uses for troubleshooting, and it does not contain home directories
- Whether there is a change ticket associated with the configuration backup

Note: The Current local configuration backups list is also on the Configuration Backup tab. Any changes to this list on the Configuration Restore tab also appear on the Configuration Backup tab.

Delete a configuration backup from the firewall

- 1 Select **Local Sidewinder**.
- 2 Select a configuration backup from the list.
- 3 Click **Delete**. Click **Yes** to confirm the deletion.

Move a configuration backup file from the Admin Console to the firewall

- 1** Select **Local Sidewinder**.
- 2** Click **Upload**. The Upload Configuration Backup window appears.
- 3** Navigate to the location on the Admin Console computer where the configuration backup file is stored.
- 4** Select the configuration backup file and click **Open**. An "Upload successful" message appears when the upload is complete.
- 5** Click **OK**. The backup file appears in the list.

Move a configuration backup file from the firewall to the Admin Console

- 1** Select **Local Sidewinder**.
- 2** Select a configuration backup file from the list.
- 3** Click **Download**. The Save Configuration Backup window appears.
- 4** Navigate to the location on the Admin Console computer you want to copy the backup file to. This can include a directory on the hard drive, removable media with a drive attached to the Admin Console computer, or a network drive the Admin Console computer is connected to.
- 5** Click **Save**. A "Download successful" message appears when the download is complete.
- 6** Click **OK**.

Compare configuration backups

- 1** Select configuration backups to compare:
 - To compare the changes made between a configuration backup and the current running configuration, select the backup file from the list.
 - To see the changes that have been made between two configuration backups, select the two backup files from the list by pressing and holding the **Ctrl** key and selecting the files.
- 2** Click **Compare**. The Compare Backups window appears.
- 3** Review the changes.
- 4** Click **Close**.

View change ticket audits

- 1** Select a Lite configuration backup file from the list.
- 2** Click **Audit**. The View Audit Data window appears.
- 3** When you are done viewing the data, click **Close**.


About the Configuration Backup: Schedule tab

Use this tab to schedule automatic configuration backups. You can back up configuration files to the Sidewinder, a USB flash drive, or a remote system.

Figure 339 Schedule tab

The screenshot shows the 'Schedule' tab of the 'Configuration Backup' configuration window. At the top, there are three tabs: 'Configuration Backup', 'Configuration Restore', and 'Schedule'. Below the tabs, there is a checkbox labeled 'Enable scheduled configuration backups' which is checked. To the right of this checkbox is an information icon and a message: 'Configuration backups are currently scheduled daily at 3:00am.' Below this, there are two main sections: 'Backup destination' and 'Backup schedule'. The 'Backup destination' section has two radio buttons: 'Local McAfee Firewall Enterprise' (selected) and 'Remote system (SCP)'. Under 'Local McAfee Firewall Enterprise', there are options for 'Location' (Disk selected, USB flash drive unselected) and 'Maintain local configuration backups' (Keep all backups unselected, Keep the last 20 backups selected). The 'Remote system (SCP)' section has input fields for Username, Password, Hostname, Port (22), and Directory. The 'Backup schedule' section has a 'Frequency' dropdown set to 'Daily' and a time selector set to '03:00 AM'. There is also a 'Custom' checkbox which is unchecked, followed by five input fields for Minute (0), Hour (3), Day of month (*), Month (*), and Day of week (*). At the bottom, there are two text boxes: 'Command' with the value 'cf config backup location=local' and 'Description' with the value 'Run automated configuration backups'.

Configuration Backup | Configuration Restore | Schedule

☒ Enable scheduled configuration backups  Configuration backups are currently scheduled daily at 3:00am.

Backup destination

☒ Local McAfee Firewall Enterprise
Location: ☒ Disk ☐ USB flash drive
Maintain local configuration backups: ☐ Keep all backups ☒ Keep the last backups

☐ Remote system (SCP)
Username:
Password:
Hostname: Port:
Directory:

Backup schedule

Frequency:

☐ Custom

Minute Hour Day of month Month Day of week
Command:
Description:

To schedule automatic configuration backups:

1 Select **Enable scheduled configuration backups**. If this check box is cleared, scheduled configuration backups will not occur.

2 Select the backup destination.

- If you select **Local Sidewinder**:
 - To save the backup file on the firewall, select **Disk**.
 - To save the backup file on a flash drive inserted in the USB port on the firewall, select **USB Flash Drive**.
 - Select whether you want to keep all backups or configure a number of backups to keep.
- If you select **Remote System (SCP)**, define the remote system that is receiving the configuration backup files:
 - In the **Username** field, enter the user name of a user on the remote system. If the remote system is a Sidewinder, this is a firewall administrator.
 - In the **Password** field, enter the password used to authenticate the user to the remote system. (The firewall does not save the password.)
 - In the **Hostname** field, enter the host name or IP address of the remote system.
 - The **Port** field default is 22.
 - In the **Directory** field, enter the directory on the remote system where the configuration files are stored. If the remote system is a Sidewinder, the administrator's home directory is the default.

This information is retained. You can change it at any time.

3 Configure the backup schedule:

- **Frequency** – From the drop-down list, select the frequency for exporting the file (hourly, daily, or weekly).
 - If you selected Hourly, enter the number of minutes after the hour.
 - If you selected Daily, enter the time for export.
 - If you selected Weekly, enter the time and day. You can select multiple days.
 - [Conditional] To define a custom frequency for exporting files, select **Custom** and complete the fields. Refer to man 5 crontab for options.
- **Command** – Displays the backup that will be executed.
- **Description** – A default description populates this field. If desired, enter a descriptive name for the task.

Activating the Sidewinder license

The Sidewinder license is automatically activated after running the Quick Start Wizard.

- When you first connect to a Sidewinder using the Sidewinder v7 Admin Console, a window appears displaying a list of features that are currently licensed for that firewall.
- To view a list of features currently available for the firewall and the status of each feature on your particular firewall, select **Maintenance > License** and click the **Firewall** tab. The Current Features pane displays the feature and license information.

For more information about each License window, see:

- [About the License: Contact tab](#)
- [About the License: Company tab](#)
- [About the License: Firewall tab](#)
- [About the License: Enrollment List tab](#)

If your firewall did not get licensed during initial configuration, the Sidewinder will operate for seven days with a trial license. These features are licensed during the trial period:

- SecureOS
- Support
- VPN
- Failover
- Strong Cryptography

If you need to license or relicense your firewall, or if you need to license a feature after initial configuration, you can perform these procedures:

- [Licensing from a firewall connected to the internet](#)
- [Licensing from a firewall on an isolated network](#)

Note: If you relicense to a new firewall ID, you must call Forcepoint support.

Licensing from a firewall connected to the internet

If you are working on a Sidewinder that is connected to the internet, perform this procedure to provide the necessary company information and obtain an activation key.

To license a Sidewinder connected to the internet:

- 1** Locate the serial number for your firewall. The serial number should appear on your Activation Certificate.
- 2** In the Admin Console, select **Maintenance > License**. The License window appears.
- 3** Click the **Contact** tab and enter your company information.
See [About the License: Contact tab](#) for details.
- 4** Click the **Company** tab and enter your company information.
See [About the License: Company tab](#) for details.
- 5** Click the **Firewall** tab and enter the firewall information:
 - a** In the **Serial Number** field, type the 16-digit alpha-numeric serial number for this firewall. The serial number is located on your Sidewinder Activation Certificate.
 - b** In the **Firewall ID** field, accept the default. Do not change the Firewall ID unless instructed by Customer Service.
See [About the License: Firewall tab](#) for details.
- 6** Click **Submit Data**. The license information is sent to the licensing web site using an encrypted HTTPS session.
If the data is complete, the request is granted and a new activation key is written to the **Activation Key** field. The Current Features list updates with the new license information.

Licensing from a firewall on an isolated network

If you are on an isolated network and do not have access to the activation server, perform this procedure to request an activation key.

To license a firewall on an isolated network:

- 1 Select **Maintenance > License**. The License window appears.
- 2 Click the **Firewall** tab.
- 3 In the **Serial Number** field, verify that it shows the 16-digit serial number located on the Activation Certificate or on your hardware platform.
- 4 In the **Firewall ID** field, accept the default. Do not change the Firewall ID unless instructed by Customer Service.
- 5 [Conditional] If your Admin Console does not have web access, move to a computer that has web access. Bring a copy of the serial number and firewall ID with you.
- 6 Open a browser and go to the Sidewinder activation web page:
<http://sidewinder.activations.forcepoint.com/activation.cfm>
- 7 Complete the form on the web site and click **Submit**. A confirmation screen appears.
- 8 Verify that the information you entered is correct. If it is not correct, use the **Back** button to return to the form and correct the information.
- 9 Click **Submit**. After a minute or so, a new web page appears displaying the activation key.
- 10 Using the on-screen instructions, save the activation key to a diskette.

Tip: You may choose to continue following the on-screen instructions for importing the file via command line, or use the Admin Console instructions given here.

- 11 Insert the diskette into the *Admin Console's diskette drive*.
- 12 Select **Maintenance > License**. The License window appears.
- 13 Click the **Firewall** tab.
- 14 Click the **Import Key...** button. The Import Key window appears.
- 15 Complete the Import Key window:
 - a Select **Local File**.
 - b In the **File** field, enter the name of the file that contains the activation key. If necessary, click **Browse...** and navigate to the file.
- 16 Click **OK**. The activation key is extracted from the file and appears in the **Activation Key** field.

Your firewall software and any features you licensed are activated.
- 17 Enter information in the License windows to complete the licensing process:
 - a Click the **Contact** tab and enter your company information.

See [About the License: Contact tab](#) for details.
 - b Click the **Company** tab and enter your company information.

See [About the License: Company tab](#) for details.
 - c Save your changes.

Configuring the Firewall License tabs

To configure license information, select **Maintenance > License**. The License window appears. The License window contains four tabs used to collect various licensing information:

- **Contact** – See [About the License: Contact tab](#) for details.
- **Company** – See [About the License: Company tab](#) for details.
- **Firewall** – See [About the License: Firewall tab](#) for details.
- **Enrollment List** – See [About the License: Enrollment List tab](#) for details.

About the License: Contact tab

Use the Contact tab to enter contact information for the administrator of this Sidewinder. This information is needed so that you can receive important customer bulletins and renewable support licenses.

Figure 340 Firewall License: Contact tab

The screenshot shows the 'Contact' tab of the 'Firewall License' window. At the top, there are four tabs: 'Contact', 'Company', 'Firewall', and 'Enrollment List'. Below the tabs is a note: 'PLEASE NOTE: This information applies to the End User of the firewall.' The form contains several fields: 'First Name' (Bruce), 'Last Name' (Wayne), 'Email' (bruce@batman.com), 'Primary Phone' (222-222-2222), '(Alternate Phone)' (222-222-2223), '(Fax)' (empty), '(Job Title)' (Administrator), '(Purchased From)' (Wayne Enterprises), and '(Comments)' (a large empty text area with a scrollbar). The 'Purchased From' field has a dropdown arrow on the right. The 'Comments' field also has a scrollbar on the right.

To enter contact information:

- 1 In the **First Name** field, enter the first name of the Sidewinder administrator.
- 2 In the **Last Name** field, enter the last name of the Sidewinder administrator.
- 3 In the **Email** field, enter the e-mail address of the Sidewinder administrator.
- 4 In the **Primary Phone** field, enter the phone number of the Sidewinder administrator, including the area code.
- 5 [Optional] In the **Alternate Phone** field, enter an alternate phone number in case the first number is unavailable.
- 6 [Optional] In the **Fax** field, enter a fax number for your organization.
- 7 [Optional] In the **Job Title** field, enter the job title of the person responsible for administering this firewall.

- 8** [Optional] In the **Purchased From** field, enter the name of the company that sold you this firewall.
- 9** [Optional] In the **Comments** field, enter miscellaneous information about your site.
- 10** Save your changes.

About the License: Company tab

Use the Company tab to enter information about the company that has purchased this Sidewinder.

Figure 341 Firewall License: Company tab

To enter company information:

- 1 In the **Company Name** field, type the full name of the company that purchased this firewall.
- 2 In the **Industry Classification** drop-down list, select the classification that most closely matches your industry.
- 3 Fill in the requested address information fields on the Company Address tab and on the Billing Address tab. If the information is the same on both tabs, enter the information on the **Company Address** tab, then switch to the **Billing Address** tab and click **Copy From Company Address**.
- 4 Save your changes.
- 5 Click the **Firewall** tab to provide the information necessary to license your firewall. The Firewall tab appears.

About the License: Firewall tab

Use the Firewall tab to enter information about the Sidewinder you are attempting to license, and to view the features available on the firewall and their licensing status.

- Do not edit the activation URL unless instructed to do so by Forcepoint support. The **Activation URL** field displays the URL of the web site to which the Sidewinder licensing information is sent.
- The Firewall ID is automatically selected. Do not change the Firewall ID unless instructed by Customer Service.

Figure 342 Firewall License: Firewall tab

Please Note: Accurate information is required to ensure uninterrupted support and product updates.

Serial Number: Firewall Version:

Firewall ID:

Activation URL:

Activation Key:

Current Features:

Feature	License Status	Expiration
SecureOS	Licensed	Expires Tue Nov 15 09:04:42 2016
Support	Licensed	Expires Thu Aug 17 07:53:57 2017
VPN	Licensed	Expires Thu Aug 17 07:53:57 2017
Failover	Licensed	Expires Thu Aug 17 07:53:57 2017
Strong Cryptography	Licensed	Expires Thu Aug 17 07:53:57 2017
Anti-Virus	Licensed	Expires Thu Aug 17 07:53:57 2017
Anti-Spam	Licensed	Expires Thu Aug 17 07:53:57 2017
IPS	Licensed	Expires Thu Aug 17 07:53:57 2017
SSL Decryption	Licensed	Expires Thu Aug 17 07:53:57 2017
IPS Signature	Licensed	Expires Thu Aug 17 07:53:57 2017
Promotion	Licensed	Expires Fri Sep 16 07:53:57 2016
Protected Host Licenses	Licensed for 500 protected hosts.	

To enter Sidewinder licensing information:

- 1 In the **Serial Number** field, type the 16-digit alpha-numeric serial number for this firewall. Include the dashes in your serial number.

The serial number is located on your Sidewinder Activation Certificate.

- 2 Click **Submit Data...** to submit the license information to the licensing web site using an encrypted HTTPS session.

If the data is complete, the request is granted and a new activation key is written to the **Activation Key** field. The Current Features list updates with the new license information.

- 3 In the **Current Features** pane, view the features currently available for the firewall and the licensing status of each feature.
 - If a feature you want to use is currently not licensed, that feature's window title in the Admin Console will include the words *Not Licensed*.
 - If a feature you want to use is currently not licensed, you must obtain a different activation key in order to enable that feature.
- 4 [Optional] If you need to import an activation key that has been saved to a file, click **Import Key...**

You will typically use this button if your Sidewinder or local network does not have access to the URL defined in the **Activation URL** field. See [Licensing from a firewall on an isolated network](#) for instructions on obtaining a license in this situation.

- 5 Save your changes.

About the License: Enrollment List tab

Use the Enrollment List tab to view and modify the Host Enrollment List.

- You can identify which IP addresses are currently counted against your protected host license cap.
- You can delete IP address entries that you do not want counted against your host cap. For example, you might do this if a connection is initiated from a test system in your lab and you do not want that system to count against the host license cap.

Note: If you have an unlimited license, all license processing is bypassed.

To display and modify the contents of the Host Enrollment List, select **Maintenance > License** and click the **Enrollment List** tab.

You can perform the following actions:

- View the number of hosts authorized by your current Sidewinder license in the **Licensed host limit** field. This is your host license "cap."
- View the current number of hosts listed in the **Number of hosts in enrollment list** field. This number is important because if it exceeds the number of hosts authorized by the Sidewinder license, you will be considered to be in violation of your license cap. If you have an unrestricted host license, the term *Unlimited* will appear in this field.
- Delete hosts from the **Host Enrollment List** by selecting the host and clicking **Delete**. To select multiple hosts to delete, hold the **Shift** key while selecting the hosts.
- Refresh the window to reflect updated information by clicking **Refresh**.

See [Protected host licensing and the Host Enrollment List](#) for an in-depth explanation of the Host Enrollment List.

Protected host licensing and the Host Enrollment List

The Host Enrollment List is a dynamic list that is used to record each unique IP address (host) that makes an outbound connection to the Internet. The Sidewinder uses this list to verify compliance with the IP address license “cap”—the portion of your Sidewinder license that dictates the number of hosts the firewall will support.

Note: You may ignore this section if you have an unlimited license. All license processing is bypassed if you have an unlimited license.

The Sidewinder strictly enforces the maximum IP address (host) license number, meaning only the number of IP addresses authorized by the protected host license will be allowed to make connections through the firewall.

- [More than 100 users] If the number of IP addresses in the enrollment list exceeds 75% of the number allowed by your protected host license, an audit will occur informing you that you are approaching the maximum number of hosts. The audit will also display the current number of hosts and the maximum number of hosts that are allowed for your license.
- If the enrollment list becomes full:
 - Additional audits will occur each time a new IP address attempts to make a connection to the Internet, and only the IP addresses already contained in the enrollment list will be allowed to make a connection to the Internet.
 - A user attempting to make a connection using a browser will receive a standard policy denial message. A user attempting to make a connection using a non-browser application (for example, FTP), will simply be blocked and they will not receive an error message.

If you reach the host enrollment maximum and you want to allow access to additional hosts, you will need to do the following:

- Modify the host enrollment list to remove hosts entries that no longer need to be listed. See [Managing the Host Enrollment List](#).
- Upgrade your license, or upgrade to a larger Sidewinder.

How hosts are calculated

Any host that contains a unique IP address and that initiates a connection from a non-Internet burb is counted as a new host and added to the Host Enrollment List.

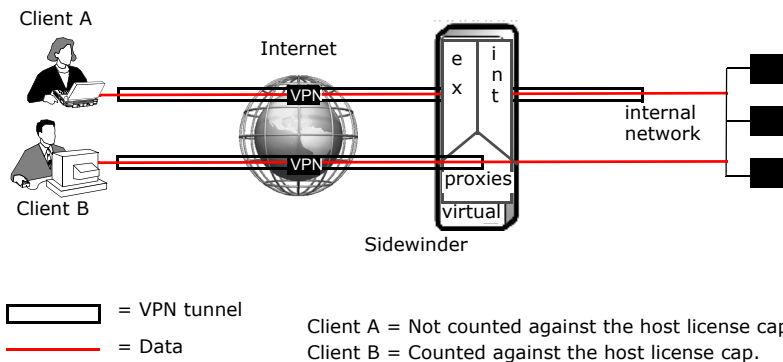
The manner in which remote hosts access the Sidewinder may affect the host count. For example:

- Remote hosts that use dynamic addressing rather than static addressing may have multiple IP addresses added to the Host Enrollment List.

Note: If you have control over the dynamic host server, you might want to control the range of IP addresses used for remote hosts, to minimize the impact of dynamic IP addresses on the enrollment list.

- Hosts accessing the firewall via a VPN will be added to the Host Enrollment List if the VPN uses proxies to move the traffic from a non-Internet burb to another burb. [Figure 343](#) illustrates this idea.

Figure 343 Determining which VPN clients count against the host license cap



The firewall counts total hosts connecting to the Internet, not the number of hosts connecting to the Internet at one time. It is important to understand the distinction. Assume the following:

- You have a 25-host license.
- You have 30 hosts.
- No more than 20 hosts are online at any one time.

You will still exceed the license cap because each host that goes online has their IP address added to the Host Enrollment List. Eventually a 25th host will be added to the Host Enrollment List, and the 26th host will be detected by the firewall, putting you over the limit.

Managing the Host Enrollment List

Use the Enrollment List tab to view and modify the Host Enrollment List.

- You can identify which IP addresses are currently counted against your protected host license cap.
- You can delete IP address entries that you do not want counted against your host cap. For example, you might do this if a connection is initiated from a test system in your lab and you do not want that system to count against the host license cap.

To display and modify the contents of the Host Enrollment List, select **Maintenance > License** and click the **Enrollment List** tab.

You can perform the following actions:

- View the number of hosts authorized by your current Sidewinder license in the **Licensed host limit** field. This is your host license “cap.”
- View the current number of hosts listed in the **Number of hosts in enrollment list** field. This number is important because if it exceeds the number of hosts authorized by the Sidewinder license, you will be considered to be in violation of your license cap. If you have an unrestricted host license, the term *Unlimited* will appear in this field.
- Delete hosts from the **Host Enrollment List** by highlighting the host and clicking **Delete**. To select multiple hosts to delete, hold the **Shift** key while selecting the hosts.

Note: You can update the contents of the Host Enrollment List field by clicking **Refresh**.

Consider the following information when deleting entries from the enrollment list:

- If the host you delete has a current connection through the firewall, that connection will be preserved.
- If the *host* severs the connection and attempts a new connection, the new connection request may or may not be approved.
- A new connection request will be permitted only if there is still room available within the enrollment list.

You can use the System Responses window to configure to have e-mails sent when the enrollment list reaches the maximum number allowed: Select **Monitor > System Responses**, then select **host license exceeded** and click **Modify**. See [About the Modify System Responses: Event tab](#) for more information.

Software management

The Sidewinder comes installed with the latest software available at the time. Use the Software Management window to keep your firewall current with updates. You can also uninstall updates or revert to a previous configuration.

Understanding software management

Updates can include improvements and enhancements to the Sidewinder software, as well as updates to optional features. The types of software updates are:

- **Base** – A major release with significant enhancements and added functionality.
- **Patch** – Released periodically and contains software fixes and/or new features. General release patches must be installed sequentially.
- **Vendor Patch** – Contains fixes, updates, or new features specific to the anti-virus add-on module. Only install the patch if you have the feature enabled. Patches are to be installed sequentially.
- **Upgrade** – Brings your firewall to a new base version and includes significant new features.
- **Hotfix** – Contains an issue-specific fix, and should be installed only if it addresses a current problem.
- **Restricted** – Distribution is restricted to identified customer(s) to solve a particular defect.
- **E-Patch** – Sent to a particular customer on an as-needed basis to determine if the fix corrects an identified defect for a particular version of the product. If the fix works, it is then converted to a hotfix or incorporated into a full product release.

Software management includes these actions:

- [Loading and installing](#)
- [Uninstalling and rolling back](#)
- [Re-installing and re-imaging](#)

Loading and installing

Periodic software updates, called *packages*, are available on our FTP site.

- **Load** – You load a package onto the firewall. This moves the package from the FTP site to the firewall, but does not install it.
 - Packages can be loaded manually, at automatic intervals, or at scheduled times using the Manage Packages and Download Packages tabs.
 - The FTP site is the default location for software packages. You might configure a different location to load packages from if you have firewalls on an isolated network, or to speed up downloads to several firewalls.
- **Install** – Packages that are loaded on the firewall can then be installed on the firewall using the Manage Packages tab.
 - Packages can be installed manually or automatically at a scheduled time. Packages can be installed individually or several at a time.
 - Key package information is provided, such as dependencies with other packages, whether a reboot is required, and whether the package can be uninstalled.

Uninstalling and rolling back

If you are not satisfied with an update, you can uninstall the package or revert to the previous configuration.

- **Uninstall** – Package uninstalls are performed on the Manage Packages tab.
 - The Uninstallable column states if a package is uninstallable.
 - Any configuration changes you make after the package was installed remain after the package is uninstalled.
 - Packages can be uninstalled manually or automatically at a scheduled time. Packages can be uninstalled individually or several at a time.
 - Uninstalled packages remain loaded on the firewall.
- **Rollback** – Use the Rollback tab to restore the firewall to a previous state. Rolling back is an option if a package is not uninstallable.
 - Any configuration changes made after the package was installed are lost. A rollback reverts the firewall to the state just before the package was installed. Therefore, rolling back is a recommended recovery option for only a short time after a package installation.
 - A rollback can be performed manually or automatically at a scheduled time.
 - A rollback always requires a reboot.
 - The tab displays what patch level the firewall will roll back to and the date and time the patch was loaded.

Re-installing and re-imaging

More serious issues might require you to re-install or re-image your Sidewinder.

- **Re-install from the virtual CD** – Use the virtual CD function to re-install the Sidewinder software. Re-installing from the virtual CD puts the firewall in its original unconfigured state.

For information and instructions, see [Re-installing your firewall from the virtual CD](#).

- **Re-install from a CD-ROM** – Re-installing from a CD-ROM erases the firewall's hard drive and re-installs the software. Re-installing from a CD-ROM is a more disruptive procedure. It should be done only when other methods cannot fix the issue and only if instructed by Forcepoint support.

For information and instructions, see [Re-installing your firewall from a CD-ROM](#).

- **Re-install from a USB drive** – Re-installing from a USB drive erases the firewall's hard drive and re-installs the software. Re-installing from a USB drive is a more disruptive procedure. It should be done only when other methods cannot fix the issue and only if instructed by Forcepoint support.

For information and instructions, see [Re-installing your firewall from a USB drive](#).


To manage Sidewinder software, select **Maintenance > Software Management**. The Software Management window appears. The Software Management window has three tabs:

- [About the Software Management: Manage Packages tab](#).
- [About the Software Management: Download Packages tab](#).
- [About the Software Management: Rollback tab](#).

About the Software Management: Manage Packages tab

To manage software packages for the Sidewinder, select **Maintenance > Software Management**. The following window appears:

Figure 344 Software Management: Manage Packages tab



Use this tab to load, install, and uninstall packages on the firewall.

You can perform these tasks:

- [Viewing available packages](#)
- [Sorting the Manage Packages table](#)
- [Loading, installing, and uninstalling packages now](#)
- [Scheduling automatic installs and uninstalls](#)
- [Viewing package information and activity logs](#)
- [Enabling e-mail notification](#)

Viewing available packages

To populate the table with packages that are available for downloading, click **Check for Updates**. Packages appear in the table with a status of *Available*. They are not yet installed.

Available packages are loaded from our FTP site, or from another site you designate on the Download Packages tab.


Use the Download Packages tab to configure automatic checking and loading, and to change the location where packages are downloaded from.

Sorting the Manage Packages table

To sort the table of packages:

- Select which package types and statuses appear in the table: Click **View Options** and make your selections in the pop-up window.
- Sort the Package Name column in ascending or descending order by clicking the column heading.
- Sort other columns by right-clicking a column heading and selecting a filter option from the pop-up list.
- Use the **Find** field to search for a specific element(s) in the list. Type your search criteria, and only packages with matching elements will appear in the list.

Figure 345 Manage Packages table find field




Loading, installing, and uninstalling packages now

Use the Manage Packages table to load, install, or uninstall packages immediately.

- The buttons that appear depend on the status of the package selected in the table:

Figure 346 Manage Packages table buttons



- You can select more than one package to manage.
 - The packages you select must have the same status.
 - Select several consecutive packages by selecting the first package, then pressing the **Shift** key while selecting the last package. To select several non-consecutive packages, press the **Ctrl** key as you select each desired package.

To download a package to the Sidewinder:

- 1 Select a package in the table with the status of *Available*.
- 2 Click **Download**. A “successfully loaded” message appears and the package status changes to *Loaded*.

To install a package:

- 1 Select a package in the table with the status of *Loaded*.
- 2 Click **Install**. The Install window appears.

If the package has dependencies, the dependencies appear in a lower table.

- If a dependency is loaded, it is installed at the same time as the selected package. If you do not want to install the dependency, click **Cancel**. The selected package and dependency will not be installed.
- If a dependency is not loaded, a message states that you must load the dependency before installing the selected package. Click **Cancel** to exit the Install window.

Caution: If you install multiple packages and one of the installations fails, the firewall will not reboot and a successful installation that requires a reboot will not be complete. You must investigate the cause of the installation failure and determine if you should manually restart the firewall or uninstall the package.

3 Select **Install now**.

4 [Conditional] If the package requires a reboot, select whether you want the reboot to occur after installation.

- If you select **Activate packages after installation**, the package is installed, the firewall reboots, and the package is activated after you click **OK**.
- If you clear **Activate packages after installation**, the package is installed after you click **OK** but the firewall does not reboot and the package is not activated. A rollback is required to activate the package after installation. You will be prompted to schedule a rollback to activate the package.

5 Click **OK**.

The package appears in the Manage Packages table with a status of *Installed*.

To uninstall a package:

1 Select a package in the table with the status of *Installed*.

2 Click **Uninstall**. The Uninstall window appears.

If the package has a dependency, the dependency appears in a lower table and will be uninstalled at the same time as the selected package. If you do not want to uninstall the dependency, click **Cancel**. The selected package and dependency will not be uninstalled.

3 Select **Uninstall now**.

4 Click **OK**.

The package appears in the Manage Packages table with a status of *Loaded*.

Scheduling automatic installs and uninstalls

You can schedule a time to install or uninstall packages.

1 On the Manage Packages tab, click **Schedule**. The Schedule Install/Uninstall window appears.

- Any package with a status of *Loaded* appears in the Select packages to install list.
- Any package with a status of *Installed* appears in the Select packages to uninstall list. (If a package cannot be uninstalled, it does not appear in the list.)

2 Select the packages you want to install or uninstall.

3 Select **Schedule for** and select a date and time for the action to take place.

4 Click **OK**.

The selected packages appear in the Manage Packages table with a status of *Install scheduled on <date>* or *Uninstall scheduled on <date>*.

To cancel a scheduled install or uninstall, select **Unschedule All** and click **OK**.

Viewing package information and activity logs

You can view information and activity for all packages listed in the Manage Packages table.

To view information and activity logs for a single package:

1 In the Manage Packages table, select a package.

2 Click **View Package Details**.

Figure 347 View Package Details button



The Details window appears.

- The Readme tab states what changes the package is making.
- The Package Log tab shows all load, install, and uninstall activities for the package.

To view all package installation activities on this firewall:

Click **View Log**. The View Log window appears.

The View Log lists detailed histories of all package installs and uninstalls on the firewall, including programs run, package parameters, and errors.

Enabling e-mail notification

To enable e-mail notification of software management activities:

- 1** Select **Enable e-mail notifications for install, uninstall, automatic load and rollback**.
- 2** In the **E-mail Address** field, enter the e-mail address of the person who will be notified. The default address is the administrator of this firewall.
- 3** Save your changes.

About the Software Management: Download Packages tab

To manually or automatically load software packages onto the Sidewinder, select **Maintenance > Software Management**, then click the **Download Packages** tab. The following window appears:

Figure 348 Software Management: Download Packages tab

The screenshot shows the 'Download Packages' tab in the 'Software Management' window. At the top, there are three tabs: 'Manage Packages', 'Download Packages' (which is active), and 'Rollback'. Below the tabs is the 'Automatic Load Configuration' section. It contains two checkboxes: 'Automatically check for and load packages' (which is checked) and 'Automatically check for available packages' (which is unchecked). To the right of these checkboxes is a 'Restore Defaults' button. Below the checkboxes is a 'Load using:' dropdown menu set to 'FTP'. Below that is a 'Directory:' text field containing 'packages/sidewinder/7.0.1'. Below the directory field are two rows of fields: 'Host:' with 'sidewinder.downloads.forcepoint.com' and 'Port:' with '21'. Below the host field are 'User Name:' with 'anonymous' and 'Password:' (empty). Below the frequency field is 'Frequency:' with a dropdown set to 'Weekly' and 'Confirm Password:' (empty). At the bottom left of the configuration section is a 'Perform Manual Load Now' button.

Use this tab to automatically or manually load packages on the Sidewinder. You can also configure a different site for the packages to be loaded from.

To configure your firewall to load packages automatically:

1 Select an automatic action:

- **Automatically check for and load packages** – To find available packages and load them on the firewall, select this option. Packages appear in the Manage Packages table with a status of Loaded. They can be installed.

This option works well if you regularly schedule package installations.

- **Automatically check for available packages** – To find available packages, select this option. A list of packages appears in the Manage Packages table with a status of Available. They can be loaded.

This option works well if downloading to your firewall is slow, or if you do not want to store and manage packages you will not use.

2 Identify the package site:

- **Load using** – Select the protocol used to transport the package.
- **Directory** – The path name on the site where the package is located.
- **Host** – The host name or IP address of the site where the package is located.
- **User Name** – The user account defined on the site.
- **Port** – The number of the port used to access the site.
- **Password** – The password to validate you to the site.
- **Confirm Password** – Verify the password.

Note: To restore the system default values to these fields, click **Restore Defaults**.

3 From the **Frequency** drop-down list, select how often you want the firewall to check for available packages.

Note: Use `cf crontab` to designate a specific time for the action to take place. See the *Command Line Interface Reference* at <https://support.forcepoint.com>.

4 Save your changes.

To manually load a package:

1 Click **Perform Manual Load Now**. The Download Packages: Manual Load window appears.

2 Identify the package and its location:

- **Load packages from** – Select the location of the package you are loading to the firewall. It can be an FTP or HTTPS site, a CD-ROM, or a file.
- **Directory** – The path name where the package is stored.
- **Packages** – The name of the package you are loading. If you are loading more than one package, separate each package name with a comma. Package names contain only alphanumeric characters. No spaces are allowed in package names.
- **Host** – [Conditional] The host name of the FTP or HTTPS site where the package is located.
- **User Name** – [Conditional] The user account defined on the FTP or HTTPS site.
- **Port** – [Conditional] The number of the port used to access the FTP or HTTPS site.
- **Password** – [Conditional] The password to authenticate to the FTP or HTTPS site.
- **Confirm Password** – [Conditional] Verify the password.

3 Click **OK**. Click **Yes** to confirm the load.

The package appears in the Manage Packages tab with a status of Loaded. The package can be installed.

About the Software Management: Rollback tab

To roll back the Sidewinder to a previous state, select **Maintenance > Software Management**, then click the **Rollback** tab. The following window appears:

Figure 349 Software Management: Rollback tab

Package Name	Status
70000	Installed on Thu Feb 22 10:02:19 2007
70000t01	Installed on Fri Feb 23 21:36:47 2007
70000t02	Loaded on Fri Feb 23 21:36:31 2007

Use the Rollback tab to revert the Sidewinder to a previous state. The table shows which packages will be on the firewall after the rollback.

- Any configuration changes made after the last package was installed are lost. A rollback reverts the firewall to the state just before the package was installed. Therefore, rolling back is a recommended recovery option for only a short time after a package installation.
- A rollback requires a reboot.

To roll back immediately:

- 1 Click **Rollback Now**. A warning message appears stating that configuration changes will be lost and that the firewall will reboot.
- 2 Click **Yes** to continue. The firewall reboots.

To schedule the rollback for a future time:

- 1 Select **Schedule rollback for** and select the desired date and time.
- 2 Save your changes. A warning message appears stating that configuration changes will be lost.
- 3 Click **Yes** to schedule the rollback.

Editing files

You might need to modify a text file or a configuration file. Although the typical UNIX editors are available (vi and emacs), you may find it easier to use the **File Editor** provided with the Admin Console. The File Editor simplifies the editing process, enabling you to perform virtually every necessary editing task from the Admin Console instead of using a command line.

The File Editor also provides some additional conveniences such as unique file backup and restore features. (UNIX aficionados are still welcome to use the editor of their choice if they prefer.) In addition, using the File Editor through the Admin Console provides a secure connection.

In general, use the Admin Console for configuration changes. If you do not have much command line experience, only edit files manually when instructed to do so by the documentation or Technical Support. If you have experience using the command line, remember that files may have been altered for security reasons and therefore may not behave as you expect. For all administrators, create a backup file before making changes so that you can, if necessary, quickly return your firewall to a functional configuration.

About editing Sidewinder files

Sidewinder files are not protected against simultaneous editing by two individuals. Whoever saves the file last usually prevails. In some cases, file corruption occurs.

Caution: An administrator should take care not to make changes to a file when another administrator is working on it.

For example, if an administrator is editing the *server.conf* configuration file using the Admin Console's File Editor while someone else is using a text editor to change that file, there may be undesirable results. If two people try editing the same file and both are using vi or both are using emacs, however, the editor will warn the users about the situation.


A frequent error to be aware of when manually editing the Sidewinder configuration files (*server.conf*, *roles.conf*, etc.) is the misuse of special characters that are used to format commands within these files. Special characters include double quotes, single quotes, brackets ([]), the pound symbol (#), and parenthesis (). Inadvertently placing special characters in the configuration files will render the files unreadable to the firewall. Enter `man sidewinder.conf` at a command prompt for details.

Note: Save any scripts you create for the firewall in the */usr/local/bin* directory. During software upgrades, the upgrade procedure will automatically save any scripts that reside in that directory.

Using the File Editor

To access the File Editor, select **Maintenance > File Editor**, and then click **Start File Editor**. The File Editor window appears.

Figure 350 File Editor window




The File Editor window contains three different menu options:

- **File** – This menu contains the basic action options. Use it to open new or existing files, and to save files. The File menu also provides two unique capabilities: it enables you to create a backup copy of a file, and it enables you to restore a file from a previously saved backup copy.
See [Creating a backup file in the File Editor](#) and [Restoring a file](#) for details.
- **Edit** – Use the functions in this menu to cut, copy, paste, and find/replace text.
See [Using the Find/Replace option](#) for information on finding and replacing text.
- **Help** – The following options are available under this menu:
 - **File Editor Help** – Displays specific information for the File Editor window.
 - **About Help** – Displays information about the current version of the Admin Console software.

Opening and saving files in the File Editor

To open a file or save a file with a different name or location: from the File menu, select **Open** or **Save As**. The Open File or Save As window appears.

Figure 351 Open File window



To open or save a file:

- 1 In the **Source** field, select where the source is located:
 - **Local File** – Indicates the file is located on the local Windows workstation or on a network connected to the workstation.
 - **Firewall File** – Indicates the file is located on the firewall.
- 2 In the **File** field, type the full path name of the file.

If you do not know the full path name, click **Browse** to browse the available directories. When you locate the file, click **OK**. The file name appears in the **File** field.
- 3 Click **OK** to open or save the file, or click **Cancel** to cancel the request.


Creating a backup file in the File Editor

When modifying firewall configuration files, it is normally a good practice to create a backup copy of the file before you begin editing the file. That way, if you make a mistake while editing the file you can revert to the original file.

The File Editor provides an easy method for creating a backup copy of a file. You can even make a backup after you begin modifying a file. The key is to create the backup before you save your changes. Once you save your changes you will not be able to create a backup file that mirrors the original file.

To make a backup copy of a file, open the file. From the **File** menu, select **Backup**. The Backup File window appears.

Figure 352 Backup File window



- 1 In the **Name of Backup File** field, specify a name for the backup file. By default, the file is given the same name as the original file but with a **.bak** extension.

The backup file will be created in the directory listed in the **Current Directory** field. This is the directory in which the original file currently resides, and cannot be modified.
- 2 Click **OK** to save the information and exit the window, or click **Cancel** to exit the window without saving the backup file.


Restoring a file

Use the Restore function to restore a file to its original contents.

- The file must be open within the File Editor.
- You must have previously created a backup copy of the file.

From the **File** menu, select **Restore**. The Restore File window appears.

Figure 353 Restore window



- 1 In the **Restore From File** field, specify the name of the backup file to use when restoring the file to its original condition. If you do not know the name of the backup file, click **Select** to browse the available files. When you locate the file, click **Open**. The file name appears in the **Restore From File** field.

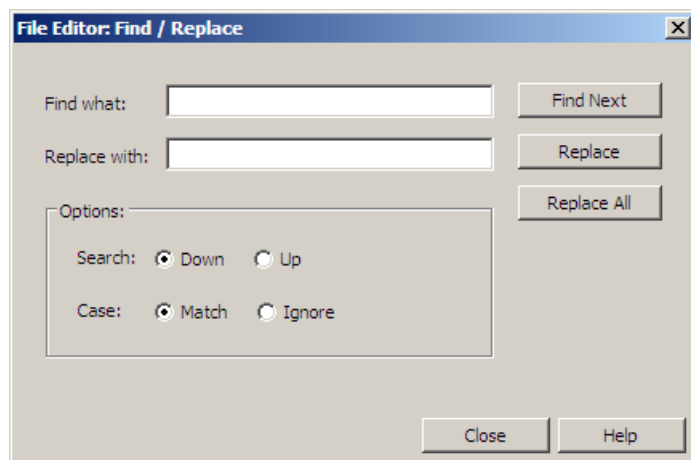
Note: If a backup file exists, it will appear in the same directory as the current file, because you are only allowed to create a backup in the same directory. The **Current Directory** field displays the name of that directory and cannot be modified.

- 2 Click **OK** to save the information and exit the window, or click **Cancel** to exit the window without restoring from the backup file.

Using the Find/Replace option

Use the Find/Replace function to locate character strings, and to replace them with different character strings. From the **Edit** menu, select **Find/Replace**. The Find/Replace window appears.

Figure 354 Find/Replace window



- 1 In the **Find what** field, specify the character string you want to search for within the file.
- 2 [Optional] If you want to replace the character string specified in the **Find what** field with a different character string, type the new string in the **Replace with** field.
- 3 In the **Search** field, specify which direction in the file the search should be performed:
 - **Down** – From your current position within the file, the File Editor will search down (forward) in the file for the specified character string.
 - **Up** – From your current position within the file, the File Editor will search up (backward) in the file for the specified character string.
- 4 In the **Case** field, specify whether the File Editor should find *any* matching character string, or if it should consider upper and lower case when performing the search:
 - **Match** – Find only those character strings that exactly match the case as specified in the **Find what** field.
 - **Ignore** – Find all matching character strings regardless of upper and lower case.
- 5 Click **Find Next** to start the character search and to locate the next occurrence within the file.
[Optional] If the character search locates a match, you can click **Replace** to replace the found character string with the character string specified in the **Replace with** field. To replace all occurrences of the character string, click **Replace All**. An Info window will appear indicating how many times the character string was replaced. Click **OK** to close the Info window.
- 6 To find additional occurrences of the character string, continue to click **Find Next** for each occurrence. When there are no additional occurrences, a message will appear telling you that the search is complete.
- 7 When you are finished searching, click **Close** to exit this window.

Checking file and directory permissions (ls command)

Standard UNIX permits access to files based on a process' user and group identifiers and the file's permissions (mode bits that indicate who can read, write, or execute a file). As described in [Chapter 1, Introduction to Sidewinder](#), Sidewinder's Type Enforcement mandatory security policy is the ultimate authority on if, and when, a given process may access files and it overrides standard UNIX permissions. A Sidewinder file that appears to be accessible based on standard UNIX permissions can be denied by the Type Enforcement (TE) policy.

To check Type Enforcement, enter the following commands:

- for files: `/bin/ls -alZ filename`
- for directories: `/bin/ls -dlZ directory_name`

You will see output similar to the following:

```
secureos/Slog:logs 3965 Feb 23 23:55 cron
```

↑ Creating Domain (such as *Admn*, *\$Sys*, *mtac*)

↑ Type (*diry* for directory, and *exec*, *file*, *conf*, etc. for files)

↑ File size

↑ Time stamp

↑ File or directory name

Changing a file's type (chtype command)

Type Enforcement assigns each file and directory a *type*. In general, you should not, and cannot, change this type. In the rare situation where you need to change the type, use the **chtype** command.

To change type on a file or directory:

- 1 At a command prompt, log in and enter the following command to switch to the *Admn* role:

```
srole
```

- 2 Copy the file or directory you want to change:

```
cp file1 newfile
```

- 3 Delete the original file:

```
rm file1
```

- 4 Change the new file to the target domain and/or file type:

```
chtype domain:filetype newfile
```

- 5 Rename the file or directory:

```
mv newfile file1
```

Creating your own scripts

While operating in either the *User* or *Admn* domains, you can create your own scripts for use on the firewall. Scripts created in the *User* domain are executable by the *Admn* and *User* domain but no other domain. Scripts created in the *Admn* domain cannot be executed by anyone except the administrator.

Registering with Forcepoint Sidewinder Control Center

Use the Control Center Registration window to register this Sidewinder to a Control Center Management Server.

- If you chose to auto-register to Control Center during initial configuration of this Sidewinder, do not register with Control Center here. Use the Sign Up Firewalls dialog in the Control Center Configuration Tool to initiate rapid deployment.
- A Sidewinder administrator account named *ccfwadmin* is automatically created when you register with Control Center. If you already have an administrator account with that name, it will be overwritten.

To register this firewall with Control Center, select **Maintenance > Control Center Registration**. The Control Center Registration window appears.

Figure 355 Control Center Registration window with Configure backup server selected

The screenshot shows the 'Control Center Registration' window. At the top, it displays 'Firewall hostname: fw11v191.example.net'. Below this is a section titled 'Control Center servers'. Inside this section, there is a 'Primary Server:' label followed by 'Name:' and 'IP Address:' fields, each with a search icon. Below the primary server fields, there is a checked checkbox labeled 'Configure backup server'. Under this checkbox, there are 'Backup Server Name:' and 'IP Address:' fields, also with search icons. At the bottom of the window, there is a button labeled 'Register with the Control Center now'.

To register this firewall:

- 1** In the **Name** field, enter the host name of the Management Server that will manage this firewall. If you are also using backup servers, use the host name of the active Management Server.
- 2** In the **IP Address** field, enter the IP address of the Management Server.
- 3** [Optional] If you are using a High Availability Management Server configuration, select the **Configure backup server** check box.
 - a** In the **Backup Server Name** field, enter the host name of the Management Server acting as a backup to the active Management Server.
 - b** In the **IP Address** fields, enter the IP address of the corresponding backup servers.
- 4** Click **Register with the Control Center now**. Click **Yes** to confirm your changes. The Authentication window appears.
- 5** Enter the user name and password of the Control Center administrator, then click **OK**.

A “registration succeeded” message appears.
- 6** Click **OK**.

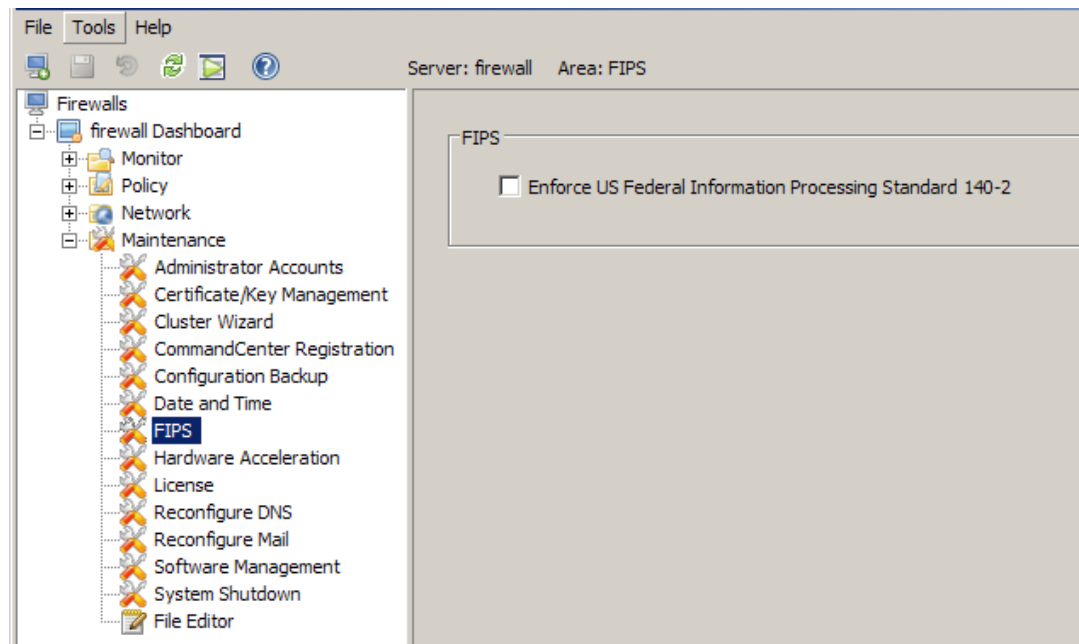
To complete registration, go to the Management Server Configuration Tool. See the *Product Guide* for Control Center for more information.

Enforcing FIPS

Federal Information Processing Standard (FIPS) 140-2 is a standard that describes the U.S. federal government requirements for a cryptographic module used in a security system. Select this option to configure settings that make a Sidewinder FIPS 140-2 compliant. For more information on how enabling this option affects the Sidewinder, see the FIPS application note at <https://support.forcepoint.com>.

Note: This option is appropriate only for organizations that are explicitly required by the U.S. federal government to be FIPS 140-2 compliant.

Figure 356 FIPS window



To enable FIPS:

- 1 Select **Maintenance > FIPS**. The FIPS check box appears in the right pane.
- 2 Select **Enforce US Federal Information Processing Standard**.
- 3 Save the configuration change.
- 4 Select **Maintenance > System Shutdown** and reboot the firewall to the Operational kernel to activate the change.


Enabling hardware acceleration

If you use SSL decryption, you can use a supported hardware accelerator card in your Sidewinder to offload decryption, increasing system performance.

- If you do not have a supported hardware accelerator card installed on your firewall and would like to use one, contact your sales representative for assistance.
- To install, consult the product documentation for the accelerator and chassis.

To enable a hardware accelerator card, select **Maintenance > Hardware Acceleration**. The Hardware Acceleration window appears.

Figure 357 Hardware Acceleration window



Hardware Accelerators:

Device Name	SSL Acceleration	Enabled
cavium	RSA, DSA, DH	on

The table displays information about the hardware acceleration card that is installed on your Sidewinder.

- Click **Enable** (green check mark) to enable the card.
- Click **Disable** to disable the card.

Configuring UPS

Many organizations connect Sidewinder to an Uninterruptible Power Supply (UPS). This allows the firewall to continue to be operational if a power outage occurs. If the power outage is long enough, however, the battery in the UPS will begin to fail. To avoid an uncontrolled shutdown, you can configure the firewall to initiate an orderly shutdown before the UPS fails. The firewall is much more likely to restart in a good condition following an orderly shutdown than from an uncontrolled shutdown.

Perform this procedure before configuring UPS:

- 1 From a console attached to the firewall, log in and enter `su role` to switch to the Admn domain.
- 2 [Conditional] If the UPS device will be connected to the firewall's COM1 port, type the following command:


```
setconsole video
```

- 3 Restart the firewall.
- 4 Enter the firewall's BIOS.
- 5 Disable console redirection for the COM port that the UPS will be connected to.
- 6 Save your changes and exit the BIOS.

To configure the firewall to use a UPS:

- 1 Connect the UPS's serial cable to the appliance's COM1 or COM2 port. (Use a simple signaling cable.)
 - You must connect the UPS device to the firewall appliance before enabling UPS in the Admin Console. If you enable UPS without an attached UPS device, the firewall might shut down immediately.
 - Each member of an HA cluster must use the same COM port.
- 2 Select **Maintenance > UPS**. The UPS window appears.

Figure 358 UPS window



- 3 Select **Enable Uninterruptible Power Supply (UPS)**.
 - If a UPS is enabled and a power outage occurs, the firewall will monitor the UPS and will perform an orderly shutdown when the UPS battery begins to run low.
 - If a UPS is not enabled, and a power outage occurs, and the appliance *is* connected to a UPS, the firewall will not monitor the UPS and will not perform an orderly shutdown when the UPS battery begins to run low.
- 4 From the **UPS Serial Port** drop-down list, select the port the UPS is connected to.
- 5 In the **Battery Time** field, specify the estimated amount of time (in seconds) that the UPS battery will last before running low. The firewall will initiate an orderly shutdown when this timer expires, regardless of the amount of battery power remaining in the UPS.
- 6 Save your changes.

22 Certificate/Key Management

Contents

[About Certificate/Key Management](#)

[Managing firewall certificates](#)

[Managing certificate authorities](#)

[Managing VPN certificates](#)

[Exporting certificates](#)

[Managing SSH keys](#)

About Certificate/Key Management

Certificates are used to verify the identity and authenticity of hosts during electronic communication. A certificate can be thought of as the digital equivalent of a driver's license. Certificates are used together with encryption to ensure that communication sent over a network is secure. They are commonly used to associate a user or device with the appropriate public/private key pair for use with public key cryptography.

Public key cryptography, also known as asymmetric cryptography, is an encryption method in which each participant has a private key that is kept secret and a public key that can be distributed to anyone. The public and private keys are mathematically related so that data that is encrypted using the public key can only be decrypted using the corresponding private key. Certificates and public key cryptography are used in the Secure Sockets Layer (SSL) protocol.

The Forcepoint Sidewinder uses certificates for:

- VPN authentication and identity management
- SSL decryption
- firewall administration services

The firewall's SSH proxy also presents SSH host keys to SSH clients. See [Managing SSH keys on page 645](#).

The following topics are covered in this chapter:

- [Managing firewall certificates on page 623](#)
- [Managing certificate authorities on page 629](#)
- [Managing VPN certificates on page 633](#)
- [Exporting certificates on page 643](#)
- [Managing VPN certificates on page 633](#)
- [Managing SSH keys on page 645](#)

Understanding Distinguished Name syntax

The Certificate Manager supports using distinguished names (DN) for a number of purposes, including identifying the subject of an X.509 certificate. DNs need to be entered using the proper syntax. As defined in the X.500 specifications, a DN is an Abstract Syntax Notation One (ASN.1) value. Within an X.509 certificate, a DN is represented as a binary value. When it is necessary to represent a DN in a human-readable format, as when entering information into the Certificate Manager, the firewall uses the string syntax defined by RFC 2253. This section summarizes the DN string syntax through a series of examples.

Note: For more information on this string syntax, visit <http://www.ietf.org/rfc.html> and search for RFC 2253, "Lightweight Directory Access Protocol (v3): UTF-8 String Representation of Distinguished Names."

A distinguished name (DN) consists of a sequence of *identity components*, each composed of a type tag and a value. The components of a DN are sets of attribute type/value pairs. The *attribute type* indicates the type of the item, and the *attribute value* holds its contents. Each type/value pair consists of an X.500 attribute type and attribute value, separated by an equal sign (=). In the example CN=Jane Smith, "CN" is the attribute type and "Jane Smith" is the value.

The attribute type/value pairs are separated by commas (,). This example shows a DN made up of three components:

```
CN=Jane Smith,OU=Sales,O=Forcepoint
```

Plan out your organization's certificate identification needs before creating any DNs. DNs have a hierarchical structure, reading from most specific to least specific. No preset hierarchy of attribute type exists, but the structure for a given organization needs to be consistent. In this example, the organization Forcepoint has organizational units, making the organizational unit attribute type more specific than the organization attribute type.

```
CN=Jane Smith,OU=Sales,O=Forcepoint
```

```
CN=Ira Stewart,OU=Engineering,O=Forcepoint
```

An attribute type is specified by a tag string associated with the X.500 attribute being represented. The Sidewinder supports the attribute tag strings displayed in [Table 69](#), which includes the most common ones recommended by RFC 2253. The tag strings are not case sensitive.

Table 69 Supported X.500 Attribute Type Tags

Tag String	X.500 Attribute Name	Character String Type	Maximum Number of Characters
C	CountryName	PrintableString	2
CN	CommonName	DirectoryString	64
Email Address	EmailAddress	IA5String	128
L	LocalityName	DirectoryString	128
O	OrganizationName	DirectoryString	64
OU	OrganizationUnitName	DirectoryString	64
SN	Surname	DirectoryString	128
ST	StateName	DirectoryString	128
Street	StreetAddress	DirectoryString	128
UID	UserID	DirectoryString	128

The attribute value holds the actual content of the identity information, and is constrained by the associated attribute type. For the supported attribute types, [Table 70](#) shows the corresponding string type (which limits the allowed set of characters) and its maximum length. For example, given "CN=Jane Smith" as a name component, the string "Jane Smith" is of type DirectoryString, and is constrained to a maximum of 64 characters. The maximum number of characters allowed in a DN (that is, the number of characters for all attribute values added together) is 1024.

Table 70 defines the allowed character set for each of the character string types used in Table 69.

Table 70 Character String Types

Character String Type	Allowed Characters
DirectoryString	All 8 bit characters without encoding. All non-8 bit characters with UTF-8 encoding
PrintableString	A–Z, a–z, 0–9, ()+-./:=?, comma (','), space (' '), apostrophe ('')
IA5String	All 7 bit characters

When representing attribute values, be careful when using special characters. The following characters have special meaning in the string syntax and must be preceded by a backslash character (\):

- comma (,)
- equal sign (=)
- plus sign (+)
- less than sign (<)
- greater than sign (>)
- pound sign (#)
- semicolon (;)
- backslash (\)
- quotation (")

All other printable ASCII characters represent themselves. Non-printable ASCII must be have a backslash preceding the ordinal value of the character in two-digit hexadecimal (for example, the BEL character, which has an ordinal value of seven, would be represented by \07). Here are some examples of the escape conventions:

```
CN=Jane Smith\, DDS, OU=Sales, O=Forcepoint
```

```
CN=\4a\61\6e\65\20Smith, OU=Sales, O=Forcepoint
```

Attribute values may optionally be contained within double-quote characters, in which case only the backslash (\), double quote ("), and non-printable ASCII characters need to be preceded by a backslash. Here the double-quotes eliminate the need to escape the CN's comma:

```
CN="Jane Smith, DDS", OU=Sales, O=Forcepoint
```

Note: Entries containing backslashes or double-quotes will appear “normalized” (without extra characters or spaces) in the GUI once they are saved.

Use this supported syntax when entering information on the Admin Console's Certificate Manager tabs.

Note: For additional information on DN syntax, see RFCs 2044, 2252, 2253, and 2256.

Selecting a trusted source

If you have decided to use certificate authentication, you must choose whether to use a single certificate or Certificate Authority root certificate. In both methods, when a key is generated, the trust point (the firewall or a trusted CA like Netscape, Entrust, etc.) places the key in an electronic envelope called an X.509 certificate. Every certificate contains a collection of information about the entity possessing the private key (the firewall or VPN client). This information may include an identity, a company name, and a residency.

Note: If you select Netscape as a CA server, note that only Netscape version 4.2 is supported at this time.

Single certificate versus Certificate Authority trusted sources

To validate this information, a certificate must be electronically verified and witnessed by a trusted source.

- A CA-based trusted source is best designed for larger deployments and allows for greater flexibility, as both the root (general authoritative certificate from the CA) and personal certificates may be retrieved online. However, a CA configuration does require managing the Certificate Authority server or paying someone else to manage it for you.
- A firewall self-signed trust source is best for very small deployments, as a separate VPN definition must be created for each client. Certificates must be exported from the firewall and then installed on each client.

Public versus private Certificate Authorities

If you are planning to use a specific Certificate Authority to validate certificates created on the firewall, or as part of a group of trusted CAs from which the firewall can directly import certificates, you should set up these CAs before you begin configuring a VPN. You can use the following types of CA servers:

- **Private CA server** – You can purchase and install your own CA server and configure this server as the trusted authority for any VPNs you establish. This is an ideal solution for companies that prefer to allow only VPNs with certificates signed by a CA server on their own protected network.

Note: Before you begin, you must install the CA server and make its URL accessible to the firewall. For details on installing and configuring a private CA server, review the manufacturer's documentation.

- **Public CA server** – You can choose to accept certificates signed by trusted CAs administered elsewhere. This option allows remote machines to use one certificate for VPNs with more than one corporate partner.

Managing firewall certificates

A firewall certificate is used to identify the firewall to a potential peer in the following scenarios:

- VPN authentication
- SSL decryption
- Firewall administration services

When creating a certificate for the firewall, you have the option to submit the certificate to a CA for validation, or have the firewall generate a self-signed certificate. You should create these certificates before you begin configuring a VPN.

Select **Maintenance > Certificate/Key Management**, then select the **Firewall Certificates** tab. The following window appears:

Figure 359 Firewall Certificates tab

The screenshot shows the 'Firewall Certificates' tab in the Forcepoint Sidewinder interface. The window has a tabbed header with the following tabs: Remote Certificates, Firewall Certificates (selected), Certificate Authorities, Remote Identities, SSL Certificates, Certificate Server, and SSH Keys. On the left, there is a list of certificates under the heading 'Certificates:'. The list contains two entries: 'Default_SSL_Cert' and 'MFE_Communication_Ce'. Below the list are 'New' and 'Delete' buttons. At the bottom of the window are 'Import ...', 'Export ...', and 'N/A' buttons. On the right side of the window, the properties for the selected certificate are displayed. The 'DN - Distinguished Name:' field contains the text 'CN=fw11v191.example.net,O=Secure Computing,C=US'. Below this are input fields for 'E-Mail Address:', 'Domain Name:', and 'IP Address:'. The 'Signature type:' is set to 'RSA' and the 'Status:' is set to 'SELF'.

Note: Use the Firewall Certificates tab to view the list of available certificates. The firewall will use a firewall certificate to identify itself to a peer in a VPN connection. To display the properties of a specific certificate, select the certificate from the list and its properties are displayed on the right portion of the window. For a description of these properties, see [Creating firewall certificates on page 624](#). You cannot modify the properties of a certificate from this window. To modify a certificate, you must delete it and then add it back using the new properties.

From this tab, you can perform the following actions:

- **Create a firewall certificate** – Click **New** to add a certificate to the Certificates list. See [Creating firewall certificates on page 624](#) for details.
- **Delete a firewall certificate** – Select the certificate and click **Delete** to remove the selected certificate from the Certificate list.
Note: A certificate cannot be deleted if it is currently used by one or more areas (for example, VPN definitions, Application Defenses, etc.).
- **Import a firewall certificate** – Click **Import** to import an existing certificate and its related private key file. See [Importing firewall certificates on page 626](#) for more information.
- **Export a firewall certificate** – Click **Export** to export the selected certificate to a file. The export function is generally used when capturing the certificate information needed by a remote partner such as a VPN client. See [Exporting certificates on page 643](#) for more details.
- **Retrieve a certificate** – If a certificate request has been submitted to be signed by a CA, click the **Query** button to query the CA to see if the certificate is approved. If yes, the Status field will change to `SIGNED` and the approved certificate will be retrieved.

If the certificate request is Manual PKCS10, click the **Load** button to load the signed certificate from a file supplied by the CA.

Note: By default, Netscape CAs and CAs that support the Simple Certificate Enrollment Protocol (SCEP) are checked every 15 minutes for any certificates waiting to be signed.

Creating firewall certificates

Use the Create New Firewall Certificate window to add a certificate to the Firewall Certificate list.

Note: The default certificate key size is 1024 bits. The default lifetime for self-signed certificates created on the firewall is five years.

To create a firewall certificate, select **Maintenance > Certificate/Key Management**, then select the **Firewall Certificates** tab and click **New**. The Create New Certificate window appears.

Figure 360 Firewall Certificates: Create New Certificate window

Firewall Certificates: Create New Certificate

Certificate Name:

Distinguished Name:

E-Mail Address:

Domain Name:

IP Address:

Submit to CA:

Signature type: ☒ RSA ☐ DSA

Other Parameters:

To add a certificate:

- 1** In the **Certificate Name** field, type a name for this certificate.
- 2** In the **Distinguished Name** field, create a distinguished name. See [Understanding Distinguished Name syntax on page 619](#) for information on the format that should be used. Note the following:
 - The order of the specified distinguished name fields must match the order listed in the certificate.
 - Some CAs will not support the optional identity types specified in [Step 3](#) through [Step 5](#).
- 3** [Optional] In the **E-Mail Address** field, type the email address associated with this firewall certificate.
- 4** [Optional] In the **Domain Name** field, type the domain name associated with this firewall certificate.
- 5** [Optional] In the **IP Address** field, type the IP address associated with this firewall certificate.
- 6** In the **Submit to CA** drop-down list, select the enrollment method to which the certificate will be submitted for signing. The valid options are:
 - **Self Signed** – Indicates the new certificate will be signed by the firewall rather than by a CA.
 - **Manual PKCS10** – Indicates the certificate enrollment request will be placed in a PKCS10 envelope and exported to the file designated in the **Generated PKCS10 File** field.
 - The name of the CA to which the certificate is submitted for signing. The CA can be either private (one you own and manage) or it can be public (a trusted CA administered elsewhere).
- 7** In the **Signature Type** field, select the encryption format that will be used when signing the certificate. Select the format supported by the remote VPN peer. Valid options are **RSA** or **DSA**.
 - RSA is faster at signature verification. It is the most commonly used encryption and authentication algorithm.
 - DSA is faster at signature generation.
- 8** [Conditional] Depending on the method you select in the **Submit to CA** field, the **Other Parameters** area may contain additional fields, as described below:
 - If you selected **Manual PKCS10** in the **Submit to CA** field, the **Generated PKCS10 File** field appears. Specify the name and location of the file that will contain the signed certificate, or click **Browse...** to browse the network directories for the location of the file you want to specify. This file contains a PKCS10 “envelope” that is used to send a certificate to a CA for signing.
 - If you selected a method that uses SCEP, you will need to provide a password in the **SCEP Password** field that appears.
- 9** [Conditional] In the **Format** field, select the appropriate format for your PKCS10 certificate request.
- 10** Click **Add** to add the certificate to the Certificates list.
- 11** Save your changes.

Importing firewall certificates

You can import a certificate to the list of firewall certificates defined on the firewall.

To import a firewall certificate, select **Maintenance > Certificate/Key Management**, then select the **Firewall Certificates** tab and click **Import**. The Import Firewall Certificate window appears.

Note: The displayed fields will vary depending on which import source you select.

Figure 361 Firewall Certificates: Import Certificate window

The screenshot shows a window titled "Firewall Certificates: Import Certificate". It contains the following fields and controls:

- Import source:** A dropdown menu with "File" selected.
- Certificate Name:** A text input field.
- Import Certificate From File:** A text input field with a "Browse..." button to its right.
- Import Private Key from file:** A text input field with a "Browse ..." button to its right.
- Password:** A text input field.
- Buttons:** "OK", "Cancel", and "Help" buttons at the bottom.

To import a certificate:

- 1 In the **Import Source** field, select either **File** or **Encrypted File (PKCS12)**.

Note: The available fields will vary based on the import source you select.

- If you select **File**, you must identify the file on the **Import Certificate From File** field.
- If you select **Encrypted File (PKCS12)**, specify the certificate and key file.

- 2 In the **Certificate Name** field, type a local name for the certificate you are importing.

- 3 In the **Import Certificate From File** or the **Import Certificate/Key** field, type the name and location of the certificate file you will import. You may also click **Browse** to browse the network directories for the location of the file(s) you want to specify.

- 4 [Conditional] In the **Private Key File** field, type the name and location of the private key file associated with this certificate, or click **Browse** to browse the network directories for the location of the file(s) you want to specify. The file can be in either PK1 or PK8 format. (This field is only available if the **Import Source** field displays **File**.)


- 5 [Conditional] In the **Password** field, enter the password to decrypt the imported file. This password must match the password given when the file was encrypted.

Note: This field is only available if the Import Source field displays **Encrypted File (PKCS12)**.

Loading manual firewall certificates

If you created a manual firewall certificate, you must retrieve the certificate after it is signed by the CA; the firewall will not retrieve it automatically. For this process, the **Load** button appears when an unsigned requested certificate name is highlighted. Clicking this button will initiate the process to retrieve and import the signed certificate. The Load Certificates for PKCS 10 Request window appears.

Figure 362 Firewall Certificates: Load Certificate for PKCS 10 Request window



Use the Load Certificate for PKCS 10 Request window to load signed certificates. It also functions to query an LDAP server for whether or not a requested certificate is signed.

To load a signed certificate:

- 1 In the **Certificate Source** field, select the source location of the certificate. The following options are available:
 - **File:** Indicates you will manually specify the location of the certificate.
 - **LDAP:** Indicates you will access the services of an LDAP (Lightweight Directory Access Protocol) directory to locate the certificate. The LDAP server can be version 2 or version 3.
 - **Pasted PEM Certificate:** Indicates you will paste or type in the certificate from another source, such as another open application window or personal communication.
- 2 [Conditional] In the **Certificate from File** field, if the certificate source is a file, type the location or **Browse** to the location.
- 3 [Conditional] In the **Manual (pasted) PEM Certificate** field, if the certificate source is a Pasted PEM Certificate, type or paste the certificate in this field.
- 4 Click **OK** to issue a query command for your requested certificate, or click **Cancel** cancel the certificate request.

If you click **OK** and the certificate is available, it will automatically be imported and the status will change to SIGNED.
- 5 Save your changes.

Assigning new certificates for Admin Console services

The default SSL certificates are unique to each firewall. However, if you would like to change your default certificate for any reason, follow the steps in this section.


Before assigning a new certificate to these services you must first create the new certificates. You should create two new certificates, one for the Admin Console service and one for the synchronization server. You create the certificates from the **Firewall Certificates** tab. Each certificate must be:

- a firewall certificate
- a self-signed certificate
- of type RSA/DSA

See [Configuring and displaying remote certificates on page 638](#) for information on creating a firewall certificate.

To assign a new certificate for the Admin Console or the synchronization server, select **Maintenance > Certificate/Key Management**, then select the **SSL Certificates** tab. The SSL Certificates tab appears.

Figure 363 SSL Certificates tab



Use this tab to assign a new SSL certificate to the Admin Console service (cobra) and other firewall services.

The SSL Certificate tab allows you to view the proxies to which you can assign new certificates and identifies the name of the certificate currently assigned to each proxy. The certificate will either be 1) the default certificate or 2) a self-signed, RSA/DSA firewall certificate that is defined on the Firewall Certificates tab.

To assign a new certificate to a selected proxy, click **Modify** and select a new certificate to assign to this proxy.

Note: You will receive a warning message if you click **Modify** and there is not at least one self-signed RSA/DSA firewall certificate currently defined on the firewall. See [Configuring and displaying remote certificates](#) for information on defining this type of certificate.

Managing certificate authorities

This section explains how to configure the Certificate Authorities tab and display the imported signed root certificate. CAs are used to validate (sign) certificates that are used in a VPN connection or firewall certificates used for SSL termination. Select **Maintenance > Certificate/Key Management**, then click the **Certificate Authorities** tab. The Certificate Authorities tab appears.

Figure 364 Certificate Management: Certificate Authorities tab

The screenshot shows the 'Certificate Authorities' tab in the Forcepoint Sidewinder interface. The interface has a top navigation bar with tabs: Remote Certificates, Firewall Certificates, Certificate Authorities (selected), Remote Identities, SSL Certificates, Certificate Server, and SSH Keys. On the left, a list of 'Cert Authorities' includes DOD_CA-11 through DOD_CA-26 and DOD_EMAIL_CA-11 through DOD_EMAIL_CA-19. Below this list are 'New' and 'Delete' buttons. The main area displays details for 'DOD_CA-11'. The 'DN - Distinguished Name' field contains 'CN=DOD_CA-11,OU=PKI,OU=DoD,O=U.S. Government,C=US'. The 'Type' is set to 'Manual'. There are input fields for 'CA Address' and 'CA ID'. Below these are buttons for adding (+), removing (-), and moving (up/down) items. An 'OCSP Responder Urls' section is also present. At the bottom, there are buttons for 'Get CA Cert', 'Export...', and 'Get CRL'. A 'Ticket:' field is visible in the bottom right corner.

Use the Certificate Authorities tab to add certificate authorities (CAs) and to view the list of available CAs. CAs are used to validate (sign) certificates that are used in a VPN connection.

To display the properties of a specific certificate, select the certificate from within the **Cert Authorities** list. Its properties are displayed on the right portion of the window. For a description of these properties, see [Adding certificate authorities on page 631](#).

From this tab, you can perform the following actions:

- **Add a new certificate authority to the list** – Click **New** and make entries in the New Certificate Authority window.

See [Adding certificate authorities on page 631](#) for details.

- **Delete a certificate authority from the list** – Select the certificate you want to delete and click **Delete**.

Note: A Certificate Authority cannot be deleted if it is currently being used by one or more definitions (the **Delete** button is disabled).

- **Retrieve a CA certificate** – Click **Get CA Cert** to query the CA and import a certificate for the selected CA. The selected CA must be either Netscape 4.2 or an SCEP (Simple Certificate Enrollment Protocol) CA.

- **Export a CA certificate** – Click **Export** to export a CA certificate from the local cache to a file and/or a screen.

See [Exporting certificate authorities on page 631](#) for details.

- **Retrieve a CRL** – Click **Get CRL** to manually retrieve a new Certificate Revocation List (CRL) for this CA.

- A CRL identifies certificates that have been revoked. CRLs expire on a regular basis, which is why you must periodically obtain a new CRL.
- You generally only need to manually get a CRL for Netscape CAs when the CA is initially added. After that CRLs are automatically updated every 15 minutes or so for Netscape 4.2 CAs.


Note: If you do not have access to a Netscape CA or if you do have access to an LDAP directory, you should disable the **Perform CRL Checking** button on the Certificate Server window.

Adding certificate authorities

Use the New Certificate Authority window to add a new Certificate Authority to the list of CAs used when authorizing certificates in a Sidewinder VPN connection.

To create a certificate authority, select **Maintenance > Certificate/Key Management**, then select the **Certificate Authorities** tab and click **New**. The New Certificate Authority window appears.

Figure 365 Certificate Authorities: New Certificate Authority window



To add a new Certificate Authority:

- 1 In the **CA Name** field, type a name for this certificate authority. Only alphanumeric characters are accepted in this field.
- 2 In the **Type** drop-down list, select the type of CA used by your location:
 - **Manual** – Indicates that the necessary files are obtained and loaded by an administrator rather than by a CA.
 - **Netscape 4.2** – Indicates that a Netscape version 4.2 CA is being defined.
 - **SCEP** (Simple Certificate Enrollment Protocol) – Indicates the CA being defined supports this widely used certificate enrollment protocol. The CA can be of any type (Netscape 4.2, Entrust, VeriSign, etc.) as long as it supports SCEP.
- 3 [Conditional] In the **File** field, type the name and location of the root certificate for the CA, or click **Browse** to browse your network directories for the location of the root certificate. The root certificate is used to verify certificates issued by this CA. (This field is available only if you select **Manual** in the **Type** field.)

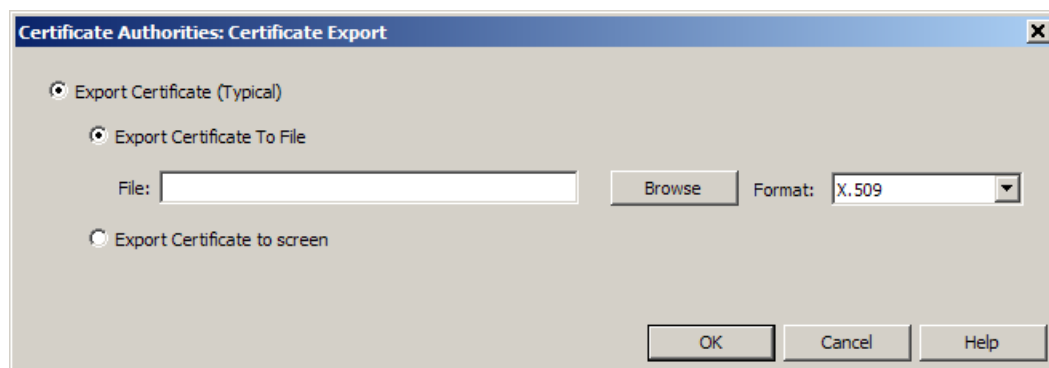
Note: Valid file formats are **X.509** and **ASN.1** (.pem or .der). For information on obtaining a root certificate, see the documentation that accompanied the CA.
- 4 [Conditional] In the **URL** field, type the URL address of the Netscape CA in the **URL** field. Certificates that need to be signed by the CA are sent to this address. (This field is available only if you select **Netscape** or **SCEP** in the **Type** field.)
- 5 [Optional] In the **CA Id** field, type the value used to identify this specific CA. Check with your CA administrator to determine the identifier to use. Many administrators use the fully qualified domain name of the CA as the identifier. (This field is available only if you select **SCEP** in the **Type** field.)
- 6 Click **Add** to add the CA to the Certificate Authority list.
- 7 Save your changes.

Exporting certificate authorities

Use the Certificate Export window to export the selected CA certificate from the firewall to a separate file and/or to the screen. The certificate can be written to a file on the hard drive of a workstation, or it can be written to a transportable medium such as a diskette.

To create a certificate authority, select **Maintenance > Certificate/Key Management**, then select the **Certificate Authorities** tab and click **Export**. The Certificate Export window appears.

Figure 366 Certificate Authorities: Certificate Export window



To export the certificate:

- 1 Select the **Export Certificate (Typical)** radio button.
- 2 Select the export destination:
 - **Export Certificate To File** – Select this option to export the certificate to a file. Continue with [Step 3](#).
 - **Export Certificate To Screen** – Select this option to export the certificate to the screen.
- 3 [Conditional] If you are exporting the certificate to file, do the following:
 - In the **File** field, type the name and location of the file to which the CA certificate will be written. If you want to overwrite an existing file, but you are not certain of the path name or the file name, click **Browse**.
 - In the **Format** field, select the appropriate format for the file.
- 4 Click **OK** to export the certificate to the desired location.

The certificate has now been exported.

Managing VPN certificates

If you are using automatic key generation and intend to use certificates for authentication, you should configure the certificate and/or Certificate Authority (CA) server information before you set up the VPN. This eliminates the need to configure certificates and CAs during the VPN process. To configure certificate or CA information, follow these general steps.

- 1 Review the section [Selecting a trusted source on page 622](#) for details on certificates and CAs.
- 2 Decide if you will use a public CA server, your private CA server, or self-signed certificates generated by the firewall (which can be used between two firewalls or between a firewall and a VPN client machine).
- 3 If you are using a public or private CA server, go to [Managing certificate authorities on page 629](#). You may also want to add remote identities to be used in conjunction with a Certificate Authority policy. See [Configuring and displaying remote identities on page 636](#).
- 4 If you are using self-signed certificates, refer to the section titled [Configuring and displaying remote certificates on page 638](#).
- 5 If you are configuring a VPN between the firewall and VPN client software, and if you are not using a CA, you must create a remote certificate, export it, then import the certificate into the VPN client. Refer to the following sections:
 - [Managing VPN certificates on page 633](#)
 - [Exporting certificates on page 643](#)

Configuring the certificate server

The Certificate server performs a number of functions, including providing support for the certificate management daemon (CMD) and for an optional external LDAP server. If the LDAP function is configured, it can be used to automatically retrieve certificates and Certificate Revocation Lists (CRLs) from a Version 2 or Version 3 Lightweight Directory Access Protocol (LDAP) server. The firewall will attempt to retrieve any certificates and (optionally) any CRLs that it needs to validate certificates in a CA-based VPN. Note that the LDAP functionality is used only for non-Netscape Certificate Authorities (for example Entrust and etc.).

Note: In addition to configuring the Certificate server, a root certificate from the Certificate Authority must be imported into the Certificate Authorities tab for a certificate issued by the CA to validate.

To configure the Certificate server, select **Maintenance > Certificate/Key Management**, then click the **Certificate Server** tab. The Certificate Server tab appears.

Figure 367 Certificate Management: Certificate Server tab

The screenshot shows the 'Certificate Server' tab in a configuration window. The tab bar at the top includes: Remote Certificates, Firewall Certificates, Certificate Authorities, Remote Identities, SSL Certificates, Certificate Server (selected), and SSH Keys. The main configuration area contains the following settings:

- ☐ Use LDAP to search for Certificates and CRLs (with a 'Restart Server' button to the right)
- LDAP Server Address: 0.0.0.0
- LDAP Server Port: 389
- LDAP Timeout: 60 (seconds)
- Maximum Validated Key Cache Size: 100 (keys)
- Certificate Key Cache Lifetime: 6 (hr) : 0 (min)
- ☒ Perform CRL Checking
- CRL Retrieval Interval for CA's: 15 minutes
- Audit Level: Normal

Use the Certificate Server tab to configure the Certificate Server.

To configure the Certificate Server tab:

- 1** To enable the LDAP feature, select the **Use LDAP to search for Certificates and CRLs** check box, and follow the sub-steps below. If enabled, the firewall will attempt to retrieve the certificates and CRLs it needs from an LDAP server.
 - a** In the **LDAP Server Address** field, type the IP address of the LDAP server.
 - b** In the **LDAP Server Port** field, type the port number on which the LDAP server listens. The port number is typically 389, but the server can be configured to listen on different ports.
 - c** In the **LDAP Timeout** field, specify the maximum time (in seconds) that CMD will wait while performing an LDAP search. The valid range is between 0 and 3600 seconds. The recommend value is between 5 and 300 seconds.
- 2** In the **Maximum Validated Key Cache Size** field, specify the maximum number of validated keys that will be stored in cache memory. Caching validated keys can increase system performance. Valid ranges are 0–500. A value of 0 indicates that no keys will be cached. For most systems, a value of 100 is sufficient.
- 3** In the **Certificate Key Cache Lifetime** field, specify the maximum amount of time a certificate can remain in the validated key cache before it must be re-validated. The valid range is 0–168 hours (1 week). A value of 0 indicates that the certificate keys must be re-validated with each use.
- 4** Select the **Perform CRL Checking** check box to enable CRL checking. If this option is disabled, CRL lists will not be consulted when validating certificates.
- 5** In the **CRL Retrieval Interval for CA's** drop-down list, specify how often a CA is queried in order to retrieve a new CRL.
- 6** In the **Audit Level** drop-down list, select the type of auditing that should be performed on this server. The options are:
 - **Error** – Logs only major errors.
 - **Normal** – (Default) This is the most common setting. It outputs major errors and informational messages.
 - **Verbose** – Use this level when initially troubleshooting VPN connectivity problems. This audit output is useful for detecting configuration issues.
 - **Debug** – Logs all errors and informational messages. Also logs debug information.

Note: Only use **Debug** and **Error** if you are an experienced administrator or under guidance from Forcepoint support. In particular, debug can overflow your audit logs if left on for an extended period of time.
- 7** Save your changes.

Configuring and displaying remote identities

Remote Identities can be created for two purposes.

- If you choose to have a Certificate Authority policy defined for a VPN (whereby a group of trusted CAs is authorized to issue certificates for access to the VPN), you will also require a list of Remote Identities. Remote Identities are used as part of a VPN definition to determine which remote certificates from a CA may be used to authenticate to a VPN.
- You may also be required to configure a remote identity to be used in a VPN definition for a software client, such as the SafeNet SoftRemote client, using pre-shared passwords.

Select **Maintenance > Certificate/Key Management**, then select the **Remote Identities** tab. The Remote Identities tab appears.

Figure 368 Remote Identities tab

The screenshot shows the 'Remote Identities' tab selected in a web-based management interface. The top navigation bar includes tabs for 'Remote Certificates', 'Firewall Certificates', 'Certificate Authorities', 'Remote Identities' (which is active), 'SSL Certificates', 'Certificate Server', and 'SSH Keys'. The main content area is divided into two sections. On the left, under the heading 'Identities:', there is a large, empty rectangular box for displaying a list of identities. Below this box are two buttons: 'New' and 'Delete'. On the right, there are four input fields for configuring a new identity: 'DN - Distinguished Name:' (a large text area), 'E-Mail Address:' (a single-line text field), 'Domain Name:' (a single-line text field), and 'IP Address:' (a single-line text field).

Use this tab to view and modify the list of available remote identities. Remote identities are used to identify the authorized users who take part in a VPN definition and either have been issued a certificate from a particular CA or use a VPN client configured with a pre-shared password. For example, as part of a remote identity you might define a Distinguished Name that authorizes only people from the Sales department of Bizco corporation.


You can perform the following actions:

- To display the properties of a specific identity, select the identity from within the list. Its properties are displayed on the right portion of the window.
- To modify an identity, make the desired changes and click the **Save** icon. For specific information on modifying the properties that appear for a remote identity, see [About the Create New Remote Identity window on page 637](#).
- To create a new remote identity, click **New**. The Create New Remote Identity window appears. See [About the Create New Remote Identity window on page 637](#) for details.
- To delete an existing identity, highlight the identity you want to delete and click **Delete**.

About the Create New Remote Identity window

The Create New Remote Identity window enables you to add a new remote identity. You can also modify an existing remote identity within the Remote Identities tab.

Figure 369 Remote Identities: Create New Remote Identity window



Tip: An asterisk can be used as a wildcard when defining the fields on this window. (Other special characters are not allowed.) For example; *, O=bizco, C=us represents all users at Bizco.

To add or modify a remote identity:

- 1 In the **Identity Name** field, type a name for this Remote Identity.
- 2 In the **Distinguished Name** field, create a distinguished name. See [Understanding Distinguished Name syntax on page 619](#) for information on the format that should be used.

Note: The order of the specified distinguished name fields must match the order listed in the certificate.

- 3 [Optional] In the **E-Mail Address** field, enter the e-mail address(es) to which you want to restrict access. Enter one e-mail address per identity or use a wildcard to indicate all e-mail addresses, such as **@example.com*.
- 4 [Optional] In the **Domain Name** field, type the specific domain name to which you want to restrict access. Enter one domain name per identity or use a wildcard to indicate all domain names, such as **.example.com*.
- 5 [Optional] In the **IP Address** field, type the unique IP address or group of IP addresses to which you want to restrict access. For example: 182.19.0.0/16 indicates that only users with IP addresses beginning with 182.19 (as contained in the certificate) will be authorized to use the VPN.
- 6 Click **Add** to add the identity to the Identities list.
- 7 Save your changes.

Configuring and displaying remote certificates

A remote certificate identifies one or more peers that can be involved in a VPN connection with a Sidewinder. The firewall can import existing certificates into its Remote Certificates database, or it can create new remote certificates. In either case, all certificates should be in place before you begin configuring a VPN.

Select **Maintenance > Certificate/Key Management**, then select the **Remote Certificates** tab. The Remote Certificates tab appears.

Figure 370 Remote Certificates tab

The screenshot shows the 'Remote Certificates' tab selected in a web interface. At the top, there is a navigation bar with tabs: 'Remote Certificates', 'Firewall Certificates', 'Certificate Authorities', 'Remote Identities', 'SSL Certificates', 'Certificate Server', and 'SSH Keys'. The 'Remote Certificates' tab is active. Below the navigation bar, the interface is divided into two main sections. On the left, under the heading 'Certificates:', there is a large, empty rectangular box for displaying a list of certificates. Below this box are two buttons: 'New' and 'Delete'. At the bottom of the left section are three buttons: 'Import ...', 'Export ...', and 'Query'. On the right side of the interface, there are several input fields and labels. The first is 'DN - Distinguished Name:' followed by a large text area with a vertical scrollbar. Below this are three single-line text input fields labeled 'E-Mail Address:', 'Domain Name:', and 'IP Address:'. At the bottom right, there are two labels with values: 'Signature type: N/A' and 'Status: N/A'.

Use the Remote Certificates tab to view the list of available remote certificates. These certificates represent the potential peers with which firewall can establish a VPN connection. To display the properties of a specific certificate, select the certificate from within the list. Its properties are displayed on the right portion of the window. For a description of these properties, see [Creating a new remote certificate on page 639](#).

Note: You cannot modify the properties of a certificate from this window. To modify a certificate you must delete it and then add it back using the new properties.

From this window, you can perform the following actions:

- **Add a new certificate to the Certificate list** – Click **New** and see [Creating a new remote certificate on page 639](#) for details.
- **Delete a certificate from the list** – Select the certificate you want to delete and click **Delete**.
- **Import certificates** – Click **Import** and see [Importing a remote certificate on page 641](#) for details.
- **Export certificates** – Click **Export** and see [Exporting certificates on page 643](#).
- **Query the CA for Certificate status** – If a certificate request has been submitted to be signed by a CA, click the **Query** button to query the CA to see if the certificate is approved. If yes, the Status field will change to **SIGNED** and the approved certificate will be retrieved.

If the certificate request is Manual PKCS10, click the **Load** button to query and retrieve the signed certificate.


Note: By default, Netscape CAs and CAs that support the Simple Certificate Enrollment Protocol (SCEP) are checked every 15 minutes for any certificates waiting to be signed.

Creating a new remote certificate

The Remote Certificates: Create New Certificate window enables you to add a certificate to the Remote Certificate list.

Note: The default certificate key size is 1024 bits. The default lifetime for self-signed certificates created on the firewall is five years.

Figure 371 Remote Certificates: Create New Certificate window



To add a remote certificate:

- 1 In the **Certificate Name** field, type a name for this certificate.
- 2 In the **Distinguished Name** field, create a distinguished name. See [Understanding Distinguished Name syntax on page 619](#) for information on the format that should be used. Note the following:
 - The order of the specified distinguished name fields must match the order listed in the certificate.
 - Some CAs will not support the optional identity types specified in [Step 3](#) through [Step 5](#).
- 3 [Optional] In the **E-Mail Address** field, type the email address associated with this remote certificate.
- 4 [Optional] In the **Domain Name** field, type the domain name associated with this remote certificate.
- 5 [Optional] In the **IP Address** field, type the IP address associated with this remote certificate.
- 6 In the **Submit to CA** drop-down list, select the enrollment method to which the certificate will be submitted for signing. The valid options are:

- **Self Signed:** Indicates the new certificate will be signed by the firewall rather than by a CA.
- **Manual PKCS10:** Indicates the certificate enrollment request will be placed in a PKCS10 envelope and exported to the file designated in the **Generated PKCS10 File** field.
- The name of the CA to which the certificate is submitted for signing. The CA can be either private (one you own and manage) or it can be public (a trusted CA administered elsewhere).

Note: The CA option is only available if a CA is already configured on the Certificate Authorities tab.

7 In the **Signature Type** box, select the encryption format that will be used when signing the certificate. Valid options are **RSA** or **DSA**.

8 [Conditional] In the **Generated PKCS10 File** field, specify the name and location of the file that will contain the signature request, or click **Browse** to browse the network directories for the file location.

This file contains a PKCS10 “envelope” that is used to send a certificate to a CA for signing. This field is available only if **Manual PKCS10** is specified in the **Submit to CA** field.

Note: To create a new file using the **Browse** button, enter the name and extension (allowed file formats are binary or .pem).

9 [Conditional] In the **Format** field, select the appropriate format for your PKCS10 certificate request.

10 [Conditional] In the **SCEP Password** field, type a password for this certificate. You will need this password if you ever need the CA to revoke this certificate. The password may not contain spaces or single quotes. This field is available only if the **Submit to CA** field displays a CA of type SCEP.


11 Click **Add** to add the certificate to the Certificates list.

12 Save your changes.

Importing a remote certificate

To import a certificate to the list of remote certificates defined on the firewall, select **Maintenance > Certificate/Key Management**, then select the **Remote Certificates** tab and click **Import**. The Import Remote Certificate Window appears.

Figure 372 Remote Certificates: Import Certificate window




To import a remote certificate:

- 1 In the **Import source** field, select the source location of the certificate:
 - **File:** Indicates you will manually specify the location of the certificate file.
 - **Encrypted File:** Indicates you will manually specify the locations of the certificate and private key file.
 - **LDAP:** Indicates that you will access the services of an LDAP (Lightweight Directory Access Protocol) directory to locate the certificate. The LDAP server can be version 2 or version 3.
 - **Paste PEM Certificate:** Indicates you will import the certificate by performing a cut and paste. The Distinguished Name field will change to become the **Manual (pasted) PEM Certificate** field. Paste the certificate into this area.
- 2 In the **Certificate Name** field, type a local name for the certificate you are importing.
- 3 [Conditional] In the **Import Certificate From File** field, type the name and location of the certificate file you will import, or click **Browse** to browse the network directories for the location. (This field is available only if the **Import source** field displays **File**.)
- 4 [Conditional] In the **Password** field, enter the password to decrypt the imported file. This password must match the password given when the file was encrypted. (This field is only available if the **Import Source** field displays **Encrypted File**.)
- 5 Click **OK** to import the remote certificate.
- 6 Save your changes.

If you created a manual remote certificate, you must retrieve the certificate after it is signed by the CA; the firewall will not retrieve it automatically. For this process, the **Load** button appears when an unsigned requested certificate name is highlighted. Clicking this button will initiate the process to retrieve and import the certificate. The Load Certificates for PKCS 10 Request window appears.

Figure 373 Remote Certificates: Load Certificate for PKCS 10 Request window



Use the Load Certificate for PKCS 10 Request window to load signed certificates. It also functions to query an LDAP server for whether or not a requested certificate is signed.

To load a signed certificate:

- 1** In the **Certificate Source** field, select the source location of the certificate. The following options are available:
 - **File:** Indicates you will manually specify the location of the certificate.
 - **LDAP:** Indicates you will access the services of an LDAP (Lightweight Directory Access Protocol) directory to locate the certificate. The LDAP server can be version 2 or version 3.
 - **Pasted PEM Certificate:** Indicates you will paste or type in the certificate from another source, such as another open application window or personal communication.
- 2** [Conditional] In the **Certificate from File** field, if the certificate source is a file, type the location or **Browse** to the location.
- 3** [Conditional] In the **Manual (pasted) PEM Certificate** field, if the certificate source is a Pasted PEM Certificate, type or paste the certificate in this field.
- 4** Click **OK** to issue a query command for your requested certificate, or click **Cancel** cancel the certificate request.

If you click **OK** and the certificate is available, it will automatically be imported and the status will change to SIGNED.
- 5** Save your changes.

Exporting certificates

Once the certificates have been generated, they need to be exported and transferred to a VPN client such as SafeNet SoftRemote or to another firewall. This section contains the following export procedures:

- **Exporting a Remote Certificate** – You are most likely to export a remote certificate if your users use a VPN client to establish a VPN connection between their computers and the firewall. The VPN client requires the use of a certificate to identify itself during the VPN connection negotiations. It is possible to use the firewall to create a self-signed certificate for the VPN client. Once it is created it may be converted to a new file format and then exported. From there it is imported to the VPN client program.
- **Exporting a Firewall Certificate** – This is used to export the firewall certificate to a remote peer. This allows the remote peer to recognize the firewall. On the remote peer the firewall certificate is imported as a remote certificate.

Note: You can also export CA certificates. See [Exporting certificate authorities](#).

To export a certificate, select **Maintenance > Certificate/Key Management**, then select either the **Remote Certificates** tab or the **Firewall Certificates** tab. Select the certificate you wish to export and click **Export**. The Export Firewall Certificate Window appears.

Note: The tab you select depends upon your reason for exporting the certificate. See the explanation in the previous paragraphs.

Figure 374 Certificate Export window

Use the Certificate Export window to export the selected certificate from the firewall to a separate file and/or to the screen. The certificate can be written to a file on the hard drive of a workstation, or it can be written to a transportable medium such as a diskette. You can export only the certificate, or both the certificate and the private key.

Exporting only the certificate

To export a certificate only:

- 1 Select the **Export Certificate (Typical)** radio button.
- 2 Select the export destination:
 - **Export Certificate To File** – To export the certificate to a file, select this option and proceed to [Step 3](#).
 - **Export Certificate To Screen** – Select this option to export the certificate to the screen.
- 3 [Conditional] If you are exporting the certificate to file, do the following:
 - In the **File** field, type the name and location of the file to which the client (or firewall) certificate will be written. If you want to overwrite an existing file, but you are not certain of the path name or the file name, click **Browse**.
 - In the **Format** field, select the appropriate format for the file.
- 4 Click **OK** to export the certificate to the desired location.

Exporting both the certificate and private key

To export both a certificate and private key:

- 1 Specify whether the certificate and private key will be exported as one file or two files by selecting one of the following options:
 - **Export Certificate and Private Key as one file (PKCS12)** – Select this option to export both the certificate and private key as a single file, and proceed to step 2. (This is the preferred method if you are using SoftRemote.)
 - **Export Certificate and Private Key as two files (PKCS1, PKCS8, X.509)** – Select this option to export the certificate and private key as two separate files. Proceed to step 3.
- 2 [Conditional] To export the certificate and private key as a single file, do the following:
 - a In the **File** field, type the name and location of the file to which the client (or firewall) certificate will be written. If you want to overwrite an existing file but you are not certain of the path name or the file name, click **Browse**. (The **Format** displays the file format.)
 - b In the **Password** field, enter the password that will be used to encrypt the certificate file.
 - c In the **Confirm Password** field, re-enter the password that you entered in the **Password** field.
 - d Click **OK** to export the certificate and private key as a single file.
- 3 [Conditional] To export the certificate and private key as two separate files, do the following:
 - a In the **Certificate File** field, type the name and location of the file to which the client or firewall certificate will be written. If you want to overwrite an existing file but you are not certain of the path name or the file name, click **Browse**. In the **Format** field, select the appropriate format for the file.
 - b In the **Private Key File** field, type the name and location of the file to which the key will be written. If you want to overwrite an existing file but you are not certain of the path name or the file name, click **Browse**. In the **Format** field, select the appropriate format for the file.

Note: If you use a transportable medium to store the private key file (for example .pk1, .pk8, or pk12), the medium should be destroyed or reformatted after the private key information has been imported to the appropriate VPN client.
 - c Click **OK** to export the certificate and private key as separate files.


Managing SSH keys

You can create, import, or export SSH host keys that the SSH Proxy present to SSH clients.

- To configure which key(s) the SSH proxy presents to SSH clients, configure the Client Advanced tab on the appropriate SSH Application Defense.
- You can also manage the SSH host keys on the SSH Proxy agent properties window.

To manage SSH host keys, select **Maintenance > Certificate/Key Management**, then select the **SSH Keys** tab.


Figure 375 SSH Keys tab



Name	SSH Fingerprint	Signature type
Default_DSA_Key	00:ba:b5:89:60:27:5e:24:1e:77:ca:78:c6:59:27:ae	dsa
Default_RSA_Key	bd:d5:3d:b0:40:05:7b:18:cb:26:13:b6:67:12:51:bb	rsa

The table lists the available SSH host keys. The SSH Fingerprint is a hashed (shortened) version of the host key to make it easier for users to compare keys.

Figure 376 SSH Keys tab toolbar



Use the toolbar to perform these actions:


Table 71 SSH Keys tab toolbar

Icon	Action
New	Create a new SSH key by clicking New and entering the key's properties in the pop-up window.
Delete	Delete an SSH key by selecting a key from the table and clicking Delete . Note: Either the Default_DSA_Key or the Default_RSA_Key must exist in order to create a new SSH Application Defense.
Import SSH Key	Import a host key generated on another device by clicking Import SSH Key and then entering the key properties in the pop-up window.
Export SSH Key	Export a host key from the firewall by selecting a host key in the table and clicking Export SSH Key and then entering the key destination in the pop-up window.

About the Create New SSH Key window

Use this window to create an SSH host key that the SSH Proxy presents to the SSH clients.

Figure 377 Create New SSH Key window



- 1 In the **Key name** field, enter a name to identify this SSH key.
- 2 From the **Key length** drop-down list, select a key length.


Note: Longer keys are more secure, but they can take longer to generate or they can slow down the encrypted communications.

- 3 Select a signature type.
- 4 Click **Add** and save your changes.

About the Import SSH Key window

Use this window to import a host key that has been generated on another device.

Figure 378 Import SSH Key window

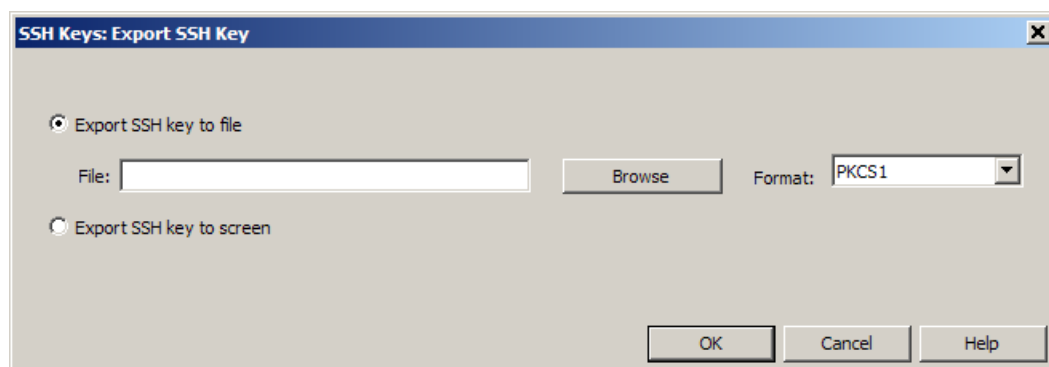


- 1 From the **Import source** drop-down list, select the source location of the host key.
 - **File** – Select this to manually specify the location of the host key.
 - **Cut and paste** – Select this to import the host key by performing a cut and paste. Paste the key in the **Paste SSH key here** box.
- 2 In the **SSH key name** field, type a local name for the host key you are importing.
- 3 [File source only] In the **Import SSH key from file** field, type the name and location of the host key you will import, or click **Browse** to browse the network directories for the location.
- 4 [Cut and paste source only] In the **Paste SSH key here** box, paste the host key that you copied from the source.
- 5 Click **OK** to import the host key.
- 6 Save your changes.

About the Export SSH Key window

Use this window to export the selected host key and import it manually on an SSH client.

Figure 379 Export SSH Key window



1 Select the export destination:

- **Export SSH key to file** – Select this option to export the host key to a file.
- **Export Certificate To Screen** – Select this option to export the host key to a Certificate Data window. You can then cut and paste the host key in the client.

2 [Export to file only] Enter the file information.

- In the **File** field, type the name and location of the file that the host key will be written to. If you want to overwrite an existing file, but you are not certain of the path name or the file name, click **Browse**.
- From the **Format** drop-down list, select the appropriate format for the file.

3 Click **OK** to export the host key to the desired location.

23 High Availability

Contents

- [How Sidewinder High Availability works](#)
- [About HA configuration options](#)
- [Configuring HA](#)
- [Understanding the HA cluster tree structure](#)
- [Managing an HA cluster](#)


How Sidewinder High Availability works

Two firewalls can be configured to work together to provide load-sharing or redundancy. This configuration is known as a *High Availability (HA) cluster*. HA clusters can be configured in either of these ways:

- Load-sharing** – Both the *primary* and *secondary* firewall actively process traffic, providing improved performance and redundancy.
- Failover** – The *standby* firewall does not process traffic unless called on to take over if the primary becomes unavailable, providing redundancy.

As shown in [Figure 380](#), configuring an HA cluster requires at least three burbs for each firewall: an internal burb, an external burb, and a heartbeat burb. The heartbeat burb isolates all HA cluster-specific traffic between the cluster firewalls so that it does not impact regular network traffic.

Figure 380 Basic HA configuration



To implement an HA cluster in your network, cluster firewalls must have interfaces that reside in the same networks and you need one additional *shared cluster* address for each network. This address represents the HA cluster rather than an individual firewall interface. The table below summarizes the IP addresses needed for this HA configuration.

Table 72 HA IP addresses

	internal burb	external burb	heartbeat burb
primary IP	10.1.1.1	1.1.1.1	192.168.1.1
secondary/standby IP	10.1.1.3	1.1.1.3	192.168.1.3
shared cluster address	10.1.1.5	1.1.1.5	192.168.1.5

*. In a load-sharing HA cluster, the shared cluster addresses (except the heartbeat) are shared between Sidewinders. In a failover HA cluster, they are assigned to the primary firewall.

In this example, all users in the internal or external network must use the cluster address (10.1.1.5 or 1.1.1.5, respectively) as the network gateway. Only system administrators should know about the other IP addresses. The same concept applies for DNS host names.

Tip: When configuring an existing single Sidewinder configuration to become an HA cluster, consider using the existing interface addresses as the cluster addresses and using new IP addresses for the actual NICs. This makes the transition from a single firewall to an HA cluster transparent to the rest of your network.

About HA redundancy

The two HA firewalls communicate on the heartbeat burb. An IPsec-authenticated heartbeat is sent by the primary and acknowledged by the secondary/standby. Failures are handled as follows:

- **Load-sharing HA** – If the primary or the secondary firewall becomes unavailable (that is, a heartbeat message or acknowledgement is not received from that firewall for the specified amount of time), the remaining firewall takes over and assumes responsibility for processing all traffic.
- **Failover HA** – If the standby determines that the primary is unavailable, the standby takes over and assumes the role of the primary and the responsibility for processing all traffic.

If one of the firewalls unexpectedly becomes unavailable and the remaining firewall takes over processing all traffic, any active proxy sessions and non-stateful IP packet filter sessions assigned to the unavailable firewall are lost. However, IP packet filter sessions that are configured for stateful session failover are preserved.

About shared cluster addresses

HA clusters use shared IP addresses to receive and transmit network traffic—to other hosts on the network, the HA cluster appears to function as a single firewall. Alias IP addresses added to an HA cluster also function this way.

Note: The individual IP addresses of the firewalls are used for administration purposes only and should not be used to pass traffic.

About HA configuration options

You can configure HA clusters in either of these ways:

- **Load-sharing** – Both the *primary* and *secondary* firewall actively process traffic. See [Load-sharing HA](#).
- **Failover** – The standby does not process traffic unless called on to take over if the primary becomes unavailable. See [Failover HA](#).

Load-sharing HA

Load-sharing HA, also referred to as *active-active HA*, consists of two firewalls that actively process traffic in a load-sharing capacity. When a secondary is registered to an HA cluster, synchronized areas are overwritten to match the primary. (To determine which areas are synchronized, see [Managing an HA cluster](#).)

Note: To configure load-sharing HA, both firewalls must have the same hardware configuration (e.g., CPU speed, memory, active NICs).

Load-sharing HA provides the following benefits:

- **Performance** – Both firewalls actively participate in passing network traffic, which increases performance, especially in heavier workloads.
- **Redundancy** – Load-sharing HA provides the same redundancy benefits as failover HA—if one firewall encounters a failure, the other firewall assumes responsibility for all traffic.
- **Outage transparency** – Load-sharing HA masks outages from your users:
 - If a failover event occurs on one of the firewalls, the outage does not affect the active connections being handled by the operational firewall.
 - If you need to shut down a firewall, you can schedule a soft shutdown, which reduces the number of sessions that are lost.

How load-sharing HA works

Each network interface on both firewalls maintains an individual IP address, the shared cluster address, and any configured aliases belonging to the same network, with one exception: only the primary is assigned the cluster address for the heartbeat burb. The firewalls share the cluster IP addresses using one of several configurable layer 2 modes, allowing them both to receive all traffic sent to the shared cluster addresses. The communication to coordinate load-sharing passes between firewalls on the heartbeat burb.

The firewalls determine how to balance the load by examining the source port of incoming connections—one firewall processes the incoming connections with even source ports while the other firewall processes the incoming connections with odd source ports. However, the following types of traffic are always processed by the primary firewall:

- All management traffic, such as Admin Console, SSH, Telnet, and SNMP
- VPN
- ICMP
- Any other type of traffic that does not include a source port

Once a connection is processed by one of the cluster firewalls, all of the packets associated with that connection are handled by the same firewall. Connections that are specifically addressed to an individual firewall address will be assigned to the specified firewall.

To share the cluster IP addresses, firewalls that are configured for load-sharing HA use a common cluster MAC address for each cluster interface. To accommodate varying switch capabilities, load-sharing clusters can be configured to use one of the following layer 2 modes:

Note: For more information on the available layer 2 modes and the configuration requirements these modes present for your switches and/or routers, see KB article [KB8877](#).

- **Unicast - mirrored** – In this mode, each cluster interface on both firewalls is assigned a shared unicast MAC address. Because a unicast MAC address normally corresponds to a single host, the cluster firewalls rely on switches to forward traffic destined for a cluster MAC address to both cluster firewalls. To accomplish this, each switch that is connected to the cluster firewalls must be configured to send traffic destined for the unicast cluster MAC address to both firewall interfaces.
- **Multicast** – In this mode, each cluster interface on both firewalls is assigned a shared multicast MAC address. If you select this mode and your routers have difficulty processing ARP replies that contain multicast MAC addresses, add static ARP entries on the routers.

Note: This mode uses layer 2 multicast which should not be confused with IP multicast (which functions at layer 3).

- **Unicast - flooded** – In this mode, each cluster interface on both firewalls is assigned a shared unicast MAC address. This mode differs from Unicast - mirrored mode because it does not require any special configuration on your switches. Instead, the firewalls prevent each switch from establishing an association between the cluster MAC address and the switch port on which it can be reached. This causes each switch to send all Ethernet frames destined for the cluster MAC address out all of its ports, allowing both firewalls to receive Ethernet frames destined for the cluster MAC address.

Note: This mode is useful if the switches that are connected to the cluster firewalls do not support Multicast mode or Unicast - mirrored mode. However, this mode increases network overhead for all devices that are connected to the switch.

Soft shutdown

If you know in advance that a firewall will need to be shut down, you can reduce the number of lost connections by scheduling the shutdown (rather than shutting down immediately). When a shutdown is scheduled for a later time, a soft shutdown will be performed to reduce the number of sessions that are lost. For information on soft shutdown, see [Scheduling a soft shutdown for a load-sharing HA cluster Sidewinder](#).

Failover HA

Failover HA consists of one firewall (the primary) actively processing traffic, with the standby acting as a hot backup. When a standby firewall is registered to an HA cluster, synchronized areas are overwritten by the HA cluster configuration. (To determine which areas are synchronized, see [Managing an HA cluster](#).)

You can configure failover HA in one of two ways:

- **peer-to-peer** – In a peer-to-peer HA cluster, both firewalls are configured as standbys with the same takeover time. The first firewall to come online becomes the primary. If the primary becomes unavailable, the peer, currently acting as the standby, takes over as the primary. This firewall remains the primary until it becomes unavailable, at which time the peer takes over as the acting primary.

This is the recommended failover HA configuration. However, to configure peer-to-peer HA, both firewalls must have the same hardware configuration.

- **primary/standby** – In a primary/standby HA cluster, one firewall is designated as the primary and always acts as the primary when it is available. The standby firewall takes over as the acting primary only if the designated primary becomes unavailable. When the primary firewall recovers and becomes available, it resumes its role and another takeover event occurs. This additional takeover event does not occur in a peer-to-peer configuration.

Use this option if you have firewalls that do not share the same hardware configuration, configuring the firewall with higher performing hardware as the primary.

Note: When a takeover event occurs, a number of netprobe events can be detected when connections take time to detect the switch of systems.

Failover HA provides the following benefits:

- Redundancy – The standby firewall acts as a hot standby, ready to take over if the primary becomes unavailable.
- Simple configuration – Failover HA requires less configuration on your switches and/or routers than load-sharing HA.
- Flexible hardware requirements – The requirements for failover HA are less stringent than load-sharing HA.
 - If your firewalls have similar hardware configurations but do not meet the requirements for load-sharing, you can configure them for peer-to-peer HA.
 - If your firewalls have mismatched hardware, you can configure them for primary/standby HA.

How Failover HA works

When the primary is brought online, it activates its individual interface addresses, the cluster addresses, and any aliases assigned to the cluster. When the standby is brought online, it only activates its individual interface IP addresses.

If the standby does not receive a heartbeat signal for a number of seconds (based on the takeover setting of the standby), it activates the shared cluster addresses on its interfaces and begins processing network traffic. In the process, the standby clears its address resolution protocol (ARP) cache and attempts to generate a *gratuitous ARP*. Most systems will immediately determine that the standby is now responsible for the addresses by which the primary is known (the cluster addresses), and new connections will be established through the new acting primary. However, there may be a number of reasons why the gratuitous ARP is not received: a remote system may not recognize the message, the message may be blocked by certain switches, it may fail due to timing issues, etc. Often this can be resolved by:

- Flushing the ARP cache on the remote system.
- Shortening the time that entries stay in the remote system's ARP cache to three to five minutes.
- Configuring systems to communicate with the new ARP address by selecting the **Force ARP Reset** option on the High Availability Advanced Network Properties window when creating an HA cluster.

Configuring HA

This section provides the basic information you need to configure an HA cluster. Before you begin, sketch a diagram showing your planned configuration for reference. Include the following items on your diagram:

- Interfaces
- IP addresses
- HA shared cluster addresses
- Subnet names

Perform these procedures to create an HA cluster:

- [Ensure HA requirements are met](#)
- [Configure the heartbeat interfaces](#)
- [Add the first Sidewinder to a new HA cluster](#)
- [Add a reservation for the second firewall in the HA cluster](#)
- [Join a Sidewinder to an existing HA cluster](#)
- [Post-configuration tasks](#)

Note: A configuration backup is automatically performed and stored on the Sidewinder when creating an HA cluster.

Ensure HA requirements are met

Before you configure HA, the following conditions must be met:

- Both firewalls must be at the same version.
- A dedicated heartbeat burb and interface must be configured on each firewall.

The heartbeat burbs of the HA pair must be directly connected with the appropriate cable:

- 100baseT NIC – Use a crossover cable.
- 1000baseTX NIC – Use a standard Cat5e or Cat6 cable.
- The following areas *must* be configured identically on both firewalls before you configure HA:
 - Number and types of interfaces
 - Number of burbs
 - Burb names (burb names are case-sensitive)
 - Burb creation order

Note: You can verify the order in which burbs were created by running the `region` command at the command line. The output from this command must match on both firewalls to configure HA.

Additional requirements for load-sharing clusters

Before you configure load-sharing HA, the following additional requirements must be met:

- The Sidewinder appliances must have identical hardware configuration.
- The interface used for the heartbeat burb must be at least as fast as the fastest load-sharing interfaces on your firewall. See [Configure the heartbeat interfaces](#) for more information.
- The switches connected to your firewalls must meet certain requirements depending on the layer 2 mode you configure for the cluster. See [How load-sharing HA works](#).
- The unicast - mirrored and unicast - flooded layer 2 modes are only supported on em NICs.
- If VLAN interfaces that share the same parent NIC or NIC group are configured to use either the Unicast - mirrored or Unicast - flooded layer 2 modes, they must meet the following requirements:
 - They must share the same cluster MAC address
 - They must use the same layer 2 mode (Unicast - mirrored or Unicast - flooded)

Note: A load-sharing HA cluster enforces these requirements, keeping the cluster MAC address and layer 2 mode of the appropriate VLAN interfaces synchronized.

Configure the heartbeat interfaces

You must configure a dedicated heartbeat burb and interface on each firewall *before* configuring an HA cluster. To configure the heartbeat interfaces, follow the steps below.

- 1 Ensure that each firewall has an interface that can be dedicated to HA traffic.

Note: Do not use a VLAN for the heartbeat burb.

- 2 In the Admin Console, connect to one of the firewalls and create a heartbeat burb.
 - a Select **Network > Burb Configuration**.
 - b Click **New**.
 - c Type a name and optional description for the heartbeat burb.
 - d Select the **Respond to ICMP echo and timestamp** check box.

Note: Do not select **Hide port unreachables** for a heartbeat burb.
 - e Click **OK** and save your changes.
- 3 Go to **Network > Interfaces** and select the heartbeat burb for an existing interface or create a new interface that includes the heartbeat burb.

See [Configuring interfaces on page 397](#) for detailed information on configuring an interface.
- 4 Save your changes.
- 5 Repeat these steps [Step 1](#) through [Step 4](#) for the other firewall that will be participating in the HA cluster.
- 6 Connect the heartbeat burbs with the appropriate cable.
 - 100baseT NIC – Use a crossover cable.
 - 1000baseTX NIC – Use a standard Cat5e or Cat6 cable.
- 7 Test the network connectivity between the two firewalls for the heartbeat interface.

Note: Network connectivity must exist between the heartbeat burbs to successfully configure HA.

Add the first Sidewinder to a new HA cluster

To add the first firewall to a new HA cluster:

- 1 In the Admin Console, connect to the firewall that will become the primary.
- 2 Verify that you have a dedicated heartbeat burb and interface configured for HA on this firewall. See [Configure the heartbeat interfaces](#) for instructions.
- 3 Select **Maintenance > Cluster Wizard**.
- 4 Click **Launch Cluster Wizard**. The Cluster Wizard window appears.
- 5 Click **Next**. The Create New or Join Existing Cluster window appears. Continue with [Using the Create New or Join Existing Cluster window](#).

Using the Create New or Join Existing Cluster window

Select **Create New Cluster** and then click **Next**. The Determining High Availability mode window appears. Continue with [Using the Determining High Availability mode window](#).

Using the Determining High Availability mode window

Select one of the following HA modes and then click **Next**. Continue with [Using the High Availability Takeover Time window](#).

- **Peer-To-Peer HA** – Select this button to create a peer-to-peer HA cluster. In this configuration, the first firewall that comes online first becomes the primary.

See [Failover HA](#) for more information.
- **Load-Sharing HA** – Select this button to create a load-sharing HA cluster. Load-sharing HA consists of two firewalls that actively process traffic in a load-sharing capacity. This is the recommended configuration.

See [Load-sharing HA](#) for more information.
- **Primary/Standby HA** – Select this button to create a primary/standby HA cluster. In this configuration, one firewall is designated as the primary. Use this option if you have firewalls that do not share the same hardware configuration.

See [Failover HA](#) for more information.

Note: To configure peer-to-peer HA or load-sharing HA, both firewalls must have the same hardware configuration.

Using the High Availability Takeover Time window

Note: This window does not appear if you selected the primary/standby HA option. For primary/standby HA, the takeover time is 3 seconds for the primary and 13 seconds for the standby by default and cannot be modified in the Cluster Wizard.

- 1 In the **Takeover Time** field, specify the number of seconds that the primary must be unavailable before the secondary/standby will begin the takeover process. The default value is 13 seconds.
- 2 Click **Next**.
 - If you are creating a load-sharing HA cluster, the High Availability Layer 2 Mode window appears. Continue with [Using the High Availability Layer 2 Mode window](#).
 - If you are creating a peer-to-peer or primary/standby HA cluster, the High Availability Shared Cluster Addresses window appears. Continue with [Using the High Availability Shared Cluster Addresses window](#).

Using the High Availability Layer 2 Mode window

Note: This window only appears if you are creating a load-sharing cluster.

Use this window to select the layer 2 mode the firewalls will use to share the cluster IP addresses, and then click **Next**. The Availability Shared Cluster Addresses window appears. Continue with [Using the High Availability Shared Cluster Addresses window](#).

- **Unicast - mirrored** – Select this mode if the switches that are connected to your firewalls can be configured to send traffic destined for single unicast MAC addresses out multiple ports.
- **Multicast** – Select this mode if the switches that are connected to your firewalls do not support Unicast - mirrored mode but do support multicast MAC addresses.
- **Unicast - flooded** – Select this mode if the switches that are connected to your firewalls do not support Multicast mode or Unicast - mirrored mode.

For more information on the available layer 2 modes and the configuration requirements these modes present for your switches and/or routers, see:

- [How load-sharing HA works](#)
- KB article [KB8877](#).

Using the High Availability Shared Cluster Addresses window

Use the High Availability Shared Cluster Addresses window to configure the shared cluster addresses for the interfaces in your HA cluster. You also specify the heartbeat burb, which is responsible for sending and receiving heartbeats.

Figure 381 High Availability Shared Cluster Addresses window

Shared Cluster IP Address	Network Address	Burb
	10.65.248.0	internal
	192.168.1.0	heartbeat
	10.65.249.0	external

Configure

Heartbeat Burb: -- Select One -- (Burb used for intra-cluster communication)

☒ Use default advanced High Availability properties and skip advanced screens

Following the completion of this wizard, any additional changes to High Availability properties may require a reboot.

Help < Back Next > Cancel

Do the following, and then click **Next**:

- 1 Select the interface row that you want to configure, and click **Configure**. The High Availability Aliases window appears.
- 2 In the **Shared cluster IP address** field, type the shared cluster IP address for the interface that will be shared between firewalls within the HA cluster. To find the IP address for a host name, type the name and click **DNS Lookup**.
 - Three interfaces in each HA cluster must have a cluster IP address. If an additional interface is used for private or management purposes or is not yet ready to be shared, you do not have to assign a cluster IP address.
 - The cluster address is the address most systems should use to communicate with or through the firewall, meaning that DNS, default routes, etc., need to be aware of this address.
- 3 Click **OK**.
- 4 Repeat [Step 1](#) through [Step 3](#) for each interface that will use HA.
- 5 From the **Heartbeat Burb** drop-down list, select the burb that HA will use to send or receive heartbeats. This must be a dedicated burb.

6 Determine if you want to skip the advanced configuration windows and use the default values.

- To skip the advanced configuration windows and use the default values, select the **Use default advanced High Availability properties and skip advanced screens** check box, click **Next**. The Cluster Wizard Summary window appears. Continue with [About the Success window](#).

If you skip the advanced configuration windows, the following configuration options will be made automatically:

- IPsec authentication password and authentication type will be automatically selected.
- HA identification cluster ID and multicast address will be automatically assigned.
- Remote test configuration options will not be configured.
- [Load-sharing HA only] The layer 2 mode you selected on the Determining High Availability mode window will be applied to all interfaces (the layer 2 mode can be configured on a per interface basis).
- To modify or configure any of these properties, clear the **Use default advanced High Availability properties and skip advanced screens** check box and click **Next** to access the Advanced General Properties and Advanced Network Properties windows. Continue with [Using the Advanced General Properties window](#).

Using the Advanced General Properties window

Use the High Availability Advanced General Properties window to configure IPsec Authentication values and High Availability identification values.

Note: This window does not appear if you selected the **Use default advanced High Availability properties and skip advanced screens** check box in the High Availability Common Addresses window.

Modify any of the following values:

- **High Availability Password** – The password used to generate the authentication key for IPsec is created automatically. This password must be the same for both firewalls because they share the same virtual firewall ID. You should not need to change this password.
- **Authentication Type** – Select one of the following:
 - **SHA1** – Select this option if using HMAC-SHA1 authentication.
 - **MD5** – Select this option if using HMAC-MD5 authentication.
- **Cluster ID** – The default value does not require modification; valid values are 1–255. If modified, each firewall within an HA cluster must be assigned the same cluster ID.
- **Multicast Address** – This field displays the address of the multicast group used for HA purposes in the heartbeat burb. The default address is 239.255.0.1. To modify the address, click **Edit Address**.

When you have finished configuring this window, click **Next**. The Advanced Network Properties Top Window appears. Continue with [Using the Advanced Network Properties Top window](#).

Using the Advanced Network Properties Top window

Use the High Availability Advanced Network Properties Top window to configure cluster interface-specific properties.

Note: This window does not appear if you selected the **Use default advanced High Availability properties and skip advanced screens** check box in the High Availability Shared Cluster Addresses window.

- 1 From the cluster interfaces list, select the interface that you want to modify.
- 2 Click **Configure**. The High Availability Advanced Properties Bottom window appears. Continue with [Using the Advanced Network Properties Bottom window](#).

Using the Advanced Network Properties Bottom window

Use the Advanced Network Properties Bottom window to modify the cluster IP address, configure interface testing, and force ARP reset properties for the selected cluster interface.

Figure 382 High Availability Advanced Properties Bottom window on a load-sharing cluster

Common Parameters: High Availability Advanced Properties Bottom Window

Cluster IP: 10.65.248.20 L2 Mode: Multicast

Burb: internal Cluster MAC: 01:00:5e:41:f8:14

Network: 10.65.248.0

Interface Test (Optional)

The Remote Test IPs will be used to determine network availability. If ALL of the addresses, below, are unreachable, the associated interface will be treated as if it is disconnected.

☒ Monitor link status

IP Address

Ping interval: 30 (seconds) Failures allowed: 3

Force ARP Reset (Optional)

ARP Reset IP addresses are not configurable on a High Availability Load-Sharing cluster.

IP Address

OK Cancel Help

Perform any of the following actions for this interface:

- Change the cluster IP address – To modify the cluster IP address, edit the **Cluster IP** field.

Note: You can only modify the host portion of the cluster IP address.

Example: if the configured cluster IP address is 10.1.1.5/24, you can specify a new IP in the range 10.1.1.1–254.

- [Load-sharing HA only] Change the layer 2 mode – If the switch that is connected to this interface does not support the layer 2 mode you previously configured, select a different mode from the **L2 Mode** drop-down list.
 - **Unicast - mirrored** – Select this mode if the switch that is connected to this interface can be configured to send traffic destined for single unicast MAC addresses out multiple ports.
 - **Multicast** – Select this mode if the switch that is connected to this interface does not support Unicast - mirrored mode but does support multicast MAC addresses.
 - **Multicast no IGMP** – Select this mode if the switch that is connected to this interface supports multicast MAC addresses and you do not want this interface to send IGMP messages advertising the cluster MAC address.

Note: This is the layer 2 mode used by load-sharing HA clusters at version 7.0.0.07 and earlier. Only select this mode if you must do so to preserve compatibility with the switch that is connected to this interface.

- **Unicast - flooded** – Select this mode if the switch that is connected to this interface does not support Multicast mode or Unicast - mirrored mode.

Note: For more information on the available layer 2 modes and the configuration requirements these modes present for your switches and/or routers, see [How load-sharing HA works](#) and KB article [KB8877](#).

- [Load-sharing HA only] Change the cluster MAC address – Edit the **Cluster MAC** field if the automatically generated cluster MAC address conflicts with a device that is attached to the same network or you are instructed to do so by Forcepoint support.

Note: Do not change the first three octets (xx:xx:xx:yy:yy:yy) of the cluster MAC address.

- Enable link monitoring – Configure the HA cluster to test whether the interface link is active by selecting **Monitor link status**. This method checks if the interface is disconnected or the NIC stops working. It does not verify that other devices can be contacted by the firewall.
- Add remote test IP addresses – Configure the HA cluster to send pings from this interface to verify that the firewall can communicate with devices upstream.
 - a In the Interface Test area, click **New**.
 - b Click the **Specify IP Address** field and type an IP address that the selected interface will send pings to.
 - c Ping addresses must be owned by highly reliable systems that are directly attached to the Sidewinder network but do not belong to either cluster member.
 - d In the **Ping interval** field, specify how often (in seconds) the firewall will ping the remote address to ensure that an interface and path are operational.
 - e In the **Failures allowed** field, specify the number of failed ping attempts that must occur before the standby takes over as the primary.
 - For load-sharing HA, if remote ping fails on one of the two cluster members, that member becomes unavailable until the remote interface is again detected. If there is only one active cluster member and a remote ping failure is detected, that member audits the failure and remains in the cluster until another member joins the cluster (without a ping failure), or until the remote system is detected.
 - If the primary becomes unavailable immediately after a ping attempt has been issued, the time it takes for a secondary/standby to take over will be slightly longer (this is because it will take close to an entire test interval before the first failure is detected).

To modify a ping IP address, double-click the address in the list and make the change.

To delete a ping IP address, select it in the list and click **Delete**.

- Disable monitoring for this interface – Clear the **Monitor link status** check box and clear any IP addresses.
- Add devices that ignore gratuitous ARPs – Configure the HA cluster to attempt to force devices that ignore gratuitous ARP requests to update their ARP tables by doing the following:

Note: This area is not available if you are configuring load-sharing HA.

- a In the Force ARP Reset area, click **New**.
- b Click the **Specify IP Address** field and type an IP address that will not respect gratuitous ARP requests.

To modify an entry, double-click the address in the list and make the change.

To delete an entry, select it in the list and click **Delete**.

When you are done, click **Next**. The Cluster Wizard Summary window appears. Continue with [About the Summary window](#).

About the Summary window

The Cluster Wizard Summary window displays a list of the actions that will be performed when you click **Execute**.

- If you want to make changes to your configuration before executing, click **Back** to navigate to the appropriate window(s) and make the necessary changes.

Note: Carefully review your changes before you click **Execute**. If you change cluster properties using the High Availability window, you may need to reboot your appliance for the changes to take effect.

- When you are satisfied with the summary of changes, click **Execute** and then click **Yes** to continue. A progress bar will appear while the configuration changes are made. If the transition is successful the Success window appears, displaying the new state.

About the Success window

This window informs you that the change you initiated was successful, and lists the firewall's new state. It informs you if a reboot has been initiated, and alerts you of any changes to how the appliance will be managed.

- 1 To close the Cluster Wizard, click **Finish**. The Admin Console disconnects.

The IP address on the connection window changes to the shared cluster address.

- 2 Connect the Admin Console to the HA cluster using the shared cluster address.

Add a reservation for the second firewall and then register it to the primary. Continue with [Add a reservation for the second firewall in the HA cluster](#).

Add a reservation for the second firewall in the HA cluster

Before joining a Sidewinder to an existing HA cluster, you must make a reservation for that firewall in the Common Parameters tab of the HA cluster. Once you have *added* the firewall to the HA cluster, you will need to *join* the firewall to the HA cluster using the Cluster Wizard.

To add a reservation for the new firewall in the existing HA cluster:

- 1 Connect to the HA cluster IP address using the Admin Console.
- 2 In the Admin Console tree, select **High Availability**. The High Availability Common Parameters tab appears.
- 3 In the **Pair Members** area, click **New**. The Add New Firewall window appears.
- 4 In the **Hostname** field, enter the name of the firewall you are adding to the HA cluster. The name must be a fully qualified host name in the same domain as the primary.
- 5 [Primary/standby only] In the **Takeover Time** field, select the number of seconds that the primary must be unavailable before the secondary/standby will begin the takeover process. The default value is 13 seconds.
- 6 In the **IP Address in Heartbeat Burp** field, enter the individual IP address (in the heartbeat burp) of the firewall that you are adding to the HA cluster.
- 7 In the **Registration Key** field, create the registration key for this HA cluster. The key must be at least one character long and may consist of alphanumeric characters, hyphens (-), and underscores (_).

Note: You will need the registration key when you join the firewall to the HA cluster using the Cluster Wizard.

- 8 Click **Add** to add the firewall to the HA cluster.

You can now join the Sidewinder to the HA cluster using the Cluster Wizard. Continue with [Join a Sidewinder to an existing HA cluster](#).

Join a Sidewinder to an existing HA cluster

To join a firewall to an existing HA cluster:

- 1 Using the Admin Console, connect to the firewall that will be joining the HA cluster.
- 2 Select **Maintenance > Cluster Wizard**.
- 3 Click **Launch Cluster Wizard**. The Welcome window appears.
- 4 Click **Next**.
- 5 Select **Join Existing Cluster** and then click **Next**.
- 6 In the Gathering information to join cluster window, configure the following fields:
 - **Partner's Heartbeat Burp IP Address** – Enter the heartbeat IP address of the HA partner.
Note: This is the individual heartbeat IP address for the HA partner, not the shared cluster address.
 - **Registration Key** – Enter the registration key for the HA cluster (the key that you created when you added this firewall to the HA cluster in the [Add a reservation for the second firewall in the HA cluster](#) procedure).
Note: The **Cluster Member Name** field contains the host name of the firewall that you are joining to the HA cluster. This field is read-only.
- 7 Click **Next**. The Cluster Summary window displays a list of the actions that will be performed.
- 8 Review the Cluster Wizard Summary window.

If you want to make changes to your configuration before executing, click **Back** to navigate to the appropriate window(s) and make the necessary changes.

Note: Carefully review the changes, as changes you make after joining this firewall to the cluster will require an additional reboot.
- 9 When you are satisfied with the summary of changes, click **Execute**, then click **Yes** to confirm.

A progress bar appears while the configuration changes are made. If the transition is successful, the Success window appears displaying the new state.
- 10 Click **Finish**.

The firewall icon you just joined to the cluster disappears from the Admin Console tree.
- 11 On the primary, select the new cluster member in any branch.

"This firewall is not registered" appears in the right pane.
- 12 Click **Check Now**.

The new cluster member is registered as part of the HA cluster and all branch contents are visible.

Post-configuration tasks

After you have configured an HA cluster, you can perform the following tasks:

- [Enabling and disabling load-sharing for an HA cluster](#)
- [Removing a Sidewinder from an HA cluster](#)

Enabling and disabling load-sharing for an HA cluster

If you have an HA cluster configured and want to enable or disable load-sharing, follow the steps below. For more information on load-sharing HA, see [Load-sharing HA](#).

Tip: We recommend scheduling downtime to perform this procedure since it involves restarting both firewalls.

- 1 Connect to the HA cluster and select **High Availability**.
- 2 Click the plus sign (+) in front of the High Availability branch to display the individual icons for each firewall that is part of the HA cluster.
- 3 Select the primary firewall's icon. The Local Parameters tab appears.

Tip: If you do not know which firewall is currently the primary, click **Cluster Status** in the Pair Members area of the High Availability window.

- 4 In the **Cluster Mode** area, enable or disable load-sharing by selecting the appropriate cluster mode as follows:
 - **Designate as part of a Load Sharing High Availability Cluster** – Select this option if you want to enable load-sharing for the HA cluster (both firewalls actively process traffic).
 - **Designate as part of a Primary/Standby High Availability Cluster** – Select this option if you want to disable load-sharing HA and convert the HA cluster to a failover HA cluster (only one firewall processes traffic, with the other firewall acting as a hot backup).
 - **Primary** – Select this option to make the selected icon the primary member of the cluster.
 - **Standby** – Select this option to make the selected icon the standby member of the cluster.
- 5 Save your changes.
- 6 Wait 60 seconds to allow the firewalls to synchronize.
- 7 Restart the cluster firewalls:
 - a Shut down the secondary/standby.
 - b Restart the primary.
 - c When the primary is finished rebooting, start up the secondary/standby.

Note: For more detailed restart instructions see [Restarting an HA cluster](#).

Removing a Sidewinder from an HA cluster

To remove a secondary/standby from an HA cluster, perform the steps below.

Removing a secondary/standby from an HA cluster

- 1 Connect to the HA cluster and select **High Availability** in the Admin Console tree. The Common Parameters window appears.
- 2 In the **Pair Members** table, select the secondary/standby and then click **Delete**.

When the firewall is removed from the HA cluster, it will transition to a standalone state.
- 3 To connect to the removed firewall, you must add it to the Admin Console tree:
 - a From the **File** menu, select **New Firewall**. The Add Firewall window appears.
 - b Enter the firewall name and IP address, then click **Add**.

Removing the primary from an HA cluster

You must remove the secondary/standby from the HA cluster before you can remove the primary from the HA cluster. After you have removed the secondary/standby from an HA cluster, follow the steps below to remove the primary from the HA cluster.

- 1** From the Admin Console, connect to the primary using the cluster address.
- 2** In the Admin Console tree, select **High Availability**. The Common Parameters window appears.
- 3** In the Pair Members area, click **Change State to Standalone**. The Cluster Wizard Welcome window appears.
- 4** Click **Next**.
- 5** Select **Change To Standalone**, and then click **Next**.

The Cluster Wizard Summary window appears.

- 6** Review the summary. When you are satisfied with the summary of changes, click **Execute** and then click **Yes**.

A progress bar appears while the configuration changes are made. If the transition is successful, the Success window appears displaying the new state.

- 7** Click **Finish**. The Admin Console disconnects. The firewall will transition to a standalone state.

Understanding the HA cluster tree structure

The Admin Console tree structure is slightly different for an HA cluster. When you administer an HA cluster, both firewalls are managed within a single Admin Console connection to the cluster IP address.


Areas of the HA cluster that are synchronized (that is, areas in which the information for both firewalls must be the same and remains in sync) will appear with a single tree option. When you modify information within those areas, the information will automatically be updated for both firewalls.

Information specific to individual firewalls within the HA cluster (such as configuration backup and restore) includes a sub-folder (indicated by a plus [+] sign) that contains an icon for each firewall that is part of the HA cluster. To modify information within these areas, expand the tree branch, select the appropriate firewall, and make the desired changes. Non-synchronized modifications to an individual firewall will be applied only to that firewall and will not be overwritten by changes made to the other firewall.

The License window is further split in an HA cluster: The Contact and Company tabs appear when you select **License** in the tree; the Firewall and Enrollment List tabs appear when you select an individual Sidewinder.

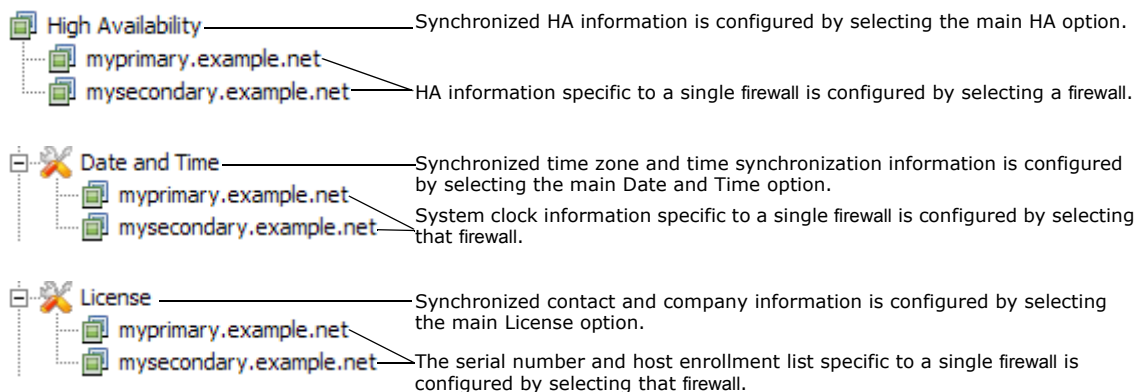
The figure below demonstrates the difference between an individually configured area of the HA cluster (Interfaces) and a synchronized area of the HA cluster (Burb Configuration).

Figure 383 Example of an individually configured area



The **High Availability**, **Date and Time**, and **License** areas within the HA cluster tree include some areas that are synchronized and some areas that are configured on an individual firewall basis, as shown in the figure below.

Figure 384 High Availability, Date and Time, and License areas



The following lists summarize the features that are synchronized and the features that are configured individually in an HA cluster.

• **Features that are synchronized within an HA cluster**

- High Availability
- IPS Attack Responses
- System Responses
- Audit (health monitoring, filters)
- Rules
- Rule Elements
- Application Defenses
- Network Defenses
- Burp Configuration
- VPN Configuration
- DNS
- sendmail
- Routing
- Administrator Accounts
- Interface Aliases

Features that are managed individually within an HA cluster

- Dashboard
- Service Status
- Interfaces
- Configuration Backup
- Date and Time
- License
- Software Management
- System Shutdown
- File Editor

Managing an HA cluster

Once you have configured an HA cluster, the HA cluster will be represented in the Admin Console tree by one combined firewall icon. When you connect to the HA cluster, you will use the HA shared cluster address that you created when you configured HA. This allows you to manage both firewalls by connecting to the HA cluster.

Caution: If you modify your hardware interface configuration, a reboot is required.

Modifying HA common parameters

Use the Common Parameters tab to configure properties that are common to the HA cluster. These parameters affect all firewalls in your HA configuration. These properties were set when you ran the Cluster Wizard (if you skipped the advanced properties windows, many of these properties were set automatically).

Note: If you make any configuration changes on this tab, both cluster firewalls must be restarted. See [Restarting an HA cluster](#).

To configure or modify common HA parameters, connect to the HA cluster using the Admin Console. In the Admin Console tree, select **High Availability**. The Common Parameters tab appears.

Figure 385 Common Parameters tab

The screenshot shows the 'Common Parameters' tab for a High Availability cluster. The left sidebar contains a tree view with 'Firewalls' expanded, showing 'HA' and 'High Availability' (selected). The main area contains the following sections:

- High Availability Identification:** Cluster ID (20), Multicast Group Address (239.255.0.1), Heartbeat Burp (heartbeat), Heartbeat Verification Burp (<None>).
- IPSec Authentication:** Authentication Type (SHA 1), Password (1c0cafd77e).
- Pair Members:** List of members (gabrielle.example.net, xena.example.net) with New, Delete, and Cluster Status buttons.
- Cluster Interfaces:** A table with columns: Cluster IP, Network Address, Burp, Remote Test IPs, and ARP Reset IPs.

Cluster IP	Network Address	Burp	Remote Test IPs	ARP Reset IPs
10.65.248.20	10.65.248.0	internal		
10.65.249.20	10.65.249.0	external		
192.168.1.20	192.168.1.0	heartbeat		

At the bottom, there is a 'Modify' button and a note: 'Interfaces can be added or deleted from the Interface screen.'

To configure or modify the HA common parameters: You can configure or modify the following areas:

High Availability Identification

- **Cluster ID** – The default value does not require modification; valid values are 1–255. If modified, each firewall within an HA cluster must be assigned the same cluster ID.
- **Multicast Group Address** – Displays the address of the multicast group used for HA purposes on the heartbeat burb. The default address is 239.255.0.1. To modify the address, click **Edit address**.
 - Do not specify an address that conflicts with other multicast groups on the heartbeat burb. Addresses in the range of 239.192.0.0 to 239.255.255.255 have been reserved by RFC 2365 for locally administered multicast addresses.
 - Boundary routers should be configured to not pass your selected address if such a feature exists.

Note: If the default is not used, you should change the reverse lookup files in DNS to allow DNS reverse resolution of the multicast address. Refer to the `/etc/namedb.u/failover.rev` file.

- **Heartbeat Burb** – Displays the burb that HA uses to send or receive a heartbeat.

A *heartbeat* is a short message that is sent out at specific intervals to verify whether a Sidewinder is operational. The heartbeat, session information, and configuration information are also transferred between the heartbeat burbs. This must be a dedicated heartbeat burb.

To change the heartbeat burb, select a new one from the drop-down list.

- **Heartbeat Verification Burb** – From the drop-down list, select the burb that HA will use to send or receive a mini-heartbeat. This should be a burb that regularly passes traffic, such as the internal burb.

This mini-heartbeat helps protect against false failover events by doing the following:

- If the firewall does not detect the heartbeat but does detect the mini-heartbeat, the HA cluster does not fail over. An audit message is generated, alerting the administrator to check the heartbeat burbs' connectivity.

Note: The switch must be able to pass IGMP (with an IP protocol number of 2) and must not decrement time to live (TTL) of the HA heartbeat.

Loss of communications on the heartbeat burb causes diminished HA services.

- For load-sharing HA, the active secondary no longer shares the session load; it goes to a standby state.
- For failover HA, the standby cannot receive updated information about new packet filter sessions established on the primary.
- If the firewall does not detect either the heartbeat or the mini-heartbeat, the HA cluster fails over.

Additional information on heartbeat verification is available in the KnowledgeBase articles [8851](#) and [9199](#).

IPSec Authentication

- **Authentication Type** – Select the type of IPSec authentication to use for HA.
 - **SHA1** – Select this option if using HMAC-SHA1 authentication.
 - **MD5** – Select this option if using HMAC-MD5 authentication.
- **Password** – Displays the auto-generated password that is used to generate the authentication key for IPSec. You should not change this password.

Pair Members

The table lists the firewalls that are in the HA cluster.

- If one firewall is in the list, you can perform these actions:
 - Click **New** to add the second firewall to the cluster. This reserves a space in the cluster for the second firewall. You must then join the second firewall to the cluster.

See [Add a reservation for the second firewall in the HA cluster](#) and [Join a Sidewinder to an existing HA cluster](#) for details.
 - Click **Change State to Standalone** to remove the primary firewall from the HA cluster.
- If two firewalls are in the list, you can perform these actions:
 - Select the standby and click **Delete** to remove the standby firewall from the cluster.
 - Click **Cluster Status** to see which firewall is the current primary and which is the secondary.
- **Auto-Recover On Reconnect** – When a monitored interface fails, it triggers a failover event that causes the primary firewall to become a standby firewall. The **Auto-Recover On Reconnect** checkbox controls how the standby firewall returns to the cluster after the connection is restored.

Note: The **Auto-Recover On Reconnect** option only applies to peer-to-peer and primary/standby HA configurations.

The effects of the **Auto-Recover On Reconnect** option are summarized in [Table 73](#).

Table 73 Auto-Recover on Reconnect summary

HA cluster configuration	Auto-Recover on Reconnect checkbox is...	
	Selected	Deselected
Peer-to-peer	The standby firewall is available for failover when the monitored interface reconnects.	An administrator must manually restart the standby node in order to make it available for failover.
Primary/standby	The standby firewall will return as the cluster primary when the monitored interface reconnects.	An administrator must manually restart the standby node in order to restore it to the cluster as the primary firewall.
Load-sharing	No effect	No effect

The remote test IP addresses that monitor interfaces are displayed and configured through the Cluster Interfaces table.

Cluster Interfaces

The table identifies the cluster address, network address, burb, remote test IP address, ARP reset address, and layer 2 mode, and cluster MAC address for each interface.

Note: Layer 2 modes and MAC addresses are listed only for a load-sharing HA. You may need the cluster MAC address to configure your router.

Select an interface and click **Modify** to make the following changes:

- Modify the cluster IP address.

Note: You can also modify the cluster IP address on the Interface Properties window. If you modify the cluster IP address in one window, it is automatically updated in the other window.
- Create, modify, or delete remote test IP addresses.
- Create, modify, or delete force ARP reset IP addresses.

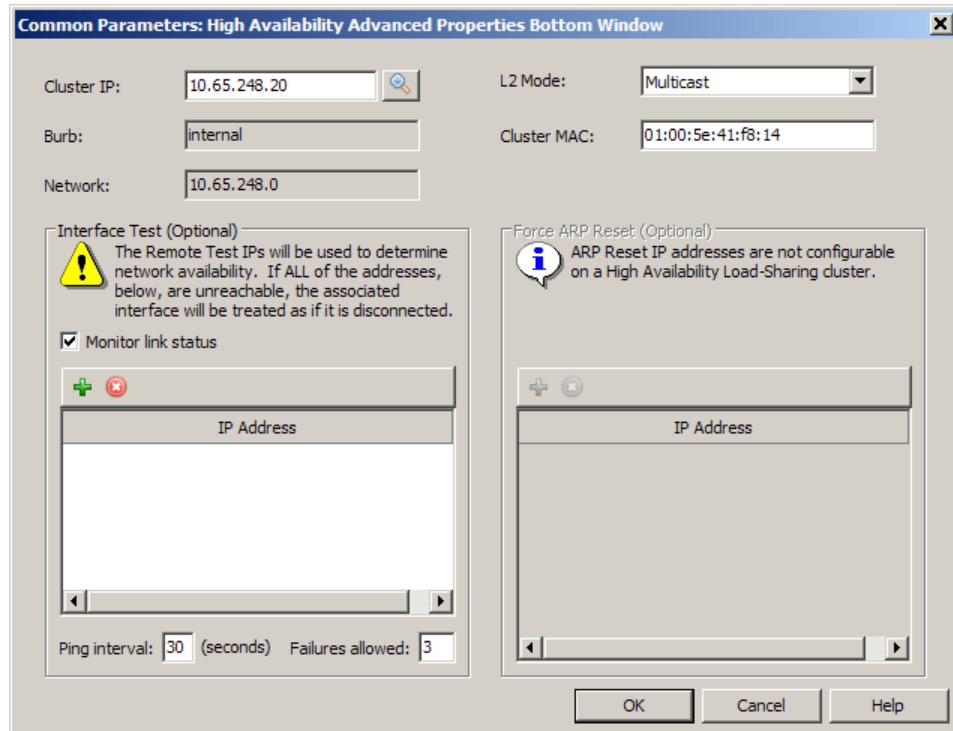
See [Modifying cluster interface properties](#) for more information.

Modifying cluster interface properties

Use the Cluster Interface Properties window to perform the following actions:

- Modify the cluster IP address.
- Create, modify, or delete remote test IP addresses.
- Create, modify, or delete force ARP reset IP addresses.

Figure 386 Cluster interface properties window on a load-sharing cluster



Modify the cluster IP address

In the **Cluster IP** field, make the desired change.

The address should be in the same network as the interface's burb. The network address and burb for the selected interface are listed.

Note: You can also modify the cluster IP address on the Interface Properties window. If you modify the cluster IP address in one window, it is automatically updated in the other window.

[Load-sharing only] Change the layer 2 mode

If the switch that is connected to this interface does not support the layer 2 mode that is currently configured, select a different mode from the **L2 Mode** drop-down list.

- **Unicast - mirrored** – Select this mode if the switch that is connected to this interface can be configured to send traffic destined for single unicast MAC addresses out multiple ports.
- **Multicast** – Select this mode if the switch that is connected to this interface does not support Unicast - mirrored mode but does support multicast MAC addresses.
- **Multicast no IGMP** – Select this mode if the switch that is connected to this interface supports multicast MAC addresses and you do not want this interface to send IGMP messages advertising the cluster MAC address.

Note: This is the layer 2 mode used by load-sharing HA clusters at version 7.0.0.07 and earlier. Only select this mode if you must do so to preserve compatibility with the switch that is connected to this interface.

- **Unicast - flooded** – Select this mode if the switch that is connected to this interface does not support Multicast mode or Unicast - mirrored mode.

For more information on the available layer 2 modes and the configuration requirements these modes present for your switches and/or routers, see:

- [How load-sharing HA works](#)
- KB article [KB8877](#).

[Load-sharing only] Change the cluster MAC address

Change the cluster MAC address by editing the **Cluster MAC** field.

- Do not modify the cluster MAC address unless it conflicts with a device that is attached to the same network or you are instructed to do so by Forcepoint support.
- Do not change the first three octets (**xx:xx:xx:yy:yy:yy**) of the cluster MAC address.

Monitor the interface link

Test whether the interface link is active by selecting **Monitor link status**. This checks if the interface is disconnected or the NIC stops working. It does not verify that other devices can be contacted by the firewall.

Create, modify, or delete remote IP addresses for Interface Test

In the Interface Test area, configure remote test IP addresses for networks that you want to periodically ping.

Ping addresses must be in highly reliable systems that are directly attached to the Sidewinder network.

- 1 Click **New**, then click the **Specify IP Address** field and type an IP address that the firewall will ping.
- 2 In the **Ping interval** field, specify how often (in seconds) the firewall will ping the remote address to ensure that an interface and path are operational.
- 3 In the **Failures allowed** field, specify the number of failed ping attempts that must occur before the standby interface takes over as the primary.

Failures are counted in increments and decrements rather than successively. This means that a failed ping adds to the failure total, and a successful ping subtracts from the failure total. The failure total is never less than zero and it is never more than the configured failures allowed.

For example, if the configured failures allowed is **3**, this is how the failure count is tallied based on the ping results:

Table 74 Example ping results

Ping result:	failure	success	success	failure	failure	success	failure	failure	<i>Failover event occurs</i>
Failure total:	1	0	0	1	2	1	2	3	

- To modify a ping IP address, double-click the address in the list and make the change.
- To delete a ping IP address, select it in the list and click **Delete**.

Note: If the primary becomes unavailable immediately after a ping attempt has been issued, the time it takes for a secondary/standby to take over will be slightly longer (this is because it will take close to an entire test interval before the first failure is detected).

[Failover HA only] Create, modify, or delete IP addresses for Force ARP Reset

In the Force ARP Reset area, configure hosts that are known to ignore gratuitous ARPs, but that need to know the new cluster alias.

Click **New**, then click the **Specify IP Address** field and type an IP address that will not respect gratuitous ARP requests.

- To modify an entry, double-click the address in the list and make the change.
- To delete an entry, select it in the list and click **Delete**.

Modifying HA local parameters

To configure local HA parameters, connect to the cluster IP using the Admin Console and select **High Availability > firewall name**. The following window appears:

Note: If you make any configuration changes on this tab, both cluster firewalls must be restarted. See [Restarting an HA cluster](#).

Figure 387 Local Parameters tab

Local Parameters

Cluster Mode:

☐ Designate as part of a Load Sharing High Availability Cluster

☒ Designate as part of a Primary/Standby High Availability Cluster:

☐ Primary

☒ Standby

Control:

☒ Enabled

☐ Disabled

Takeover Time:

13 Seconds

The Local Parameters tab specifies the parameters that are unique to a particular firewall in your HA configuration. Follow the steps below.

1 In the **Cluster Mode** area, select one of the following options:

- **Designate as part of a Load Sharing High Availability Cluster** – Select this option if you want to configure load-sharing HA (both firewalls actively process traffic).
- **Designate as part of a Primary/Standby High Availability Cluster** – Select this option if you want to configure failover HA (only one firewall processes traffic, with the other firewall acting as a hot backup).

Note: To configure load-sharing HA or peer-to-peer failover HA, the firewalls must have the same hardware configuration. For more information on each HA configuration option, see [About HA configuration options](#).

2 [Conditional] If you selected **Primary-Standby** in the previous step, select one of the following options in the **Cluster Mode** area:

- **Primary** – Select this option if this will be the primary in your network. (This option is only used for the dedicated primary-standby HA configuration.)
- **Standby** – Select this option if this firewall is a standby in your network, or if you are configuring peer-to-peer HA.

Note: For peer-to-peer HA, you must configure **each** firewall as a standby.

3 In the **Control** field, select **Enabled** to enable HA for this firewall. (To disable HA, select **Disabled**.)

Note: You must reboot both firewalls before the HA configuration will take effect.

- 4 [Conditional] In the **Takeover Time** field specify the number of seconds that the primary must be unavailable before the secondary/standby will begin the takeover process.

Note: If the primary in an HA cluster goes into failure mode and the secondary/standby is not available, the primary will remain as the primary, but the Takeover Time value for that firewall will change to one, ensuring that if a secondary/standby becomes available, it can take over as the primary.

The secondary/standby Takeover Time value will differ depending on the type of HA configuration you are using:

- **Load sharing Takeover Time** – The takeover time for load-sharing HA cluster firewalls must be the same for EACH firewall that is participating in the HA configuration. The default value is 13 seconds for load-sharing configurations.
- **Peer-to-peer Takeover Time** – The takeover time for peer-to-peer HA cluster firewalls must be the same for EACH firewall that is participating in the HA configuration. The default value is 13 seconds for peer-to-peer configurations.
- **Primary-standby Takeover Time** – The takeover time for the primary is 3 seconds by default and cannot be modified. This value ensures that the designated primary will become the actual primary when it is activated. The default for the standby is 13.

Note: If you assign a standby Takeover Time value that is too close to 3 seconds, the standby may attempt to take over as the primary during periods when the primary is too busy processing data traffic to send the heartbeat.

Scheduling a soft shutdown for a load-sharing HA cluster Sidewinder

When a Sidewinder that belongs to an HA cluster is shutdown by an administrator (for example, to perform scheduled maintenance), a *soft shutdown* will automatically occur (assuming the shutdown time is not immediate). A soft shutdown provides a buffer period before the actual shutdown occurs, allowing the firewall to stop accepting new connections, while allowing most existing connections to complete before the firewall actually shuts down. IP packet filter processing is also transferred to the remaining firewall.

Note: A peer must be available in order to perform a soft shutdown.

By default, the soft shutdown process will begin 30 minutes prior to a scheduled shutdown. If the shutdown is scheduled to occur in less than 30 minutes, the soft shutdown process will begin immediately and will remain in effect until the actual shutdown time occurs. You can also manually increase or decrease the length of the soft shutdown period.

For example, suppose you configure the firewall to shutdown in two hours using the default soft shutdown of 30 minutes. The firewall continues to accept and process connections for 1.5 hours. When the firewall is 30 minutes from the shutdown time, it stops accepting new connections (the other firewall processes all new connections). Existing connections will have 30 minutes to complete. After the soft shutdown period completes, the firewall will shut down and will be unavailable until it is rebooted.

The soft shutdown feature is specified via command line. If you schedule a shutdown using the Admin Console, the default soft shutdown time will be applied. The following bullets provide examples of configuring an HA cluster firewall for shutdown:

- If you want the soft shutdown process to begin immediately, use the following command (the firewall must be shut down or manually rebooted once the soft shutdown process is complete):

```
cf failover softshutdown
```

- To configure soft shutdown to occur for a specific amount of time, as follows:

```
shutdown -s [soft_shutdown_time] [shutdown_time]
```

The *soft_shutdown_time* specifies that amount of time that soft shutdown will occur. The *shutdown_time* specifies the time at which the actual shutdown will occur. Each variable can be specified either as a number of minutes or as an exact date and time. If you are specifying the number of minutes, you must include a plus (+) sign in front of the minutes.

For example, if you want the firewall to shut down on Saturday, June 12, 2004 at 11:00 am with a 15 minute soft shutdown period, you would enter the following command:

```
shutdown -s +15 0406121100
```

In this case, the soft shutdown process would begin at 10:45 am, and the firewall would shutdown at 11:00 am on the specified day.

If you want the firewall to begin the soft shutdown at 6:00 am with an actual shutdown at 6:20 am, you would enter the following command:

```
shutdown -s 0600 0620
```

Note: For a complete listing of shutdown options, refer to the *shutdown* man page.

You can cancel a scheduled shutdown at anytime prior to the final 30 minute period by entering the **shutdown -c** command. However, once the firewall has entered soft shutdown mode, this command will no longer cancel the soft shutdown process. When the soft shutdown process is complete, you will need to reboot the firewall before it will properly function as part of the HA cluster.

Re-establishing an HA cluster if a cluster member fails

If a member of an HA cluster is no longer functional and must be re-installed, you can re-establish the cluster by restoring a configuration backup. You can use the configuration backup from either the failed cluster member or the remaining cluster member.

To re-establish an HA cluster:

- 1 Re-install the failed firewall or install a new system with the same host name.
- 2 Add the new or re-imaged firewall to the Admin Console tree.
- 3 Use the Admin Console to connect to the new or re-imaged firewall.
- 4 Restore the configuration file to the new system: Select **Maintenance > Configuration Backup**.

The firewall restores to the configuration backup and then restarts. When it finishes starting, it rejoins the HA cluster.

Restarting an HA cluster

If you make configuration changes to the following areas, you must restart both cluster firewalls:

- High Availability Common Parameters
- High Availability Local Parameters
- Heartbeat Interface

Caution: We recommend that you do not change the heartbeat interface parameters after creating an HA pair. See Knowledge Base article [9384](#) for details.

Note: A reboot is not required for a non-heartbeat interface configuration change.

- Non-interface fields on the HA screen

Note: If you make any changes to the fields in the High Availability Identification, IPSec Authentication, and Pair Members areas, a reboot is required.

To restart both firewalls in your HA cluster:

- 1 Shut down the secondary/standby:
 - a Select **Maintenance > System Shutdown > secondary icon**.
 - b Select **Halt System**.
 - c Select **Shutdown Immediately**.
 - d Click **Perform Shutdown**.
- 2 Shut down the primary:
 - a Select **Maintenance > System Shutdown > primary icon**.
 - b Select **Reboot to Operational Kernel**.
 - c Select **Shutdown Immediately**.
 - d Click **Perform Shutdown**.
- 3 Restart the primary.
- 4 Start up the secondary/standby.

SECTION 6

Troubleshooting

[Appendix A, Basic Troubleshooting](#)

[Appendix B, Re-install and Recovery Options](#)

A Basic Troubleshooting

This appendix provides basic troubleshooting information for Forcepoint Sidewinder.

Contents

[Troubleshooting rules](#)

[Troubleshooting logging in](#)

[Troubleshooting system status](#)

[Troubleshooting network status](#)

[Troubleshooting licensing problems](#)

[Troubleshooting High Availability](#)

[Troubleshooting NTP](#)

[Troubleshooting VPNs](#)

[Troubleshooting transparent \(bridged\) mode](#)

Troubleshooting rules

The following sections provide information on troubleshooting basic rule problems. For additional information on troubleshooting rules, refer to the `cf_policy` man page.

- [Failed connection requests](#)
- [Monitoring allow and deny rule audit events](#)
- [Active rules and DNS](#)

Failed connection requests

If the firewall rejects a connection request that you think should have succeeded, you can take steps to determine why the connection was rejected. Use the following steps to locate and correct rule configuration errors. They will also help you gain a better understanding of how your rules work.

- 1 Verify that the rule is configured correctly: Select **Policy > Rules**.

Verify that the rule in question specifies the correct source and destination burbs and endpoints. If the source burb is incorrect, the service will not listen for incoming connections on the correct network. Check all attributes closely, particularly port settings, application defense settings, and assigned authentication method.

- 2 Verify the position of the rules within the Active Rules window: Select **Policy > Rules**, and then click **Active Rules**.

The order of the rules in the Active Rules window is important. The attributes of a connection request sometimes may match multiple rules. If the traffic is inadvertently matching a similar rule, move the correct rule before the incorrect rule or adjust some of the properties on the incorrect rule.

3 Check the audit log information.

If the connection still fails, scan the audit log to determine which rule denied the connection. See [Chapter 11, Auditing](#) for details on viewing audit.

The text below displays a common scenario—a connection that failed to match a rule:

```
Feb  1 16:57:40 2007 CST  f_telnet_proxy a_proxy t_attack p_major
pid: 44770 ruid: 0 euid: 0 pgid: 44770 logid: 0 cmd: 'tnauthp'
domain: Atnx edomain: Atnx hostname: a.example.com
category: policy_violation event: ACL deny attackip: 10.10.1.155
attackburb: internal srcip: 10.10.1.155 srcport: 4584 srcburb: internal dstip: 1.2.3.4
dstport: 23 dstburb: external protocol: 6
service_name: all-tcpudp user_name: (null) auth_method: (null)
acl_id: Deny All cache_hit: 0 reason: Access Control List denial message
```

4 Turn on verbose auditing of rule (ACL) checks.

To determine why no proxy rule matched the connection request, type the following command to turn on verbose auditing of rule checks:

```
cf acl set loglevel=4
```

This increases the level of rule audits from the default level 2 (fatal and major errors) to level 4 (fatal, acl allows, and major errors).

When the next connection attempt is rejected, the service will generate a more verbose audit message as shown below:

```
Feb  1 16:45:28 2007 CST  f_ssh_server a_aclquery t_info p_trivial pid: 44689 ruid: 0
euid: 0 pgid: 44689 logid: 0
cmd: 'sshd' domain: ssh2 edomain: ssh2 hostname: a.example.com
Skipped 'ssh to 1.2.3.4': dest IP addr 10.10.1.15 did not match ('ipaddr', '1.2.3.4').
Skipped 'ssh from 10.10.1.150': source IP addr 10.10.1.155 did not match ('ipaddr',
'10.10.1.150').
Skipped 'ssh to DMZ': query dest burb internal != rule's: DMZ.
Skipped 'ssh to forcepoint.com': query dest burb internal != rule's: external.
AUTHENTICATION REQUIRED: Access tentatively allowed by rule 'ssh-internal'.
```

Use this output to determine why each rule failed to match the connection request. Locate the rule that you thought should have matched. Then inspect and correct the rule.

5 When you are done troubleshooting, type the following command to lower the level of rule audits back to the default:

```
cf acl set loglevel=2
```

Note: If you do not set the loglevel back to 2, you may run out of disk space.

The traffic should now match the correct rule.

Monitoring allow and deny rule audit events

Another troubleshooting tool is the rule monitoring tool (acat_acls). This real time monitoring tool enables you to display allow and deny rule audit events as they occur on the firewall. Because the rule audit events are displayed in real time, this tool provides a firewall administrator a unique window by which to view firewall rule activity. You can use the tool to determine if your rule database is properly configured, or to simply view how your rules are being used on a live system.

For example:

- If you are not certain whether your Telnet rule is properly configured, you can start the monitoring tool, attempt your Telnet connection and see (in real time) whether the connection is allowed or denied.
- If you want to see (in real time) which rules are currently the most heavily used, start the monitoring tool and watch as the current rule audit events scroll by within a command window.

The remainder of this section provides information on using the monitoring tool. Information can also be found by typing **man acat_acls** at a firewall command prompt.

Starting the rule monitoring tool (acat_acls)

To start the rule monitoring tool, enter the following commands at a firewall command prompt:

```
acat_acls -a -d
```

where:

- **-a** = display allow rule audit events
- **-d** = display deny rule audit events

If you want to view only allow rule audit events or only deny rule audit events, simply omit the undesired option (**-a** or **-d**).

Viewing the output from the rule monitoring tool

Each rule audit event is displayed on a single 80-character line. The source burb and the destination burb fields display the burb index number, not the burb name. The following example shows both an allow rule audit event and a deny rule audit event:

Table 75 Sample rule audit events

Action	Date	Time	Source burb	Source IP	Dest. burb	Dest. IP	Service
DENY	02/05/07	02:41:04	2	192.168.179.76	1	192.168.180.87	ping
ALLOW	02/05/07	02:42:32	2	192.168.179.76	1	192.168.180.87	telnet

Adjusting rule monitoring tool output

If the output from the monitoring tool is scrolling by too quickly, you can temporarily halt the output by pressing the following key combination:

Ctrl+S

To resume output, press the following key combination:

Ctrl+Q

You can also add **|more** or **|less** to a command to control how much output to view at a glance, or redirect the output to a file to view at another time.

Stopping the rule monitoring tool

To stop the rule monitoring tool, press the following key combination:

Ctrl+C

Active rules and DNS

If you create a proxy rule that contains a host name or a domain name, that rule will consult the Domain Name System (DNS) in order to translate the name to its corresponding IP address. Because of this, there are some facts related to DNS that you should consider when setting up your security policy.

Sidewinder can be configured to use transparent DNS, or one DNS server (known as single or unbound DNS), or two DNS servers (known as split DNS). The split DNS scenario is the most secure, as one DNS server is dedicated to your Internet burb and the second DNS server services your remaining burbs. This essentially isolates the two DNS servers from each other, protecting your non-Internet burbs from attacks by malicious persons on the Internet.

However, it is theoretically possible for attackers on the Internet to feed false information to your Internet DNS server. Therefore, you should be careful when using rules with endpoints that are resolved using DNS (domain and host network objects) to allow or deny access to specific hosts on the Internet.

When dealing with outside connections, there are steps that you can take to increase the level of assurance:

- 1 Use IP addresses in your rule instead of host names or domain names. This avoids having to depend on external DNS.
- 2 Make the rule demand strong authentication (for example, SafeWord).
- 3 Make the rule demand encryption of the connection (for example, VPN).

Tip: For additional protection, you should do a combination of the above.

Troubleshooting logging in

If you forget your administrator password, you can boot the firewall into emergency maintenance mode (EMM) and reset your password.

Note: By default, the EMM does not require authentication. However, if you have configured your system to require authentication to enter that mode, you will need to temporarily disable EMM authentication before you can enter that mode and change your password. For information on disabling EMM when you have forgotten your password, see [Changing authentication requirements for emergency maintenance mode](#).

Restoring access to the firewall

If an administrator accidentally alters the rule set in a way that prevents an administrator from logging into the firewall (for example, moving the Deny All rule to the first position or deleting certain access rules), use this procedure to restore access.

To regain access to both the local console and the Admin Console:

- 1 At the local console, reboot the firewall.
- 2 At the Sidewinder boot menu, select **Boot in Emergency Maintenance Mode**. A prompt appears stating:
Enter full pathname of shell or RETURN for /bin/sh:
- 3 Press **Enter**. A system prompt appears.
- 4 Restore console access by entering the following:

```
cf policy restore_console_access
```

This command recreates the default local console and the Admin Console rules. The rules are added to the beginning of the rule set.

- 5 Reboot to the Operational kernel by entering the following:

```
shutdown -r now
```

You can now log in at the local console or with an Admin Console session initiated on the firewall's internal burb.

Changing a forgotten password

You must be at the local console to run this procedure.

To change your administrator account password:

- 1 Reboot the system.
- 2 At the Sidewinder boot menu, select **Boot in Emergency Maintenance Mode**.
- 3 Press **Enter** when asked what shell patch to use. The system prompt appears.
- 4 Enter the following command to change your password:

```
cf adminuser modify user=name password=newpassword
```

- 5 Reboot to the Operational kernel by entering the following command:

```
shutdown -r now
```

You can now log into the firewall using your new password.

Manually clearing an authentication failure lockout

If you have enabled the authentication failure lockout option and are locked out of your system, have another administrator log in using the Admin Console and clear the lock (see [Configuring an authenticator](#)).

However, if you do not have another administrator who can clear your lock for you, you can still manually clear your lock by successfully logging in at the firewall's local console. When you successfully log in, the lock automatically clears and you can now log into the Admin Console as usual.

Changing authentication requirements for emergency maintenance mode

To change authentication for entering the emergency maintenance mode:

- 1 Log into the Admin Console, and select **Maintenance > File Editor**.
- 2 Click **Start File Editor**.
- 3 Select **File > Open**.
- 4 In the Source field, select **Firewall File**.
- 5 In the **File** field, type `/etc/ttyd` and click **OK**.
- 6 Edit the following line:

```
console none unknown off secure
```

- To require authentication, change the value to **insecure**.
- To disable authentication, change the value to **secure**.

- 7 Save your changes and close the file editor.

The authentication requirements are now changed.

Troubleshooting system status

You can use the commands in the following sections to display information on the current status of your network connections and view what is happening on the system.

CPU usage

CPU usage allows you to obtain information on system performance. To view CPU usage information, enter each of the following commands at a Sidewinder command prompt:

```
vmstat
uptime
top
```

Process status

To view the status of all processes currently running on the firewall, enter either of these commands at a Sidewinder command prompt:

```
ps -axd
pss processname
```

This information is useful for tasks such as determining which processes are using excessive CPU time. The **pss** command allows you to look at information about the processes running on the system. This command is a variation on the standard **ps** process status command in that it includes information on the Sidewinder domains. To display process information, enter:

```
ps -d
```

This command lists process information as well as information on the domains in which processes are operating.

In addition to the standard information displayed with the **ps** command, the **-d** switch provides the following additional information:

LABEL	PID	TT	STAT	TIME	COMMAND
secureos/Dmnd	189	con	Ss+	0:01.30	/usr/libexec daemond
secureos/Admn	1360	p	0R+	0:02.05	ps -d

where:

- **LABEL** – domain name
- **PID** – process identification number
- **TT** – terminal line from which the process was initiated
- **STAT** – current status of the process
- **TIME** – total amount of CPU time used by the process
- **COMMAND** – command line used to start the process

Disk usage

To view statistics about the amount of free disk space on a file system, enter the following command at a Sidewinder command prompt:

```
df
```

This information is useful to determine which file systems are using the most disk space.

Viewing administrator activity

To view which administrators are currently logged onto your Sidewinder, enter the following command at a Sidewinder command prompt:

```
who
```


When you use this utility, you can see the administrator's login name, console name, the date and time of their login, and their host name if it is not a local host.

```
lloyd      tty??    Feb 23    22:11    (a.example.com)
lloyd      tty0      Feb 23    21:34    (10.1.1.1)
```

Troubleshooting network status


You can use the Admin Console or commands to check firewall network status.

you can use the following commands to check the system and network status, and basic DNS information

Checking network status using the Admin Console

Use the **Tools** menu on the Admin Console's toolbar to gather network information for troubleshooting purposes.

Figure 388 Tools menu on the Admin Console toolbar



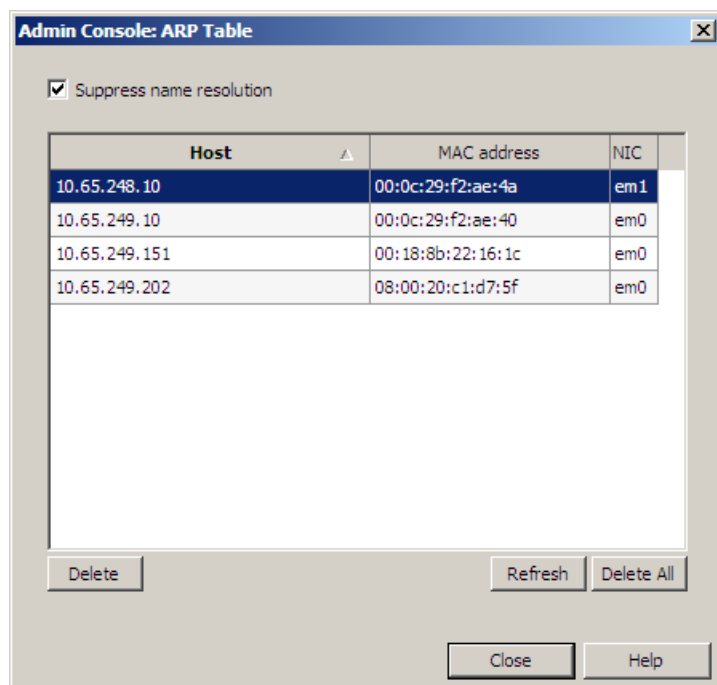
See the following for information about each tool:

- [About the ARP Table](#)
- [About the Get Route window](#)
- [About the DNS Lookup window](#)
- [About the Ping Test window](#)
- [About the TCP Dump window](#)
 - [About the TCP Dump Parameters window](#)
 - [About the Running TCP dump window](#)
- [About the Traceroute window](#)

About the ARP Table

Use this window to view the association between each MAC address on the firewall and its corresponding IP address, as well as the NIC used to reach the IP address.

Figure 389 ARP Table



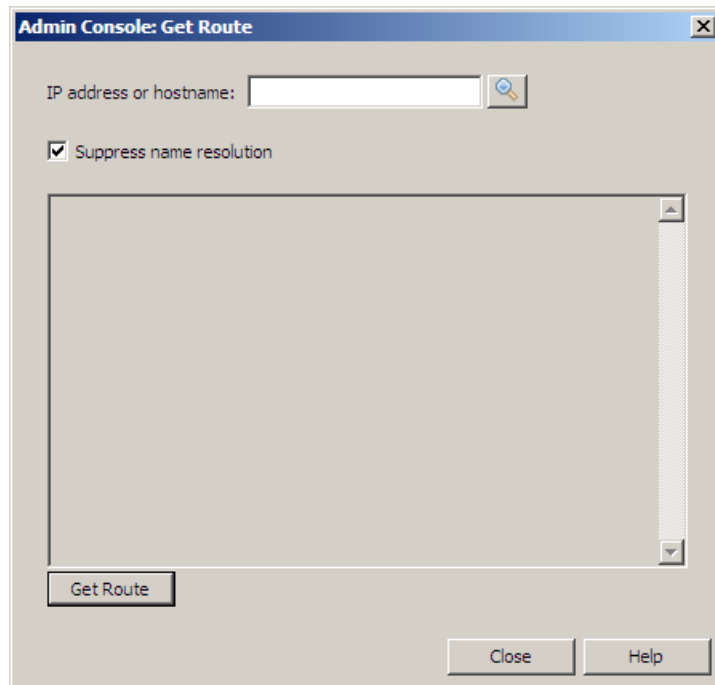
You can perform the following actions:

- To view an updated table, click **Refresh**.
- To delete an entry in the table, select the appropriate entry and click **Delete**. You can select multiple entries to delete. You can delete all entries by clicking **Delete All**.
- To show the host domain name as well as the IP address, clear the **Suppress name resolution** check box and click **Refresh**.

About the Get Route window

Use this window to find the first gateway in the route from the firewall to a stated destination.

Figure 390 Get Route window




- 1 In the **IP address or hostname** field, enter a destination. To find the IP address for a host name, type the name and click **DNS Lookup**.
- 2 Click **Get Route**. The route information appears in the lower field.
 - **route to** is the destination
 - **gateway** is the first gateway in the route to the destination
 - **interface**, **if address**, and **burp** belong to the firewall's interface

Note: To show **gateway** and **if address** as domain names, clear the **Suppress name resolution** check box.

About the DNS Lookup window

Use this window to find the IP address for a host name.

Figure 391 DNS Lookup window




Enter a host name in the **Host name** field and then click **Lookup**. IP addresses associated with the host name appear in the lower pane.

About the Ping Test window

Use this window to test interface connectivity.

Figure 392 Ping Test window




- 1 In the **IP address to ping** field, enter an IP address or fully qualified domain name that the ping will be sent to. To find the IP address for a host name, type the name and click **DNS Lookup**.
- 2 [Optional] In the **Commandline flags** field, enter parameters for the ping test. You can enter alphanumeric characters, dashes (-), and underscores (_).
- 3 Click **Start Ping**. The button changes to **Stop Ping** and the ping results appear in the window.
- 4 Click **Stop Ping** to stop the test.
- 5 Click **Close**.

About the TCP Dump window

Use this window to manage tcpdump parameters and to start the tcpdump process.

Figure 393 TCP Dump window



You can perform the following actions:

- Create, modify, or delete tcpdump.
 - Click **New** and enter interface, output, and filtering information in the TCP Dump Parameters pop-up window. The entry appears in the table on the TCP Dump window.
 - Select a table entry and click **Modify** to make changes to the parameters.
 - Select a table entry and click **Delete** to remove it from the table.
- Enable or disable an entry by selecting or clearing the check box in the Enabled column. Tcpcdumps are performed only for enabled entries.
- Save a tcpdump entry by clicking **Save**. The entry will remain in the table for later uses.
- Start the tcpdump by clicking **Start TCP dumps**. The Running TCP dump window appears and the tcpdump begins for all enabled entries.

About the TCP Dump Parameters window

Use this window to select an interface to capture network traffic for and to set other parameters for a tcpdump.

Figure 394 TCP Dump Parameters window

The screenshot shows the 'TCP Dump: TCP Dump Parameters' dialog box. It includes fields for 'Name', 'Interface' (set to 'em0'), 'Burb' (set to 'external'), 'Enabled' (checked), 'Promiscuous mode' (unchecked), 'Bytes to capture' (set to 'Header bytes only'), and 'Limit output' (set to '1 Gigabyte'). A 'Filtering' section contains a 'Template' dropdown set to 'No filtering'. A message box states 'No filtering will be done. All packets will be captured.' The bottom of the window has 'OK', 'Cancel', and 'Help' buttons.


You can make the following entries:

- **Name** – Enter a name to identify this set of tcpdump parameters.
- **Interface** – From the drop-down list, select an interface to capture network traffic for.
- **Enabled** – Select this check box to enable these tcpdump parameters. Tcpdumps are performed only for enabled entries.
- **Promiscuous mode** – Select this check box to capture packets that do not belong to this interface.
- **Bytes to capture** – Select from the drop-down list:
 - **Header bytes only** captures only packet headers.
 - **Full packet capture** captures packet headers and data.
- **Limit output** – From the drop-down list, select a size limit for the tcpdump files. If the limit is reached, the data is deleted and the dump starts over.
- **Filtering** – Select a filter option from the **Template** drop-down list.
 - **No filtering** captures all traffic.
 - **Single host** captures traffic for a designated host, and all other traffic is discarded. You can enter a host IP address, a port number, or both. To find the IP address for a host name, type the name and click **DNS Lookup**.
 - **Connection** captures all traffic for two designated hosts. You can enter IP addresses and a port. To find the IP address for a host name, type the name and click **DNS Lookup**.
 - **Custom** – Enter filters that conform to the tcpdump format. See the tcpdump man page on the firewall for more information.

About the Running TCP dump window

Use this window to see the progress of a tcpdump, to stop the tcpdump, and to save and view the dump files.

Figure 395 Running TCP dump window




- The table shows the tcpdump parameters and file size for this occurrence.
- To stop the tcpdump process, click **Stop TCP dumps**.
- To save the dump files, click **Save** and then use the pop-up window to navigate to the desired location. A pop-up window appears for each file. When all files are saved, the Running TCP dump window closes.
- To view the dump files, click **Launch pcap viewer**.
 - The Admin Console and firewall do not have a viewer. You must install a third-party tool such as Wireshark to view the dump files.
 - When you click **Launch pcap viewer**, you are first prompted to save the files. Use the pop-up window to navigate to the desired location. A pop-up window appears for each file. When all files are saved, your default viewer appears.

About the Traceroute window

Use this window to see all of the gateways that traffic passes through on a round trip between the firewall and a destination.

Figure 396 Traceroute window



To trace a route:

- 1 In the **Traceroute IP address** field, enter the IP address of the destination. To find the IP address for a host name, type the name and click **DNS Lookup**.
- 2 [Optional] In the **Commandline flags** field, add parameters to the trace. See the `traceroute` man page for more information.
- 3 Click **Start Trace**. The gateways and other route information appear in the lower window.

Traceroute finished appears in red when complete. You can click **Stop Trace** to stop the process earlier.

Checking network status using the command line

You can use the following commands to display information on the status of your network connections, routing tables, and network utilities. These commands can provide snapshots of different aspects of your system with command line outputs.

- [ping](#)
- [traceroute](#)
- [Active listens](#)
- [Network interfaces](#)
- [Routing tables](#)
- [route get](#)
- [dig](#)

ping

The **ping** command checks whether an Internet system is running by sending packets that the remote system should echo back. As output, **ping** lists how much time it took for the message to travel to the other system and back, the total number of packets sent and received, the percent of packets lost, and the average and maximum time it took for a round trip. To view this information, enter:

```
ping -c 5 ipaddress
```

traceroute

The **traceroute** command provides information on the gateways an IP packet must pass through to get to a destination. As input, the command needs the host name or IP address of the destination system. It then sends these IP packets from your Sidewinder to that address. As output, it lists the host names and IP addresses of each system the packets were handed off to and how long it took to send each packet back and forth.

To view this information, enter:

```
traceroute -m 50 ipaddress
```

To run traceroute that crosses burb boundaries, follow it to the firewall and then initiate a second traceroute from the firewall itself.

Note: Traceroute is not allowed through the firewall. In addition to security risks, NAT prevents most sites from getting a return from the external network (Internet) because of non-routable addresses.

Active listens

To view the status of all active listens, enter:

```
netstat -na|grep LISTEN
```

To view the status of all open connections, enter:

```
netstat -na|grep ESTAB
```

To view the status of all connections waiting for a termination request, enter:

```
netstat -na|grep FIN_WAIT_1
```

To view the status of all connections waiting for enough time to pass to be sure the remote TCP received the acknowledgment of its connection termination request, enter:

```
netstat -na|grep TIME
```

Network interfaces

To view the status of network interfaces on the firewall, enter:

```
netstat -in
```

To view statistics of network interfaces on the firewall, enter:

```
netstat -s
```

Routing tables

To view the status of the Operational kernel's available routes and their status, enter the following command at a Sidewinder command prompt:

```
netstat -r
```

For the same results without DNS data, enter:

```
netstat -nr
```

route get

The **route get** command looks up the route for a destination and displays the route in the window. To view this information, enter the following command at a Sidewinder command prompt:

```
route get ipaddress
```

The following shows sample output for this command:

```
# route get 10.136.78.11
route to: firewall11.ext.rack20.dfb.test
destination: default
mask: default
gateway: router.ext.rack6.dfb.test
interface: dc0
if address: firewall11.ext.rack6.dfb.test
region: 1
flags: <UP,GATEWAY,DONE,STATIC>

  recvpipe    sendpipe    ssthresh    rtt,msec    rttvar    hopcount    mtu     expire
  0           0           0           0           0         0           1500    0
```

dig

You can use the **dig** command to display DNS information.

The **dig** (Domain Information Groper) command gathers information from DNS based on a hostname or an IP address. The command queries servers based on type (NS for name servers, MX for mail servers, etc.) and has many advanced options. This command is more powerful than **nslookup**.

```
dig hostname
```

```
dig -x ipaddress
```

Here is an example of dig output:

```
; <<>> DiG 9.3.2 <<>> example.com
;; global options:  printcmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 15043
;; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 1, ADDITIONAL: 0
;; QUESTION SECTION:
;example.com.          IN      A
;; ANSWER SECTION:
example.com.          86400   IN      A      192.168.1.123
;; AUTHORITY SECTION:
example.com.          3600    IN      NS      stpdc02.scur.com.
;; Query time: 7 msec
;; SERVER: 10.65.240.246#53(10.65.240.246)
;; WHEN: Mon Jan  8 19:06:52 2007
;; MSG SIZE  rcvd: 80
```

Troubleshooting licensing problems

If the firewall comes up in failure mode because it did not license during the reboot, check the following:

- Try to obtain the license by entering:

```
cf license get
```

- Verify that there is a default route by entering:

```
netstat -nr
```

If there is not a default route, add it back with

```
cf static add route=default gateway=aaa.bbb.ccc.ddd
```

where *aaa.bbb.ccc.ddd* is the next hop router for the default route.

- Verify that DNS is resolving by entering:

```
dig www.example.com
```

- Obtain the license by doing one of the following:

- If DNS is resolving, enter `cf license get`.
- If DNS is not resolving, you will need to get the license using the activation server's IP address by entering the following on a single line:

```
cf license get  
activation_url=http://sidewinder.activations.forcepoint.com/activation.cfm
```

- Reboot the firewall again by entering:

```
shutdown -r now
```

The firewall should now be correctly licensed and fully functional.

Troubleshooting High Availability

This section provides information to determine whether High Availability is functioning properly.

Viewing configuration-specific information

The **cf failover query** command gives you configuration-specific information, as shown in the following example:

```
failover set autoreset=on backup_heartbeat_burb='' enabled=on firewall_id=1 \  
    heartbeat_burb=heartbeat load_sharing=off multicast_group=239.255.0.1 \  
    ping_wait=1 priority=245  
failover set password=b602b701bb type=SHA1  
failover modify network=10.65.248.0 failover_ip=10.65.248.20 \  
    arp_reset_addrs='' monitor_addrs='' monitor_allowed_failures=3 \  
    monitor_interval=30 monitor_link=yes  
failover modify network=10.65.249.0 failover_ip=10.65.249.20 \  
    arp_reset_addrs='' monitor_addrs='' monitor_allowed_failures=3 \  
    monitor_interval=30 monitor_link=yes  
failover modify network=192.168.1.0 failover_ip=192.168.1.20 \  
    arp_reset_addrs='' monitor_addrs='' monitor_allowed_failures=3 \  
    monitor_interval=30 monitor_link=yes
```

The **cf cluster status** command gives an overview of the whole cluster, as shown in the following example:

```
HA Cluster Status Information  
=====
```

Primary Host:	gabrielle.example.net	
Primary IP Address:	192.168.1.3	
Cluster Burb:	'heartbeat'	
Cluster Cert:	'Default_Enterprise_Certificate'	
Cluster CA:	'Default_Enterprise_CA'	

Member Name	State	IP Address
-----	-----	-----
gabrielle.example.ne	registered	192.168.1.3
xena.example.net	registered	192.168.1.2


```
Policy and Peer Connection Status  
=====
```

gabrielle.example.net (primary)	

Connection State	: Localhost
Policy Version	: 14-1205148992.63-1205071621

FW Version : 70006
Status : Up to date - Current

xena.example.net (peer)

Connection State : Currently Connected
Last Dispatch : 2008-03-10 07:39:21.957857
Policy Version : 14-1205148992.63-1205071621
FW Version : 70006
Status : Ready

Viewing status information

The `cf failover status` command gives you information on whether or not HA is active, what state the system is in (primary or secondary/standby), and useful statistical information.

Viewing status information for a primary

The following example shows sample results for a primary in a peer-to-peer HA configuration:

```
This system is operating as primary.

Failover is running in burb 3
IP alias 192.168.1.20 assigned to interface fxp0
IP alias 10.65.249.20 assigned to interface em0
IP alias 10.65.248.20 assigned to interface xl0
This system was configured as a standby with priority 245 for firewall ID 1.

Failover interface status:
fxp0    Interface is up
em0     Interface is up
xl0     Interface is up

IP Filter tracking state as primary

Active firewall list:

A backup heartbeat interface is not configured

Statistics for failover

Failover running since Sun Mar  9 10:07:02 2008

Failover allowing 3 seconds for interface swap (default)

Number of advertisements sent                = 10982
Number of received advertisements           = 77430
Number of rcvd advertisements since primary = 77430
Number of times this system has become primary = 1
Number of release messages received         = 1
Number of release messages sent             = 0
Number of failed takeover attempts          = 0
Number of possible duplicate primary messages = 0
Number of heartbeat ack messages received   = 0
Number of heartbeat ack messages sent       = 77430
Number of messages received with errors     = 0
Number of same priority advertisements rcvd = 77430
```

Viewing status information for a secondary

The following example shows sample results for a secondary that is configured for load-sharing HA:

This system is operating in load sharing mode as secondary.

This system is node 1.

The primary is node 0 (192.168.1.3).

Failover is running in burb 3

cluster heartbeat address 192.168.1.20 assigned to interface fxp0

shared cluster address 10.65.249.20 assigned to interface em0

shared cluster address 10.65.248.20 assigned to interface xl0

Failover interface status:

fxp0 Interface is up

em0 Interface is up

xl0 Interface is up

IP Filter tracking state as load sharing peer

Active firewall list:

node address

0 192.168.1.3 (primary)

A backup heartbeat interface is not configured

Statistics for failover

Failover running since Tue Mar 11 09:17:28 2008

Failover allowing 3 seconds for interface swap (default)

Number of advertisements sent	= 0
Number of received advertisements	= 114
Number of rcvd advertisements since primary	= 114
Number of times this system has become primary	= 0
Number of release messages received	= 0
Number of release messages sent	= 0
Number of failed takeover attempts	= 0
Number of possible duplicate primary messages	= 0
Number of heartbeat ack messages received	= 0
Number of heartbeat ack messages sent	= 114
Number of messages received with errors	= 0
Number of same priority advertisements rcvd	= 0

Identifying load-sharing addresses in netstat and ifconfig

Output for `netstat -i` queries will display load sharing addresses with a plus (+) sign. The following example displays the results for the `netstat -i` command with load sharing enabled.

Name	Mtu	Network	Address	Ipkts	Ierrs	Opkts	Oerrs	Coll
fxp0	1500	<Link#1>	00:a0:c9:9d:99:a1	3059	0	2869	0	0
fxp0	1500	10.65.249/	2410.65.249.20	129	-	145	-	-
fxp0	1500	10.65.249/	24a	98	-	77	-	-
xl0	1500	<Link#2>	00:10:5a:98:51:26	0	0	10	0	0
xl0	1500	10.65.248/	2410.65.248.20	0	-	0	-	-
xl0	1500	10.65.248/	24a.example.com	0	-	0	-	-
em0	1500	<Link#3>	00:0c:f1:c7:ba:ea	506253	0	356438	0	0
em0	1500	10.65.247/	2410.65.247.20	95	-	642	-	-
em0	1500	10.65.247/	2410.65.247.2	611	-	248	-	-
lo0	16384	<Link#4>		7478	0	7478	0	0
lo0	16384	127	localhost	3951	-	7478	-	-
lo2	16384	<Link#5>		114	0	114	0	0
lo2	16384	127	localhost_2	57	-	114	-	-
lo1	16384	<Link#6>		0	0	0	0	0
lo1	16384	127	localhost_1	0	-	0	-	-
lo3	16384	<Link#7>		0	0	0	0	0
lo3	16384	127	127.3.0.1	0	-	0	-	-

Troubleshooting NTP

Use the commands in this section to verify that NTP traffic is passing as expected. If NTP is experiencing problems, use the following commands to help determine the cause.

Verify that NTP is running and using the correct time

- Check the NTP daemon is running on the time server that is serving time to the firewall:

```
ps -df |grep ntpd
```

If ntpd is not running, correct this problem before proceeding.

Run the ntpdate command to show that the NTP daemon is responding to the request:

```
ntpdate -q -u server_ip
```

where *server_ip* is the IP address of the time server. This command also reports if the queried server time is different from the firewall's system time. To check the time on a firewall-hosted NTP server, use the IP address of the burb where the NTP server is running.

Note: To run the above commands, you need command line access to the NTP server.

- Check the process to see that the firewall NTP server is enabled and running:

```
pss ntpd
```

```
cf server status ntp
```

- If you have NTP properly configured and enabled, you should be able to monitor NTP packets being sent/received on the appropriate firewall interfaces. To do so, enter the following commands:

```
tcpdump -npi if_name udp port 123
```

where *if_name* is the interface and number that you are troubleshooting (for example **em0**, **em1**, etc.)

In the **tcpdump** output, check to make sure that NTP packets are being both sent and received. If traffic is not flowing both ways, verify the routing connectivity between your firewall and the NTP server.

- To check the firewall's exact system time, enter the **date** command and compare it to a known good clock source (for example, **www.time.gov**).

Troubleshooting NTP when the firewall is an NTP client

Here is an example of the output of the **ntpd** command, where the firewall is configured as an NTP client and the IP address 100.1.1.199 is the IP address of the firewall's interface that is running NTP:

```
fw:Admn {1} % ntpdc -s 100.1.1.199

remote          local          st poll reach  delay  offset  disp
=====
*cisco1-mhk.kans 100.1.1.199  2  128  377 0.06488  0.013563 0.00798
.triangle.kansas 100.1.1.199  2  128  377 0.06522  0.013999 0.00902
.ns.nts.umn.edu  100.1.1.199 16  128  377 0.07059  0.002498 0.00983
```

This is a list of the time servers that the firewall is configured to query for time. The stratum field will tell you whether the firewall is able to communicate properly with the time servers. In this case, the first two entries have a stratum of 2, while the third has a stratum of 16. A stratum of 16 always indicates that the firewall is not synchronizing successfully with a time server. In this case, the firewall was able to communicate properly with the first two servers, but not with the third.

Troubleshooting NTP when the firewall is an NTP server

Here is an example of the output of the `ntpd` command when the firewall is configured as an NTP server and the IP address 100.100.2.200 is the IP address of the firewall's interface that is running NTP:

```
Fw2:Admn {1} % ntpdc -s 100.1.2.200

remote      local      st poll reach  delay   offset   disp
=====
100.1.2.50  100.1.2.200  3  128  377 0.06488  0.013563 0.00798
```

This is a list of the NTP clients that are configured to get time from the firewall. The stratum field will tell you if the firewall is able to communicate properly with the NTP clients. In this case, the client 100.1.2.50 has a stratum of 2, which means it is communicating properly. A stratum of 16 would indicate that the client was not synchronizing successfully with the firewall. In most situations, you can look at the stratum to determine whether or not NTP is configured properly and working. There is additional information in the `ntpd` output that may be useful in troubleshooting NTP issues. However, that information is out of the scope of this document.

Synchronization problems

If you are having synchronization problems, enter the following command:

```
ntpd
```

Use `man ntpdc` to see detailed information.

Possible reason NTP stopped

NTP is designed to automatically quit whenever the client's time deviates from the server's signal by more than 15 minutes. When a deviation of this magnitude occurs, NTP writes a message to file `/var/log/messages` before quitting.

To restart NTP, first set the firewall's clock manually (refer to [Setting the system date and time](#)) and then follow the directions in the next section for restarting NTP.

Restarting NTP from a command line prompt

If the NTP process stops, you can restart the NTP process by doing the following:

- 1 To start the NTP time server, enter the following command:

```
cf daemon restart agent=ntp
```

- 2 [Optional] Verify the state of the NTP servers by entering the following command:

```
cf server status ntp
```

Troubleshooting VPNs

In addition to standard logging, the firewall also performs auditing of certain system events which allows you to generate information on VPN connections. [Table 76](#) shows some useful commands you can use to track VPN connections in real time mode and check VPN settings/configuration.

Table 76 Basic VPN troubleshooting commands

Command	Description
<code>tcpdump -npi ext_if port 500 or proto 50 or proto 51</code>	Show IPsec, ESP, and AH traffic arriving at the firewall.
<code>tcpdump -npi if_name udp port 4500</code>	Show NAT-T traffic arriving at the firewall.
<code>cf ipsec q</code>	Review VPN policies.
<code>cf ipsec policydump</code>	Determine if VPN is active. The presence of SPI and transform numbers indicates the secure connection is functioning.
<code>showaudit -vk</code>	<p>Show detailed audit trace information for VPN in real time.</p> <p>To enable a more detailed auditing level, adjust the ISAKMP server's audit level:</p> <ol style="list-style-type: none">1 In the Admin Console, select Policy > Rule Elements > Services.2 Select isakmp, and click Properties.3 Set the audit level to Verbose.4 Click OK on both windows.5 Save your changes.

Troubleshooting transparent (bridged) mode

For information on troubleshooting transparent mode, see <https://support.forcepoint.com>

Basic Troubleshooting

Troubleshooting transparent (bridged) mode

B Re-install and Recovery Options

This appendix provides re-install and recovery options for Forcepoint Sidewinder.

Contents

[About re-install and recovery](#)

[Recovery options](#)

[Re-install options](#)

About re-install and recovery

There are several options available when your firewall experiences a severe software or hardware problem. Solutions range from a rollback to a previous restore point to fix short-term problems to re-installing a full backup onto new hardware. There are two types of solutions:

- [Recovery options](#) – Use one of the recovery options when the firewall has experienced a short-term problem. For example, you can use the rollback feature to recover from an issue caused by a recently installed patch.
- [Re-install options](#) – Use one of the re-install options when the firewall has experienced a serious hardware failure, or when the firewall is to be repurposed.

The following table lists common problems and their recommended solution.

Table 77 Problems and solutions

Problem	Solution
Patch upgrade failed or caused unexpected behavior	Use the uninstall or rollback option.
Software or configuration disaster	Select the least disruptive option: configuration restore, uninstall or rollback, or re-install.
Repurpose system in some way, such as moving a firewall from one network to another	Re-install and create a new configuration.
Hard disk failure or system replacement	Re-install and restore an existing configuration.

Recovery options

Occasionally, you may experience a situation where you need to restore your firewall to a working configuration. Sidewinder has several recovery options to fit different severities and types of situations. Do the following to maximize recovery success:

- Create configuration backups on a regular basis.
- Create disaster recovery backups after installing a new patch.
- When selecting a recovery option, always attempt the least disruptive option first and determine if that solves your problem.

The following sections describe the available recovery options:

- [Configuration restore](#)
- [Uninstall](#)
- [Rollback](#)
- [Disaster recovery](#)

Configuration restore

The policy configuration restore replaces the firewall's current policy with a saved configuration file. Create configuration backups frequently to ensure you have an up-to-date configuration backup from when the policy was known to be configured correctly.

Use this option when misconfiguration or data corruption renders the current policy unacceptable, or to return to a recent configuration after completing a re-install.

See [Configuration file backup and restore](#) for this procedure.

Uninstall

The uninstall option removes a patch but maintains the current configuration. Patches can be uninstalled individually or several at a time. Patches that can be uninstalled are listed with an Uninstall status of **Yes** on the Software Management window. They are often smaller patches such as vendor patches or hotfixes that do not include features or substantial changes.

Use this option when a newly installed uninstalleable patch fails to install or introduces behavior that is incompatible with your policy.

This procedure is an option in the **Maintenance > Software Management** area. See [About the Software Management: Manage Packages tab](#) for information on uninstalling patches.

Rollback

The rollback option reverts your firewall to the previous restore point, which is a snapshot of the patch level and policy configuration just before the most recent patch was installed. If multiple patches were installed at the same time, then the restore point is before the earliest patch.

Use this option when a newly installed patch fails to install or introduces unexpected behavior that is incompatible with your policy. With the rollback option, all configuration changes made between when the problematic patch was installed and when the rollback was initiated are lost.

This procedure is an option in the **Maintenance > Software Management** area. See [About the Software Management: Rollback tab](#) for information on scheduling and initiating a rollback.

Disaster recovery

This option automatically saves the necessary configuration information, patches, and hotfixes to a USB drive. Use this option when recovering from a failed hard drive or when configuring a replacement firewall. Create backups often to ensure that you can quickly return to the correct patch level and configuration.

About disaster recovery

The disaster recovery restore option restores your firewall to a previous patch level and configuration by restoring files from a USB drive containing a disaster recovery backup. This backup includes all installed patches, an initial configuration, and a standard configuration backup file. This option is an efficient way to return your firewall to its previous patch level and a known configuration. If you have a configuration backup that is more recent than your disaster recovery backup, you may want to restore it after the recovery completes.

Use this option to recover after replacing a hard disk or entire system, or if the policy becomes seriously misconfigured or corrupt.

Note the following:

- Disaster recovery backups can be saved only on a USB drive. See [Selecting a USB drive](#).
- Do not alter the disaster recovery backup file.
- The disaster recovery files are intended to be restored on the same hardware that they were created on. If you attempt to restore the backup to different hardware, expect to make significant adjustments.

Caution: Failure to follow the guidelines listed above could corrupt the disaster recovery backup.

Selecting a USB drive

Your USB drive must meet the following requirements:

- Supported USB drive sizes are 1 GB, 2 GB, 4 GB, and 8 GB. The USB drive must be large enough to hold the configuration and patches. In general, a 1 GB USB drive should be sufficient for most firewalls. Configurations with large home directories may require a larger size.
- Your USB drive must be formatted in MS-DOS.
- The firewall's USB port must be enabled in the BIOS settings.

Creating the disaster recovery USB drive

Create a disaster recovery USB drive using the Create Disaster Recovery Backup option on the Configuration Backup window (**Maintenance > Configuration Backup**). This recovery USB drive includes the configuration and installed patches that were on the firewall when the recovery media was made.

See [Using the Configuration Backup window](#).

Using the Configuration Backup window

To create a disaster recovery USB drive:

- 1 Insert a USB drive into one of the firewall's USB ports.

Caution: This process will overwrite previous disaster recovery backups contained on this USB drive.

- 2 Select **Maintenance > Configuration Backup**.

- 3 Click **Create Disaster Recovery Backup**. The Configuration Backup: Disaster Recovery window appears.

- 4 [Optional] Enter a key to encrypt the disaster recovery backup. Valid values include alphanumeric characters, periods (.), dashes(-), and underscores (_).

- This key will not be saved. You must remember it. You will not be able to restore the disaster recovery backup without this key.
- You do not have to enter an encryption key. If you click **OK** without entering an encryption key, the backup continues.

Enter the key again to verify.

- 5 Click **OK**. A warning message appears.

- 6 Click **Yes** to confirm the backup.

A progress bar appears while the files are backed up to the USB drive.

Note: Do not remove the USB drive from the firewall until the "Disaster recovery successful" message appears.

When the backup is complete, a "successful" message appears.

- 7 Click **OK**.

The disaster recovery USB drive has been created.

Restoring the backup

To restore the disaster recovery configuration:

- 1 Verify you have the necessary tools:
 - One of the following types of firewall installation media:
 - *Sidewinder Installation CD*
 - *Sidewinder Installation USB Drive*
 - USB drive containing the disaster recovery backup
- 2 Insert the disaster recovery USB drive into the firewall's USB port.
- 3 Follow the appropriate re-installation method.
 - If the firewall has a CD drive, see [Re-installing your firewall from a CD-ROM](#).
 - If the firewall does not have a CD drive, see [Re-installing your firewall from a USB drive](#).
- 4 Immediately after the re-install completes, remove the installation media from the firewall.
- 5 Leave the disaster recovery USB drive in the firewall's USB port until the recovery is complete. You can remove the drive after the Admin prompt appears, or leave the USB drive in the firewall for future use.

The firewall is now restored to the patch level and configuration saved in that backup.

Re-install options

Use one of the re-install options when the firewall has experienced a serious hardware failure, such as a failed hard drive, or when the firewall is to be repurposed, such as moving it from one network to another.

- **Re-installing from the virtual CD (VCD)**

This re-install option restores your firewall using the virtual CD (VCD). During the re-install process, you select which patches to install from a list of packages that were available on the firewall before starting the re-install procedure. This option provides flexibility in what version the firewall will be after the recovery. Make sure you know which patches must be installed to match the configuration backup you intend to restore.

See [Re-installing your firewall from the virtual CD](#).

Note: Using this option with disaster recovery media is not recommended, as there is a chance that the selected patches will not match the patches in the disaster recovery media.

- **Re-installing from a CD-ROM**

You must insert the *Sidewinder Installation CD* into the firewall. This option is less flexible than restoring from the VCD, but requires less interaction with the firewall. When using the physical CD, you can automatically install all previously installed patches using disaster recovery media.

See [Re-installing your firewall from a CD-ROM](#).

- **Re-installing from a USB drive**

You must insert the *Sidewinder Installation USB Drive* into the firewall. This option is less flexible than restoring from the VCD, but requires less interaction with the firewall. When using the physical USB drive, you can automatically install all previously installed patches using disaster recovery media.

See [Re-installing your firewall from a USB drive](#).

Before re-installing, determine which process can best create the final configuration that meets your needs. See the following table for options.

Table 78 Re-install options

Re-install process	Recovery options
Re-install from the virtual CD	<ul style="list-style-type: none">• Load or create an initial configuration and restore a configuration backup.• Load or create an initial configuration and create a new policy.
Re-install from the physical CD or USB drive	<ul style="list-style-type: none">• Restore from a disaster recovery backup created using the Create Disaster Recovery Backup option. You may also want to restore a recent configuration backup.• Load or create an initial configuration, download and install the necessary patches, and restore a configuration backup.• Load or create an initial configuration, and create a new policy.

Re-installing your firewall from the virtual CD

Use this procedure to re-install your Sidewinder from the virtual CD, which contains a copy of all installed patches. Re-installing from the VCD is appropriate when the firewall needs to be repurposed or needs to recover from a software or configuration disaster, and the hard drive is still functional. This option gives you the opportunity to select which packages to install.

Note the following:

- This method does not require installation media.
- Plan your restore method before you begin. Options include:
 - Load or create an initial configuration, and create a new policy.
 - Load or create an initial configuration and restore a configuration backup. Know which patches need to be installed to match the configuration backup you intend to restore.
- If you are using a serial terminal for this procedure, use the hot keys to make selections. Using the arrow keys may have unexpected results.

To re-install from the virtual CD:

- 1 Power on or reboot the firewall.

After the standard boot information completes, the menu for the virtual CD appears.

- 2 Press **F1** to enter the Virtual CD.

- 3 Select a Sidewinder boot option:

- To accept the default installation option, press **Enter**.
- If you intend to use a serial console, type **4** and press **Enter**.

Note: The **Emergency Maintenance Mode** option, the **Boot with ACPI** option, and the **Escape to Loader Prompt** option are generally only to be used on instruction from Technical Support and are not appropriate for this procedure.

The VCD Sidewinder Main Menu appears, displaying the following options:

```
Select and Install - Select and install Sidewinder Packages
Maintenance Shell - VCD Maintenance Shell Menu
Advanced Menus - Advanced Install Menus
[RTN To select] X Exit Install
```

- 4 Press **S** to select `Select and Install`, and then press **Enter**.

A menu appears listing all available packages.

- 5 [Conditional] If selecting packages:

- Use the hot keys or the up and down arrows to move among the packages. (If using a serial connection, use only the hot keys.)
- Use the space bar to select a package or clear a previous selection.
- Press **O** for OK when ready to install the selected packages.
- Press **C** to cancel and return to the previous menu.

- 6 [Conditional] After pressing **O** for OK, you are prompted to confirm your selection. Press **Y**.

The packages install. When installation is complete, the VCD Main Menu appears again.

- 7 Press **X** to exit. You are prompted to confirm your decision and reminded to remove all media from the floppy and CD drives.
- 8 Press **Y**. The firewall reboots.
- 9 At the Forcepoint menu, accept the default, which is the Operational System.
- 10 At the boot menu, select a boot method.
- 11 Provide the appropriate initial configuration using one of these methods:
 - Insert a USB drive containing a disaster recovery backup. (See [Disaster recovery](#) for more information.)
License information is included in the backup file.
 - Run your chosen Quick Start method. (See the *Setup Guide* for more information.)
The firewall tries to send the license activation request to an Internet activation server for one minute. If the activation is not successful in that time, you must activate your firewall using the Admin Console.
- 12 Connect to your firewall using the Admin Console.
- 13 [Conditional] If you need to restore a configuration backup, select **Maintenance > Configuration Backup** and restore your firewall configuration data. See [Configuration file backup and restore](#) for more information.

Your firewall is now re-installed.

Re-installing your firewall from a CD-ROM

Use this procedure to re-install your Sidewinder from the *Sidewinder Installation CD*. This procedure should only be used as a last resort, such as when the firewall has had a hardware failure or has been completely replaced. Generally, you should only run this procedure under guidance from Technical Support.

Note the following:

- Plan your restore method before you begin. Options include:
 - Restore a disaster recovery backup from a USB drive. You may also want to restore a recent configuration backup.
 - Load or create an initial configuration, download and install the necessary patches, and restore a configuration backup.
 - Load or create an initial configuration, and create a new policy.
- If your firewall does not have a CD-ROM drive, see [Re-installing your firewall from a USB drive](#).
- You will need your *Sidewinder Installation CD*.

To re-install from a CD:

- 1 Boot the firewall from the *Sidewinder Installation CD*:
 - If the firewall is on, insert the CD and then reboot.
 - If the firewall is off, power it on and then quickly insert the CD.The firewall boots from the CD and displays standard boot-up information.
- 2 [Conditional] By default, the boot order is set to check the CD drive first. If the boot order has been altered and does not check the CD drive first, reboot and enter the BIOS to adjust the boot order accordingly.
- 3 At the Forcepoint menu, accept the default, which is the Operational System. The boot menu appears.
- 4 Select a Sidewinder boot option:
 - To accept the default installation option, press **Enter**.
 - If you intend to use a serial console, type **4** and press **Enter**.

Note: The **Emergency Maintenance Mode** option, the **Boot with ACPI** option, and the **Escape to Loader Prompt** option are generally only to be used on instruction from Technical Support and are not appropriate for this procedure.

The firewall continues starting and then `Auto install will start in 10 seconds, ^C to interrupt` appears.

- If you want to manually select which firewall packages to install, press **Ctrl+c**. The Sidewinder Main Menu appears. Continue with [Step 5](#).
 - If you want to install all of the firewall packages that are included on the *Sidewinder Installation CD*, take no action. The firewall loads and installs the Sidewinder packages. Continue with [Step 6](#).
- 5** [Optional] Manually select which firewall packages to install:
- a** Use the arrow keys to select **Select and Install** and then press **Enter**. The Select Sidewinder Packages menu appears.
 - b** Use the arrow keys and the spacebar to select or clear packages for installation.
 - c** When you are finished making your selections, press **Tab** to select **OK** and then press **Enter**. A confirmation prompt appears.
 - d** Select **Yes** and then press **Enter**.
- The firewall loads and installs the firewall packages.
- 6** When the "Installation complete" message appears, remove the CD from its drive.
- 7** Press **R** to reboot the firewall, and then press **Enter**. The firewall boots and displays standard boot-up information.
- 8** At the Forcepoint menu, accept the default, which is the Operational System.
- 9** At the boot menu, select a boot method.
- 10** Provide the appropriate initial configuration using one of these methods:
- Insert a USB drive containing a disaster recovery backup into one of the firewall's USB ports. License information is included in the backup file.
 - Run your chosen Quick Start method. (See the *Setup Guide* for more information.)
- The firewall tries to send the license activation request to an Internet activation server for one minute. If the activation is not successful in that time, you must activate your firewall using the Admin Console. If the system cannot retrieve its license key, the firewall comes up with a temporary license. You must obtain your license within 7 days.
- 11** Connect to your firewall using the Admin Console.
- 12** [Conditional] If you need to install additional patches, select **Maintenance > Software Management**. See [Software management](#) for more information.
- 13** [Conditional] If you need to restore a configuration backup, select **Maintenance > Configuration Backup** and restore your firewall configuration data. See [Configuration file backup and restore](#) for more information.

Your firewall is now re-installed.

Re-installing your firewall from a USB drive

Use this procedure to re-install your Sidewinder from the *Sidewinder Installation USB Drive*. This procedure should only be used as a last resort, such as when the firewall has had a hardware failure or has been completely replaced. Generally, you should only run this procedure under guidance from Technical Support.

Note the following:

- Plan your restore method before you begin. Options include:
 - Restore a disaster recovery backup from a USB drive. You may also want to restore a recent configuration backup.
 - Load or create an initial configuration, download and install the necessary patches, and restore a configuration backup.
 - Load or create an initial configuration, and create a new policy.
- If your firewall has a CD-ROM drive, see [Re-installing your firewall from a CD-ROM](#).
- You will need your *Sidewinder Installation USB Drive*.

To re-install from a USB drive:

1 Insert the *Sidewinder Installation USB Drive* into one of the firewall's USB ports.

2 Configure your firewall to boot from USB.

- a Restart or power on the firewall and adjust the boot order accordingly.

Note: Boot options vary by firewall model.

- b If necessary, save your changes.

The firewall boots from the USB drive and displays standard boot-up information.

3 At the Forcepoint menu, accept the default, which is the Operational System. The boot menu appears.

4 Select a Sidewinder boot option:

- To accept the default installation option, press **Enter**.
- If you intend to use a serial console, type **4** and press **Enter**.

Note: The **Emergency Maintenance Mode** option, the **Boot with ACPI** option, and the **Escape to Loader Prompt** option are generally only to be used on instruction from Technical Support and are not appropriate for this procedure.

The firewall continues starting and then `Auto install` will start in 10 seconds, `^C` to interrupt appears.

- If you want to manually select which firewall packages to install, press **Ctrl+c**. The Sidewinder Main Menu appears. Continue with [Step 5](#).
- If you want to install all of the firewall packages that are included on the *Sidewinder Installation USB Drive*, take no action. The firewall loads and installs the Sidewinder packages. Continue with [Step 6](#).

5 [Optional] Manually select which firewall packages to install.

- a Use the arrow keys to select **Select and Install** and then press **Enter**. The Select Sidewinder Packages menu appears.
- b Use the arrow keys and the spacebar to select or clear packages for installation.
- c When you are finished making your selections, press **Tab** to select **OK** and then press **Enter**. A confirmation prompt appears.
- d Select **Yes** and then press **Enter**.

The firewall loads and installs the firewall packages.

6 When the "Installation complete" message appears, remove the *Sidewinder Installation USB Drive* from your firewall.

7 Press **R** to reboot the firewall, and then press **Enter**. The firewall boots and displays standard boot-up information.

8 At the Forcepoint menu, accept the default, which is the Operational System.

9 At the boot menu, select a boot method.

10 Provide the appropriate initial configuration using one of these methods:

- Insert a USB drive containing a disaster recovery backup into one of the firewall's USB ports.
License information is included in the backup file.
- Run your chosen Quick Start method. (See the *Setup Guide* for more information.)

The firewall tries to send the license activation request to an Internet activation server for one minute. If the activation is not successful in that time, you must activate your firewall using the Admin Console. If the system cannot retrieve its license key, the firewall comes up with a temporary license. You must obtain your license within 7 days.

11 Connect to your firewall using the Admin Console.

12 [Conditional] If you need to install additional patches, select **Maintenance > Software Management**. See [Software management](#) for more information.

Re-install and Recovery Options

Re-install options

13 [Conditional] If you need to restore a configuration backup, select **Maintenance > Configuration Backup** and restore your firewall configuration data. See [Configuration file backup and restore](#) for more information.

Your firewall is now re-installed.

Glossary

A

activation key A string of numbers and characters that allows the operation of the software.

Admin Console The graphic user interface (GUI) used to configure and manage software.

B

burb A set of one or more interfaces and the group of systems connected to each interface that are to be treated the same from a system security policy point of view.

C

client A program or user that requests network service(s) from a server.

CGI (common gateway interface) Any server-side code that accepts data from forms via HTTP. The forms are generally on Web pages and submitted by end users.

D

daemon A software routine within UNIX that runs in the background, performing system-wide functions.

domain (1) Relative to networking, the portion of an Internet address that denotes the name of a computer network. For instance, in the e-mail address *jones@example.sales.com*, the domain is *example.sales.com*. (2) Relative to Type Enforcement, an attribute applied to a process running on SecureOS that determines which system operation the process may perform.

DNS (domain name system) A TCP/IP service that maps domain and host names to IP addresses. A set of connected name servers and resolvers allows users to use a host name rather a 32-bit internet address.

E

editor A program that can be used to create or modify text files. See also *File Editor*.

external DNS External DNS provides a limited external view of the organizational domain. No internal information is available to the external DNS and only the external DNS can communicate with the outside. Therefore, no internal naming information can be obtained by anyone on the outside. The external DNS cannot query the internal DNS or any other DNS server inside a Sidewinder.

F

failover See *High Availability*.

failure mode Also known as safe mode, a Sidewinder operating state that allows system administration while not allowing network traffic to pass through. A firewall can enter this mode after such conditions as: (a) a failed license check, (b) a reboot during which the system detects a problem with an installed patch, (c) a reboot during which the system failed to start a critical service, or (d) the audit partition has overflowed.

File Editor The program available directly in the Admin Console that can be used to create or modify text files. The File Editor communicates with the Sidewinder using a secured connection.

FreeBSD The operating system used as a base for developing SecureOS. See also *SecureOS*.

G

gateway A network component used to connect two or more networks that may use dissimilar protocols and data transmission media.

H

High Availability (HA) A feature that allows a second Sidewinder to be configured either in a load sharing capacity or in "hot backup" (secondary or standby) mode.

host Any computer connected to a network, such as a workstation, router, firewall, or server.

I

ICANN (Internet Corporation for Assigned Names and Numbers) An organization that oversees IP addresses and domain names on the Internet. The goal of ICANN is to ensure that all public IP addresses are unique and that valid addresses are accessible to all Internet users.

internal DNS Manages information only available to internal machines. The internal name server cannot receive queries from external hosts since it cannot communicate directly with the external network. Although it is unable to communicate directly with external hosts, it is able to send queries and receive the responses via the external DNS.

Internet A worldwide system of computer networks. A public, cooperative, and self-sustaining network of networks accessible to hundreds of millions of people worldwide. The Internet is technically distinguished because it uses the TCP/IP set of protocols.

IPv4 address A 32-bit address assigned to TCP/IP network devices. An IP address is unique to each machine on the Internet.

IPv6 Internet Protocol version 6. A replacement for the aging IPv4, which was released in the early 1980s. IPv6 will increase the number of available Internet addresses (from 32 to 128 bits), resolving a problem associated with the growth of the number of computers attached to the Internet.

M

mail server A network computer that serves as an intermediate station for electronic mail transfers.

MAT (multiple address translation) The ability for a single interface to support multiple external IP addresses so that inbound connections can be directed based on IP addresses and service. MAT allows proxies to be directed to different destinations for the same service by the IP address to which it was connected.

man page Short for manual page, refers to the online help that is available within the UNIX operating system. For example, entering `man ls` at the UNIX prompt displays a description of the UNIX `ls` command.

MX (mail exchanger) records DNS entries that define where e-mail addresses within domain names get delivered.

N

name server A network computer that maintains a relationship between IP addresses and corresponding domain names.

NAS (network access server) A system that provides dial-up connectivity to your IP network from a bank of dial-up modems.

NAT (network address translation) Changing the source address of a packet to a new IP address specified by the administrator.

net mask The way that computers know which part of a TCP/IP address refers to the network, and which part refers to the host range.

NIC (network interface card) Hardware, like a computer circuit board, that contains a port or jack that enables a computer to connect to network wiring (for example, ethernet cable, phone line, etc.).

O

operational kernel The SecureOS kernel that provides the normal operating state, including Type Enforcement controls. When this kernel is running, the firewall can connect to both the Internet and the internal network, and all configured services are operational.

P

ping A command that sends an ICMP message from one host to another host over a network to test connectivity and packet loss.

port The number that identifies the destination application process for transmitted data. Port numbers range from 1 to 65535. (For example, Telnet typically uses port 23, DNS uses 53, etc.)

primary name server The DNS server for a domain where the name information is stored and maintained.

protocol A set of rules by which one entity communicates with another, especially over a network. This is important when defining rules by which clients and servers talk to each other over a network. Important protocols become published, standardized, and widespread.

proxy A software agent that acts on behalf of a user requesting a network connection through the firewall. A proxy accepts a connection from a user, make a decision as to whether or not the user or client IP address is permitted to use the proxy, optionally does additional authentication, and then completes a connection on behalf of the user to a remote destination.

Q

Quick Start Wizard A Windows-based program that allows you to create an initial configuration for your Sidewinder.

R

RAID (redundant array of individual disks) Information is stored on multiple hard disks to provide redundancy. Using RAID can improve performance and fault-tolerance.

registration The process of authenticating one Sidewinder system to an HA cluster. This process establishes an encrypted, trusted connection between the two systems.

registration key Character string used for authentication during the registration process.

router A network device that forwards data between two or more networks, delivering them to their final destination or to another router.

S

safe mode See *failure mode*.

secondary name server DNS server that downloads and records a backup copy of domain information from a primary DNS server.

SecureOS The UNIX-based operating system used in Sidewinder. SecureOS is built upon FreeBSD and includes Type Enforcement security mechanisms.

server A computer system that provides services (such as FTP) to a network, or a program running on a host that offers a service to other hosts on a network.

SMTP (simple mail transport protocol) The TCP/IP protocol that transfers e-mail as it moves through the system.

subnet A network addressing scheme that separates a single network into a number of smaller physical networks to simplify routing.

T

Telnet A TCP/IP protocol that directs the exchange of character-oriented data during a client-to-server session.

Type Enforcement Security technology that protects against intruders by preventing someone from taking over the UNIX operating system within Sidewinder and accessing critical files or doing other damage.

U

URL (universal resource locator) Indicates the address of specific documents on the Web. Every Internet file has a unique URL indicating the name of the server, the directory, and the specific document.

USB drive A portable flash memory card that plugs into the computer's USB port.

Index

A

- A record (address record) [485, 488](#)
- acat_acls [681](#)
- accept certificate [35](#)
- Access Table mail file [509](#)
- account
 - administrator [44](#)
 - changing password [47](#)
- activation process [586](#)
- Active Directory [91](#)
- active policy [272](#)
- active session mode [85, 87](#)
- add-on modules
 - anti-virus [134](#)
 - IPS [118](#)
 - SSL decryption [201](#)
- address pools [521](#)
- Admin Console
 - about [29](#)
 - adding a firewall [34](#)
 - administration options [30](#)
 - exiting [32](#)
 - File Editor [606](#)
 - logging in [35](#)
 - server [173, 177](#)
 - setting system date and time [568](#)
 - starting [31](#)
 - tips when using [33](#)
 - valid port values [340](#)
- admin role
 - file access [19](#)
 - tasks [44](#)
- administration
 - remote via SSH [48](#)
 - remote via telnet [54](#)
- administration tool [29](#)
- administrator
 - account [44](#)
 - cautions when editing UNIX files [606](#)
- administrator accounts
 - monitoring activity [684](#)
- adminro role [44](#)
- Admn domain [19](#)
- agents
 - packet filter [163](#)
 - proxy [154, 178](#)
 - server [172](#)
- alarms see [IPS Attack Responses and System Event Responses](#)
 - see
- algorithms with VPN [544](#)
- alias
 - mail [514](#)
 - NAT [263, 275, 283, 286](#)
 - redirection [264](#)
 - root [514](#)
- allow traffic [274](#)
- allowed get communities [387](#)
- allow-query option [479, 482](#)
- allow-transfer option [480, 482](#)
- allow-update option [482](#)
- anomaly detection [307](#)
 - see also [attacks](#)
- anti-relay controls [502](#)
- Anti-virus filtering
 - see [alsovirus scanning](#)
 - for FTP [230–233](#)
 - for Mail [217](#)
 - for Web [206](#)
 - overview [502](#)
- aol proxy [160](#)
- Application Defense groups [60](#)
- Application Defenses
 - about [22](#)
 - Citrix [227](#)
 - FTP [228–233](#)
 - groups [254](#)
 - H.323 [236](#)
 - HTTP and HTTPS [198–212](#)
 - IIOp [234](#)
 - in rules [277](#)
 - Mail (sendmail) [212–220](#)
 - Oracle [237](#)
 - Packet Filter [252](#)
 - selecting in rules [287](#)
 - SIP [244](#)
 - SMTP proxy [220–226](#)
 - SNMP [242](#)
 - SOCKS [240](#)
 - SSH [247–251](#)
 - T.120 [235](#)
- application notes [13](#)
- ARP [32](#)
 - gratuitous [653](#)
 - Network Defense [375](#)
- ARP table [686](#)
- attack audits
 - ICMP [373](#)
 - IP [369](#)

- IPsec [377](#)
- IPv6 [379](#)
- TCP [367](#)
- UDP [371](#)
- attack protection tools [22](#), [121](#)
- attack responses [347](#)
- attacks by service [309](#)
- audit
 - *.gz files [314](#)
 - *.raw files [314](#)
 - about [311–314](#)
 - adjusting [274](#), [314](#), [363](#)
 - attack responses [347](#)
 - change tickets [332](#), [336](#), [579](#), [583](#)
 - configuring [347](#)
 - dashboard [307–310](#)
 - events [312](#)
 - IPS Attack Responses [347](#)
 - levels in rules [274](#)
 - sending SNMP traps [360](#), [383](#)
 - sending to syslog [337](#)
 - SNMP traps [383](#)
 - System Event Responses [347](#)
 - understanding messages [314](#)
 - viewing [308](#)
 - viewing a service's audit errors [345](#)
- audit data
 - filtering [318–325](#)
- Audit Management
 - Audit Options pane [331](#)
 - Crontab Editor [334](#)
 - export entry [332](#)
 - exporting log files [334](#)
 - exporting on request [335](#)
 - Logfile Options pane [332](#)
 - rolling log files [334](#)
 - rolling on request [335](#)
 - schedule export [334](#)
 - schedule rolling [334](#)
 - signing export files [334](#)
 - statistics [332](#)
 - tab [331](#)
- Audit Options pane [331](#)
- audit records [310](#)
 - exporting [328](#)
- Audit Viewing [316](#)
 - predefined filters [322](#)
- audit.raw file [312](#), [495](#)
- auditbotd [312](#)
- auditd [312](#)
- auditdbd [312](#)
- authentication
 - alternate methods [81](#)
 - clear locks [683](#)
 - Common Access Card [95](#)
 - LDAP [91](#)
 - Passport [85](#)
 - password [90](#)
 - RADIUS [97](#), [99](#)
 - SafeWord [100](#)
 - selecting in rules [277](#), [289](#)
 - SNMP message header [383](#)
 - SSH login [49](#)
 - Windows Domain [96](#), [97](#)
 - with VPN [518](#)
- authentication failure lockout [114](#), [683](#)
- authentication failure traps [386](#)
- auto-recover on reconnect [669](#)
- B**
- backup
 - configuration files [575–577](#)
 - disaster recovery [578](#)
- BGP
 - configuring [453–456](#)
- bgpd
 - server [173](#), [177](#)
- bibliography [13](#)
- bi-directional packet filter rules [171](#)
- binary characters [216](#)
- BIND [473](#)
- blackhole [353](#)
 - as IPS Attack Response [352](#)
 - with signature-based IPS [118](#)
- blackhole list [507](#)
- blackholed IPs
 - managing [300](#), [303](#)
 - TrustedSource and [144](#)
 - with signature-based IPS [124](#)
- Boot with ACPI [709](#), [710](#), [712](#)
- broadcast address [551](#)
- burb groups
 - about [59](#)
- burbs
 - configuring [393–396](#)
 - network stack separation [20](#)
 - viewing services enabled in [343](#)
- C**
- ccmd server [177](#)
- certificate accept window [35](#)
- Certificate Authority (CA)
 - adding [631](#)
 - defined [522](#)
 - exporting [631](#)
 - managing [629](#)
 - public versus private [622](#)
 - querying [624](#), [639](#)
 - retrieving [630](#)
- certificate server [634](#)
- certificates

- Admin Console [628](#)
- exporting [643](#)
- firewall [623](#)
- importing [626](#), [640](#)
- managing [619](#)
- remote [638](#)
- VPN [633](#)
- change password server [103](#), [172](#), [177](#)
- change ticket [332](#), [335](#)
- changing admin password [47](#)
- chtype command [611](#)
- Citrix
 - Application Defense [227](#)
 - proxy (ICA) [161](#)
- class types [118](#), [122](#), [124](#)
- client address pools [521](#)
- CNAME record [488](#), [489](#)
- command line interface [29](#), [30](#)
- commands
 - dig [694](#)
 - mail queues [511](#)
 - netstat [693](#)
 - ping [693](#)
 - process [684](#)
 - route [694](#)
 - showaudit [314](#), [703](#)
 - tcpdump [701](#), [703](#)
 - top [684](#)
 - traceroute [693](#)
 - uptime [684](#)
 - vmstat [684](#)
- Common Acces Card authentication [95](#)
- community names [383](#)
- compare configuration backups [579](#), [583](#)
- configuration
 - auditing [347](#)
 - DNS [469](#), [471](#)
 - files [606](#)
 - mail host [500](#)
 - Strikeback [347](#)
- configuration backup
 - Lite [332](#)
- configuration backups
 - compare [579](#), [583](#)
- configuration files
 - backing up [575–577](#)
 - restoring [581–582](#)
- connection type [159](#)
- conventions used in this guide [14](#)
- copying rules [273](#)
- CPU use tab [299](#)
- CPU, time by process [684](#)
- creating an alternate policy [268](#)
- CRL [630](#)
- Crontab Editor window [336](#)
- Crontab Editor, Audit Management [334](#)
- Custom LDAP [91](#)
- customizing the Rules window [271](#)
- D**
- daemon [339](#)
- dashboard
 - about [293–294](#)
 - audit [307–310](#)
 - device information [295–298](#)
 - HA management [666](#)
 - monitord [293](#)
 - network traffic [300–306](#)
 - summary of statistics [307–310](#)
- date (setting) [568](#)
- decryption [518](#)
- default route [428](#)
 - IPv6 [433](#)
- default route failover [430](#)
- Deny All rule [267](#), [268](#)
- deny traffic [274](#)
- Deployment options [23–28](#)
- destination burb [275](#)
- DHCP [408](#)
- DHCP relay agent
 - configuring [174–176](#)
 - server [173](#)
- dig command [694](#)
- disaster recovery
 - creating a backup [578](#), [706](#)
 - creating an alternate policy [268](#)
 - re-installing [707](#)
- disk use information [298](#), [684](#)
- Distinguished Names [619–621](#)
- DNS
 - A record (address record) [485](#), [488](#), [489](#)
 - active rules and [682](#)
 - advanced server options [479](#)
 - advanced zone options [482](#)
 - BIND [473](#)
 - CNAME record [488](#), [489](#)
 - configuration [469](#), [471](#), [474](#), [475](#)
 - configuration utility [490](#)
 - disabling servers [472](#)
 - editing configuration files [469](#)
 - enabling servers [472](#)
 - file types [495](#)
 - Firewall Hosted [471](#)
 - Firewall hosted [475](#)
 - forward zones [481](#)
 - forwarders [477](#)
 - HINFO [488](#), [489](#)
 - hosts [487](#)
 - if turned off [472](#)
 - logging [495](#)
 - mail exchanger records [486](#)

- master zone [481](#)
- master zone attributes [482](#)
- master zone contents [487](#)
- MX record [473](#), [488](#), [489](#)
- name servers table [486](#)
- proxy [160](#)
- query [473](#)
- reconfigure [490](#)
- reverse zones [481](#)
- rules [268](#)
- serial number [484](#)
- servers for VPNs [549](#)
- slave zone [481](#)
- SOA record [483](#)
- split DNS mode [471](#), [472](#)
- sub-domain [486](#)
- transparent [470](#), [474](#)
- TTL value [484](#)
- zone [480](#)

- DNS lookup [32](#), [688](#)
- documentation [13](#)
- domain definition table [18](#)
- domain object [65](#)
- domains
 - access [18](#)
 - Admn [19](#)
 - creator [611](#)
 - defined [18](#)
 - file access [18](#)
 - mail [500](#)
- drop traffic [274](#)
- duplicating rules [273](#)
- dynamic IP addressing
 - Adding a new VPN [537](#)

E

- editing UNIX files [606](#)
- editors
 - Admin Console File Editor [606](#)
- Emergency Maintenance Mode [42](#), [709](#), [710](#), [712](#)
- encrypted files, scanning for viruses [138](#)
- encryption, with VPN [518](#)
- endpoint
 - destination [276](#)
 - in rules [282–286](#)
 - source [275](#)
- entrelayd [172](#), [177](#)
- Escape to Loader prompt [709](#), [710](#), [712](#)
- etc/sidewinder.daemon.conf [339](#)
- event analysis [293](#)
- executables, installing [18](#)
- expected connections [158](#)
- export entry, Audit Management [332](#)
- Export File window [332](#)
- exporting audit records [328](#)
- exporting certificates [643](#)

- exporting log files [334](#)
 - signature [334](#)
- Extended Authentication [524](#)

F

- failed connection request [679](#)
- failover see High Availability
 - see
- failure mode see safe mode
 - see
- fast path sessions [155](#), [159](#)
- Federal Information Processing Standard [614](#)
- File Editor [606–610](#)
- file permissions [611](#)
- file size range, virus scanning [137](#)
- file type
 - checking [611](#)
 - DNS files [495](#)
- files
 - configuration [606](#)
- filtering
 - mail [213](#)
 - Web [206](#)
- filtering audit data [318–325](#)
- filters
 - access control [21](#)
 - global properties [149](#)
 - see packet filter
- finger proxy [160](#)
- FIPS [614](#)
- firewall certificate
 - creating [624](#)
 - managing [623](#)
- firewall certificates
 - importing [626](#)
- Firewall Hosted DNS [471](#)
- firewall license [586](#)
- Firewall Policy Report [61](#)
- fixed IP [551](#)
- forward zones [481](#)
- FTP proxy
 - agent [160](#)
 - command filtering [229](#)
 - server response configuration [187](#)
 - virus/spyware filtering [230](#)
- fwregisterd server [172](#), [177](#)
- fwregisterp proxy [160](#)

G

- gatekeeper [190–192](#)
- general system information [293](#)
- Geo-Location
 - database updates [146](#)
 - database version [146](#)
 - e-mail notification [146](#)
 - network objects [66](#)

Get route [687](#)
 get route [32](#)
 global properties [149](#)
 gopher proxy [160](#)

H

H.323
 Application Defense [236](#)
 gatekeeper [190–192](#)
 NetMeeting [192](#)
 overview [189](#)
 proxy [160](#)
 HA
 ping [658](#)
 HA see High Availability
 see
 halt system [42](#)
 halting services [341–345](#)
 header stripping [507](#)
 heartbeat burb [649](#), [654](#)
 heartbeat verification burb [668](#)
 help (online) [13](#)
 High Availability
 cluster addresses [650](#)
 cluster interface properties [670](#)
 cluster tree structure [665](#)
 Cluster Wizard [655](#)
 configuration options [651](#)
 configuring [653](#)
 heartbeat burb [649](#), [654](#)
 layer 2 modes [652](#)
 load-sharing [651](#)
 load-sharing soft shutdown [675](#)
 managing an HA cluster [667](#)
 packet filter stateful session failover [168](#), [171](#)
 peer-to-peer [653](#)
 primary-standby [653](#)
 redundancy [650](#)
 requirements [654](#)
 restarting a cluster [676](#)
 SNMP considerations [385](#)
 troubleshooting [696–700](#)
 HINFO [488](#), [489](#)
 Host Enrollment List [594](#)
 host name [295](#)
 host object [67](#)
 Host Resources MIB [385](#)
 hosted DNS
 on firewall [475](#)
 single [471](#)
 split server [471](#)
 HTTP and HTTPS Application Defenses [198–212](#)
 HTTP proxy
 agent [161](#)
 URL translation [178–181](#)
 HTTPS proxy [161](#)

Hybrid mode [28](#)

I

ICA (Citrix) proxy [161](#)
 ICMP
 in burbs [395](#)
 Network Defense [373](#)
 proxy [161](#)
 ICMP message types [171](#)
 ident proxy [161](#)
 IDS [118](#), [124](#)
 signature groups [127](#)
 IETF [518](#)
 IIOp Application Defense [234](#)
 IIOp proxy [161](#)
 IKE [518](#)
 IKEv1 [518](#), [542](#)
 IKEv2 [518](#), [542](#)
 imap proxy [161](#)
 importing certificates
 firewall [626](#)
 remote [640](#)
 in-addr-arpa [481](#)
 inbound rules [58](#)
 inetd [339](#), [340](#)
 inoffensive,TrustedSource [144](#)
 Installation CD [709](#)
 installing executables [18](#)
 interface properties [418](#)
 interface status
 ping [302](#)
 viewing in the dashboard [301](#)
 interfaces report [693](#)
 interfaces, configuring [397–418](#)
 Internet Key Exchange [518](#)
 Internet server [472](#)
 intra-burb forwarding [169](#)
 IP address object [68](#), [69](#)
 IP Network Defense [369](#)
 IP sniffing [17](#)
 IP spoofing [17](#)
 iPlanet [91](#)
 IPS
 about [22](#)
 about signature-based inspection [117–120](#)
 and IPS Attack Responses [121](#)
 response mappings [122–124](#)
 selecting in rules [120](#), [277](#), [288](#)
 signature file updates [121](#), [132](#)
 signature groups [125–127](#)
 IPS Attack Responses
 about [22](#), [347](#)
 and signature-based IPS [121](#)
 attack descriptions [350](#)
 ignoring network probe attempts [359](#)
 managing [348–352](#)

- response settings [353](#)
- TrustedSource and [144](#)
- IPsec
 - defined [518](#)
 - Network Defense [377](#)
 - troubleshooting [703](#)
- IPv6
 - default route [433](#)
 - Network Defense [379](#)
- irc proxy [161](#)
- Ironmail proxies [161](#)
- ISAKMP
 - server [177](#)
- ISAKMP server [173](#), [533](#), [552](#)
- K**
- keys
 - managing [619](#)
- keys (SSH) [645](#)
- keys (VPN)
 - defined [518](#)
 - encryption and decryption [518](#)
 - generating [518](#)
- keyword search [502](#)
- Knowledge Base [13](#)
- L**
- LDAP
 - authentication [91](#)
 - Certificate Authorities [634](#)
 - proxy [161](#)
- license
 - Host Enrollment List [594](#)
 - how to [586](#)
 - troubleshooting [695](#)
- listens [693](#)
- Lite configuration backup [332](#)
- load average [299](#)
- load-sharing HA [651](#)
- locked out [682–683](#)
- log in, Admin Console [35](#)
- logcheck [311](#)
- Logfile Options pane [332](#)
- Logfile Options, toolbar [332](#)
- logging DNS [495](#)
- Login Console [172](#), [177](#)
- loglevel [680](#)
- loopback address [481](#)
- lotus proxy [161](#)
- ls -yalZ command [611](#)
- M**
- M4 Config File [513](#)
- M4 config file [506–510](#), [512](#)
- m4 macros [502](#)
- MAC address [686](#)
- mail
 - advanced configuration [506–510](#)
 - aliases [514](#)
 - configuration basics [502–505](#)
 - domains [500](#)
 - internal server [500](#)
 - local delivery [501](#)
 - local program [500](#)
 - local server [500](#)
 - mailertables [505](#)
 - postmaster [500](#)
 - reconfiguring [498](#)
 - redirecting [514](#)
 - rules for transparent [499](#)
 - servers [500](#)
 - SMTP [498](#)
 - SMTP hosted [497](#)
 - transparent SMTP [497](#)
 - Type Enforcement restrictions [501](#)
 - viewing mail on the firewall [514](#)
- mail exchanger records [473](#), [483](#), [485](#), [486](#)
- mail filtering
 - keyword search filter [213](#)
 - MIME/Anti-Virus filter [213](#)
 - size filter [213](#), [214](#)
- mail host [500](#)
- mail host configuration [500](#)
- mail queues
 - about [501](#)
 - managing [512–513](#)
 - viewing [511](#)
- mailertable files [505](#)
- management information base (MIB) [383–385](#)
- manuals [13](#)
- masquerade_domain [510](#)
- master zone
 - about [481](#)
 - attributes [482](#)
 - contents (DNS) [487](#)
- messages
 - audit [314](#)
 - DNS [495](#)
 - in mail queues [511](#)
 - postmaster [500](#)
- mib2 [384](#)
- MIME filtering
 - for mail [217](#)
 - for Web [206](#)
- monitord [293](#)
- monitoring
 - attacks [307](#), [308](#)
 - network traffic [300](#), [306](#)
 - system events [307](#)
 - system resources [295–302](#)
 - system status [294](#)
 - VPN status [304](#), [535](#)

msn proxy [161](#)
 mssql proxy [161](#)
 mta domains [500](#)
 mtac domain [500](#)
 Multicast Group Address [668](#)
 multiple instantiations of a proxy agent [158](#)
 MX record [473](#), [488](#), [489](#)

N

name servers table [486](#)
 name servers, configuring [471](#)
 NAT
 about [263](#)
 in packet filter rules [166](#)
 selecting in rules [275](#)
 NAT Traversal (NAT-T) [544](#)
 NetBIOS proxy [161](#)
 netmap
 defining [70](#)
 member [70](#)
 netstat
 listens [693](#)
 load-sharing [700](#)
 Network Defenses
 about [22](#), [363–366](#)
 ARP [375](#)
 ICMP [373](#)
 IP [369](#)
 IPsec [377](#)
 IPv6 [379](#)
 TCP [367](#)
 UDP [371](#)
 network interfaces [397](#)
 network interfaces report [693](#)
 network objects
 creating [63](#)
 network probe attempts [359](#)
 networks
 interfaces report [693](#)
 process status [684](#)
 routing tables [693](#)
 services [340](#)
 stack separation [20](#)
 neutral, TrustedSource [144](#)
 NIC
 restart [301](#)
 NIC group [418](#)
 NNTP proxy [161](#)
 non-transparent
 connection type [159](#)
 rules for non-transparent mail [499](#)
 non-transparent proxies [155](#)
 notify option [479](#), [482](#)
 NSS [340](#)
 NTP
 about [569–572](#)

proxy [161](#)
 reasons for having stopped [702](#)
 references [572](#)
 restarting [702](#)
 servers and clients [570](#)
 troubleshooting [700](#)
 version number [569](#)

O

OID, editing [243](#)
 online help [13](#)
 OpenLDAP [91](#)
 Oracle
 Application Defense [237](#)
 proxy [161](#)
 ordering your rules [59](#)
 OSPF
 configuring [446–453](#)
 overview [445–446](#)
 processing [446–453](#)
 ospf
 server [173](#), [177](#)
 ospf IPv6
 server [173](#), [177](#)
 outbound rules [58](#)

P

packages
 downloading [600](#)
 installing [597](#), [600](#)
 loading [597](#)
 scheduling installs and uninstalls [601](#)
 types [597](#)
 uninstall [598](#), [601](#)
 viewing available [600](#)
 viewing current version [295](#)
 packet filter
 agent properties [168](#)
 agents [168](#)
 agents and services [163](#)
 High Availability stateful session failover [168](#)
 intra-burb forwarding [169](#)
 NAT [166](#)
 preserving source port [167](#)
 redirection [166](#)
 reserved port range [169](#)
 service properties [171–172](#)
 session maximums [169](#)
 stateful packet inspection [163–166](#)
 Packet Filter Application Defense [252](#)
 packet filter rules, direction [171](#)
 packet filters
 see alsofilters
 Passport [85](#), [??–91](#), [268](#)
 active session mode [85](#), [87](#)
 redirect delay [87](#)

- Passport authenticator
 - server [173](#), [178](#)
 - password
 - authentication [90](#)
 - changing [47](#), [102](#), [103](#)
 - how users change their own [103](#)
 - if you forget [682](#)
 - peer-to-peer HA [653](#)
 - performance report [684](#)
 - PIM-SM
 - configuring [457](#)
 - overview [457](#)
 - server [173](#), [177](#)
 - ping [32](#), [161](#), [301](#), [658](#), [693](#)
 - default route [430](#)
 - ping test [302](#), [688](#)
 - policy
 - planning guidelines [59](#)
 - POP proxy [162](#)
 - port redirection [276](#), [286](#)
 - ports
 - address translation [166–167](#)
 - address translation and [275](#), [276](#), [282](#), [283](#), [286](#)
 - restricting source [171](#)
 - selecting in packet filter services [171](#)
 - selecting in proxy services [159](#)
 - valid values [340](#)
 - viewing configurations for services [343](#)
 - postmaster [500](#)
 - powering down [42](#)
 - preserve source port [167](#), [275](#), [282](#), [283](#), [286](#)
 - pre-shared password, defined [522](#)
 - primary name server [471](#)
 - primary-standby HA [653](#)
 - printer proxy [162](#)
 - process
 - CPU time [684](#)
 - displaying information [684](#)
 - domain [684](#)
 - domain access [18](#)
 - file access [18](#)
 - report [684](#)
 - status [684](#)
 - process command [684](#)
 - process use information [297](#)
 - protocol anomaly detection, see anomaly detection
 - see
 - proxies
 - additional agent properties [178](#)
 - agents and services [154–158](#)
 - connection type [159](#)
 - fast path sessions [155](#)
 - general properties [158](#)
 - list of agents [160–162](#)
 - non-transparent [155](#)
 - overview [21](#)
 - timeouts [159](#)
 - transparent [155](#)
 - proxy agents
 - expected connections [158](#)
 - global properties [149](#), [157](#)
 - multiple instantiation [158](#)
 - proxy connection information [305](#)
 - ps command [684](#)
- ## Q
- Quality of Service
 - overview [419](#)
 - profiles [421](#)
 - queues [422](#)
 - scenarios [426](#)
 - queue interval [513](#)
 - Quick Start Wizard
 - configurations set during [37](#)
 - Management Tools CD [29](#)
- ## R
- RADIUS authentication [97](#), [99](#)
 - read-only administrators [33](#), [44](#)
 - RealMedia [162](#)
 - RealTime Blackhole List [507](#)
 - rebooting [42](#)
 - reconfigure
 - DNS [490](#)
 - mail [498](#)
 - recovery options [706–708](#)
 - redirect delay [87](#)
 - redirection
 - about [264](#)
 - in packet filter rules [166](#)
 - in rules [276](#)
 - reference material
 - online help [13](#)
 - RFCs [14](#)
 - re-installing
 - from a USB drive [711–713](#)
 - from CD-ROM [710–711](#)
 - from the VCD [709–710](#)
 - options [708](#)
 - remote administration
 - via SSH [48](#)
 - via telnet [54](#)
 - remote certificates [638](#)
 - adding [639](#)
 - importing [640](#)
 - remote identities [636](#)
 - reports
 - mail queues [511](#)
 - network interfaces [693](#)
 - routing tables [693](#)
 - reputation classes, TrustedSource [144](#)

- reserved port range [169](#)
- response mappings [122–124](#)
- responses see IPS Attack Responses and System Event Responses
 - see [347](#)
- restart NIC [301](#)
- restarting services [341–345](#)
- restoring configuration files [581–582](#)
- restoring console access [682](#)
- reverse zones [481](#)
- RFCs [14](#)
- RIP
 - configuring [437–444](#)
 - overview [436](#)
 - processing [437–442](#)
 - server [173](#), [177](#)
 - trace and log information [467](#)
- RIP-unbound
 - server [173](#), [177](#)
- rlogin proxy [162](#)
- roles
 - about [44](#)
 - admin [19](#)
- rollaudit [335](#)
- rollaudit.conf file [335](#)
- rollback [598](#), [605](#)
- rolling log files [334](#)
- root [17](#), [19](#)
- rotating files [336](#)
- route commands [694](#)
- Routed mode [23](#)
- routes
 - default [428](#)
 - static [428](#), [435](#)
- routing tables report [693](#)
- RSH proxy [162](#)
- RTSP proxy [162](#)
- rule groups
 - about [60–61](#)
 - managing [269](#), [279](#)
- rules
 - about the Rules window [269–272](#)
 - default policy [268](#)
 - displaying [271](#)
 - example of [260](#)
 - failed connection request [679](#)
 - IPS and [120](#)
 - managing [273–278](#)
 - monitoring tool [681](#)
 - NAT [263](#)
 - order [265–267](#)
 - overview [257–262](#)
 - redirection [264](#)
 - troubleshooting [679](#)
 - types of [57](#)
 - viewing the active policy [272](#)

S

- safe mode [340](#)
- SafeWord authentication [100](#)
- scan buffer size, virus scanning [137](#)
- scanners, virus scanning [137](#)
- scanning encrypted files [138](#)
- SCC MIB files [385](#)
- sccMibSw [384](#)
- SCEP [624](#), [631](#), [639](#), [640](#)
- schedule, exporting log files [334](#)
- schedule, rolling log files [334](#)
- scripts
 - creating your own [611](#)
- secondary name server [471](#)
- secure shell (SSH) [48](#)
- SecureOS [17](#), [20](#)
- Security Parameters Index (SPI), using manual key exchange [543](#)
- sendmail
 - see also SMTP
 - advanced configuration [506–510](#)
 - allow/deny mail on a user basis [509](#)
 - configuration basics [502–505](#)
 - header stripping [507](#)
 - m4 macros [502](#)
 - masquerade_domain [510](#)
 - RealTime Blackhole list [507](#)
 - server [173](#), [177](#)
 - version [502](#)
- serial number (DNS) [484](#)
- server agents, global properties [149](#)
- server response configuration, FTP [187](#)
- server rules [58](#)
- server.conf file [467](#), [606](#)
- servers
 - about [172](#)
 - access control [22](#)
 - agent properties [172](#)
 - DNS [472](#)
 - list of [177](#)
 - sendmail [500](#)
 - telnet [54](#)
- service groups
 - about [59](#)
 - creating [153](#)
 - modifying [153](#)
- service status [341–345](#)
- services
 - see also proxies, filters, and servers
 - creating [152](#)
 - modifying [152](#)
 - overview [149](#)
 - packet filter service properties [171](#)
 - proxy service properties [158](#)
 - selecting in rules [274](#), [280](#)
- session maximums [169](#)

- sfredirect server [177](#)
 - sftp [48](#)
 - showaudit command [314](#), [703](#)
 - shutdown [42](#), [43](#)
 - sighup command [340](#)
 - signature categories [118](#), [125](#), [131](#)
 - signature groups [125–127](#)
 - signature vulnerabilities [128](#), [131](#)
 - signatures
 - disable [129](#), [131](#)
 - enable [129](#), [131](#)
 - signing export files [334](#)
 - SIP
 - Application Defense [244](#)
 - proxy [162](#)
 - size filter [214](#)
 - slave zone [481](#)
 - SMTP
 - Application Defense [220](#)
 - mail queues [511–513](#)
 - secure split servers [497](#)
 - transparent mail [497](#)
 - SMTP proxy
 - agent [162](#)
 - strip source routing [188](#)
 - SNMP
 - agent [173](#), [177](#), [381](#)
 - allowed get communities [387](#)
 - Application Defense [242](#)
 - authentication header [383](#)
 - community names [383](#)
 - configuration agents [386–388](#)
 - configuring rules [389](#)
 - High Availability and [385](#)
 - management information base (MIB) [383–385](#)
 - passing to and through the firewall [390](#)
 - proxy [162](#), [381](#)
 - response trap [383](#)
 - trap destinations [388](#)
 - traps [352](#), [357](#), [360](#)
 - snmptrap [382–383](#)
 - SOA record [483](#)
 - SOCKS Application Defense [240](#)
 - SOCKS proxy [162](#)
 - SoftRemote [643](#)
 - software management
 - rollback [706](#)
 - uninstall [706](#)
 - source burb [274](#), [282](#)
 - source routing, SMTP [188](#)
 - spam, TrustedSource [144](#)
 - SPI (Security Parameters Index), using manual key exchange [543](#)
 - split DNS [471](#), [472](#)
 - SSH
 - Application Defense [247](#)
 - client [51](#)
 - configuring [48](#)
 - keys [645](#)
 - known host keys [182–186](#)
 - server [173](#), [177](#)
 - SSH proxy
 - about [182](#)
 - agent [162](#)
 - configuring [185–186](#)
 - known host keys [182–186](#)
 - SSL decryption [198](#), [623](#)
 - ssod server [173](#), [178](#)
 - stacks [20](#)
 - stateful inspection [21](#)
 - stateful packet inspection [163–166](#), [171](#)
 - stateful session failover [171](#)
 - static route [428](#), [435](#)
 - statistics, Audit Management [332](#)
 - status process [684](#)
 - status reports, routing tables [693](#)
 - streamworks proxy [162](#)
 - Strikeback [347](#)
 - Strong Cryptography [201](#)
 - sub-domain (DNS) [486](#)
 - subnet object [72](#)
 - sunrcp proxy [162](#)
 - super-user [17](#), [19](#)
 - suspicious, TrustedSource [144](#)
 - sysbase proxy [162](#)
 - syslog
 - audit messages [337](#)
 - syslog proxy [162](#)
 - syslogd [336](#)
 - system calls [18](#)
 - System Event Responses
 - about [347](#)
 - managing [354–357](#)
 - response settings [358](#)
 - system resources [295](#), [298](#), [302](#)
 - CPU use [299](#)
- ## T
- T.120
 - Application Defense [235](#)
 - proxy [162](#), [189](#)
 - TCP
 - checksum offload [397](#)
 - Network Defense [367](#)
 - state information [306](#)
 - TCP dump [689](#)
 - tcpdump [32](#), [689](#), [701](#), [703](#)
 - TE see Type Enforcement
 - see
 - Telnet
 - defined [54](#)
 - proxy [162](#)

- server [54](#), [172](#), [178](#)
 - temporarily disabling services [341–345](#)
 - time (setting) [568](#)
 - time periods
 - creating [76](#)
 - in rules [274](#)
 - selecting in rules [281](#)
 - timeouts
 - packet filter [171](#)
 - proxy [159](#)
 - toolbar, Logfile Options [332](#)
 - Tools menu
 - ARP Table [686](#)
 - DNS lookup [688](#)
 - Get route [687](#)
 - Ping test [688](#)
 - TCP dump [689](#)
 - Traceroute [692](#)
 - top command [684](#)
 - traceroute [32](#), [161](#), [692](#), [693](#)
 - traffic basics [57](#)
 - transparent
 - connection type [159](#)
 - DNS [470](#), [474](#)
 - mail (SMTP) [497](#)
 - proxies [155](#)
 - rules for transparent mail [499](#)
 - Transparent (bridged) mode
 - about [25](#)
 - scenarios [25–28](#)
 - Transparent firewall [25](#)
 - transparent interface properties [409](#)
 - transport mode [521](#), [538](#)
 - trap destinations [388](#)
 - traps [352](#), [360](#), [382–383](#), [386](#)
 - troubleshooting
 - no admin access [682](#)
 - NTP [700](#)
 - rules [679](#)
 - TrustedSource [138](#), [258](#), [276](#)
 - configuring [142](#)
 - reputation classes [144](#)
 - subscription [144](#)
 - Web site [145](#)
 - TTL value (DNS) [484](#)
 - tunnel mode [521](#), [538](#)
 - TXT record [485](#), [488](#), [489](#)
 - Type Enforcement
 - about [17](#)
 - defined [18](#)
 - effects [19](#)
 - file types [611](#)
 - how it works [17](#)
 - restore [18](#)
 - sendmail [501](#)
 - types of traffic [57](#)
- ## U
- UDP
 - checksum [171](#)
 - Network Defense [371](#)
 - unbound DNS server [472](#)
 - uni-directional packet filter rules [171](#)
 - Uninterruptible Power Supply [616](#)
 - UNIX
 - editing files [606](#)
 - security [17](#)
 - text editors [606](#)
 - unverified, TrustedSource [144](#)
 - updates
 - IPS signature files [121](#), [132](#)
 - virus scanning signature files [135–136](#)
 - UPS
 - configuring [616](#)
 - uptime command [684](#)
 - uptime statistics [295](#)
 - USB drive [707](#)
 - user groups
 - displaying [107](#)
 - selecting in rules [289](#)
 - users
 - changing password [47](#)
 - displaying [107](#)
 - users and user groups [107](#), [115](#)
- ## V
- var/log/audit.raw file [495](#)
 - var/log/daemon.log file [495](#)
 - var/spool/mqueue.X [501](#), [511](#)
 - viewing audit [316](#)
 - viewing items [33](#), [44](#)
 - virtual burb [525](#)
 - virtual CD [708](#), [709–710](#)
 - virus scanning
 - about [134](#), [502](#)
 - advanced features [137–138](#)
 - engine updates [135–136](#)
 - file size range [137](#)
 - signature file updates [135–136](#)
 - vmstat command [684](#)
 - VPN
 - AH keys [543](#)
 - algorithms [544](#)
 - authentication [523](#)
 - basic steps [533](#)
 - certificate server [634](#)
 - certificates [633](#)
 - client address pools [521](#)
 - commands [703](#)
 - definitions [535–545](#)
 - Extended Authentication [524](#)

- firewall certificate [623](#)
- fixed IP [551](#)
- IPsec [518](#)
- ISAKMP server [533](#)
- key types [518](#)
- LDAP [641](#)
- local authentication [542](#)
- managing client address pools [546](#)
- managing the ISAKMP server [552](#)
- managing VPN definitions [535](#)
- NAT Traversal (NAT-T) [544](#)
- ordering VPN definitions [525](#)
- planning [519–532](#)
- remote authentication [539](#)
- remote certificates [638](#)
- remote identities [636](#)
- scenarios [555–563](#)
- SPI [543](#)
- status of active [304](#)
- transport mode [521](#)
- troubleshooting [703](#)
- tunnel mode [521](#)
- understanding [517](#)
- user interface reference [535–554](#)
- virtual burb [525](#)
- XAUTH [539](#)
- vulnerabilities [128](#), [131](#)

W

- wais proxy [162](#)
- Web sites
 - activation [588](#)
 - TrustedSource [145](#)
- whois command [162](#)
- whois proxy [162](#)
- Windows Domain authentication [96](#), [97](#)
- wins proxy [162](#)
- WINS server [549](#)

X

- X500 proxy [162](#)
- XAUTH [539](#)
- Xwindows proxy [162](#)

Z

- zones [480](#)