

WebShield® appliance

version 3.0



COPYRIGHT

Copyright © 2004 Networks Associates Technology, Inc. All Rights Reserved.

No part of this publication may be reproduced, transmitted, transcribed, stored in a retrieval system, or translated into any language in any form or by any means without the written permission of Networks Associates Technology, Inc., or its suppliers or affiliate companies. To obtain this permission, write to the attention of the McAfee legal department at: 5000 Headquarters Drive, Plano, Texas 75024, or call +1-972-963-8000.

TRADEMARK ATTRIBUTIONS

Active Firewall, Active Security, ActiveSecurity (and in Katakana), ActiveShield, AntiVirus Anyware and design, Clean-Up, Design (Stylized E), Design (Stylized N), Enterecept, Enterprise SecureCast, Enterprise SecureCast (and in Katakana), ePolicy Orchestrator, First Aid, ForceField, GMT, GroupShield, GroupShield (and in Katakana), Guard Dog, HomeGuard, Hunter, IntruShield, Intrusion Prevention Through Innovation, M and Design, McAfee, McAfee (and in Katakana), McAfee and Design, McAfee.com, McAfee VirusScan, NA Network Associates, Net Tools, Net Tools (and in Katakana), NetCrypto, NetOctopus, NetScan, NetShield, Network Associates, Network Associates Coliseum, NetXray, NotesGuard, Nuts & Bolts, Oil Change, PC Medic, PCNotary, PrimeSupport, RingFence, Router PM, SecureCast, SecureSelect, SpamKiller, Stalker, ThreatScan, TIS, TMEG, Total Virus Defense, Trusted Mail, Uninstaller, Virex, Virus Forum, Virusscan, Virusscan, Virusscan (And In Katakana), Webscan, Webshield, Webshield (And In Katakana), Webstalker, WebWall, What's The State Of Your IDS?, Who's Watching Your Network, Your E-Business Defender, Your Network. Our Business. are registered trademarks or trademarks of McAfee, Inc. and/or its affiliates in the US and/or other countries. Red in connection with security is distinctive of McAfee® brand products. All other registered and unregistered trademarks herein are the sole property of their respective owners.

LICENSE INFORMATION

License Agreement

NOTICE TO ALL USERS: CAREFULLY READ THE APPROPRIATE LEGAL AGREEMENT CORRESPONDING TO THE LICENSE YOU PURCHASED, WHICH SETS FORTH THE GENERAL TERMS AND CONDITIONS FOR THE USE OF THE LICENSED SOFTWARE. IF YOU DO NOT KNOW WHICH TYPE OF LICENSE YOU HAVE ACQUIRED, PLEASE CONSULT THE SALES AND OTHER RELATED LICENSE GRANT OR PURCHASE ORDER DOCUMENTS THAT ACCOMPANIES YOUR SOFTWARE PACKAGING OR THAT YOU HAVE RECEIVED SEPARATELY AS PART OF THE PURCHASE (AS A BOOKLET, A FILE ON THE PRODUCT CD, OR A FILE AVAILABLE ON THE WEB SITE FROM WHICH YOU DOWNLOADED THE SOFTWARE PACKAGE). IF YOU DO NOT AGREE TO ALL OF THE TERMS SET FORTH IN THE AGREEMENT, DO NOT INSTALL THE SOFTWARE. IF APPLICABLE, YOU MAY RETURN THE PRODUCT TO MCAFEE OR THE PLACE OF PURCHASE FOR A FULL REFUND.

Attributions

This product includes or may include:

- Software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>). • Cryptographic software written by Eric A. Young and software written by Tim J. Hudson. • Some software programs that are licensed (or sublicensed) to the user under the GNU General Public License (GPL) or other similar Free Software licenses which, among other rights, permit the user to copy, modify and redistribute certain programs, or portions thereof, and have access to the source code. The GPL requires that for any software covered under the GPL which is distributed to someone in an executable binary format, that the source code also be made available to those users. For any such software covered under the GPL, the source code is made available on this CD. If any Free Software licenses require that McAfee provide rights to use, copy or modify a software program that are broader than the rights granted in this agreement, then such rights shall take precedence over the rights and restrictions herein. • Software originally written by Henry Spencer, Copyright 1992, 1993, 1994, 1997 Henry Spencer.
- Software originally written by Robert Nordier, Copyright © 1996-7 Robert Nordier. • Software written by Douglas W. Sauder. • Software developed by the Apache Software Foundation (<http://www.apache.org/>). A copy of the license agreement for this software can be found at www.apache.org/licenses/LICENSE-2.0.txt.
- International Components for Unicode ("ICU") Copyright © 1995-2002 International Business Machines Corporation and others. • Software developed by CrystalClear Software, Inc., Copyright © 2000 CrystalClear Software, Inc. • FEAD[®] Optimizer[®] technology, Copyright Netopsystems AG, Berlin, Germany.
- Outside In[®] Viewer Technology © 1992-2001 Stellent Chicago, Inc. and/or Outside In[®] HTML Export, © 2001 Stellent Chicago, Inc. • Software copyrighted by Thai Open Source Software Center Ltd. and Clark Cooper, © 1998, 1999, 2000. • Software copyrighted by Expat maintainers. • Software copyrighted by The Regents of the University of California, © 1989. • Software copyrighted by Gunnar Ritter. • Software copyrighted by Sun Microsystems[®], Inc. © 2003.
- Software copyrighted by Gisle Aas. © 1995-2003. • Software copyrighted by Michael A. Chase, © 1999-2000. • Software copyrighted by Neil Winton, © 1995-1996. • Software copyrighted by RSA Data Security, Inc., © 1990-1992. • Software copyrighted by Sean M. Burke, © 1999, 2000. • Software copyrighted by Martijn Koster, © 1995. • Software copyrighted by Brad Appleton, © 1996-1999. • Software copyrighted by Michael G. Schwern, © 2001.
- Software copyrighted by Graham Barr, © 1998. • Software copyrighted by Larry Wall and Clark Cooper, © 1998-2000. • Software copyrighted by Frodo Looijard, © 1997. • Software copyrighted by the Python Software Foundation, Copyright © 2001, 2002, 2003. A copy of the license agreement for this software can be found at www.python.org. • Software copyrighted by Beman Dawes, © 1994-1999, 2002. • Software written by Andrew Lumsdaine, Lie-Quan Lee, Jeremy G. Siek © 1997-2000 University of Notre Dame. • Software copyrighted by Simone Bordet & Marco Cravero, © 2002. • Software copyrighted by Stephen Purcell, © 2001. • Software developed by the Indiana University Extreme! Lab (<http://www.extreme.indiana.edu/>). • Software copyrighted by International Business Machines Corporation and others, © 1995-2003. • Software developed by the University of California, Berkeley and its contributors.
- Software developed by Ralf S. Engelschall <rse@engelschall.com> for use in the mod_ssl project (<http://www.modssl.org/>). • Software copyrighted by Kevin Henney, © 2000-2002. • Software copyrighted by Peter Dimov and Multi Media Ltd. © 2001, 2002. • Software copyrighted by David Abrahams, © 2001, 2002. See <http://www.boost.org/libs/bind/bind.html> for documentation. • Software copyrighted by Steve Cleary, Beman Dawes, Howard Hinnant & John Maddock, © 2000. • Software copyrighted by Boost.org, © 1999-2002. • Software copyrighted by Nicolai M. Josuttis, © 1999. • Software copyrighted by Jeremy Siek, © 1999-2001. • Software copyrighted by Daryle Walker, © 2001. • Software copyrighted by Chuck Allison and Jeremy Siek, © 2001, 2002. • Software copyrighted by Samuel Kremp, © 2001. See <http://www.boost.org> for updates, documentation, and revision history. • Software copyrighted by Doug Gregor (gregod@cs.rpi.edu), © 2001, 2002. • Software copyrighted by Cadenza New Zealand Ltd., © 2000. • Software copyrighted by Jens Maurer, © 2000, 2001.
- Software copyrighted by Jaakko Järvi (jaakko.jarvi@cs.utu.fi), © 1999, 2000. • Software copyrighted by Ronald Garcia, © 2002. • Software copyrighted by David Abrahams, Jeremy Siek, and Daryle Walker, © 1999-2001. • Software copyrighted by Stephen Cleary (shammah@voyager.net), © 2000. • Software copyrighted by Housemarque Oy <<http://www.housemarque.com>>, © 2001. • Software copyrighted by Paul Moore, © 1999. • Software copyrighted by Dr. John Maddock, © 1998-2002. • Software copyrighted by Greg Colvin and Beman Dawes, © 1998, 1999. • Software copyrighted by Peter Dimov, © 2001, 2002.
- Software copyrighted by Jeremy Siek and John R. Bandela, © 2001. • Software copyrighted by Joerg Walter and Mathias Koch, © 2000-2002.

PATENT INFORMATION

Protected by US Patents 6,496,875; 6,499,109; 6,513,122; 6,668,289.

Contents

Preface	15
Audience	15
Conventions	16
Getting information	17
Contacting McAfee Security & Network Associates	18
 1 Basic Concepts	 19
What is the WebShield appliance?	19
WebShield features	19
What are its default settings?	20
How secure is the appliance?	20
Important considerations	21
Choosing the right operational mode	22
Protocol support	22
Handling protocol traffic	23
Recommended network topologies	23
Restrictions	24
Accessing the appliance	24
Using a web browser	24
Using the WebShield client application	24
Making the appliance more secure	25
Specifying your inside and outside networks	25
Using anti-virus scanning	26
Controlling SMTP e-mail access	28
Handling unwanted content in SMTP e-mail messages	28
Scanning SMTP e-mail messages for spam	28
Using policies	29
Sharing scanning resources	29
Monitoring the appliance	30
Maintaining the appliance	30
Troubleshooting the appliance	30

2 Which Operational Mode?	31
Explicit or transparent?	32
Explicit Proxy mode	34
Positioning the appliance	34
Configuration example	36
Transparent Router mode	38
Positioning the appliance	39
Configuration example	40
Transparent Bridge mode	41
Positioning the appliance	41
Changing operational modes	43
 3 Explicit Proxy Mode	 45
SMTP Scenarios	46
One site	46
Directional scanning	47
Multiple sites	49
Demilitarized zone (DMZ)	50
International organization	52
Handling network failure	53
Fail-over	53
Fail-closed	54
Fail-open	55
FTP scenarios	56
Outbound FTP communication	56
Inbound FTP with one server	57
Inbound FTP with multiple servers	58
HTTP scenarios	59
Outbound HTTP with internal web cache	59
Outbound HTTP with external web cache	60
Outbound HTTP without web cache	62
Inbound HTTP	62
POP3 scenarios	63
Using a generic connection	63
Using a dedicated connection	64
POP3 with multiple servers	65
Load sharing	66

4	Transparent Router Mode	67
SMTP scenarios		67
One site		68
Multiple sites		69
Demilitarized zone (DMZ)		70
International organization		71
FTP scenarios		72
Outbound FTP		72
Inbound FTP		73
HTTP scenarios		73
Outbound HTTP with internal web cache		74
Outbound HTTP with external web cache		75
Outbound HTTP without web cache		76
Inbound HTTP		77
POP3 scenarios		78
Load sharing		79
5	Transparent Bridge Mode	81
SMTP scenarios		81
One site		82
Dedicated appliances		83
Multiple sites		84
Demilitarized zone (DMZ)		85
International organization		86
FTP scenarios		87
Outbound FTP		87
Inbound FTP		88
HTTP scenarios		88
Outbound HTTP with internal web cache		89
Outbound HTTP with external web cache		90
Outbound HTTP without web cache		91
Inbound HTTP		92
POP3 scenarios		93
Load sharing		94

6 Initial Configuration	95
Accessing the appliance	95
User name and default password	95
Initial configuration	95
Appliance name	96
Domain name	96
Default gateway	96
Operational mode	97
Bridge Priority	97
Protocols	97
Interface addresses	98
Inside and Outside Networks	98
What should be in the Inside Networks list?	98
What should be in the Outside Networks list?	99
Adding domains and networks	99
Resolving conflicts	100
DNS servers	101
Static routes	101
Dynamic routes	101
Load sharing	102
Password settings	102
Date and time settings	102
Operational language	102
7 Policies	103
What is a policy?	103
Policy actions	103
Issuing alerts and notifications	104
Global policies	104
Non-global policies	104
Inheriting global settings	105
Ordering non-global policies	105
Adding time-specific settings to non-global policies	106
Multiple instances of rules and settings	106
Handling overlapping time restrictions	107
Specifying time restrictions	107
Policy Groups	108
Content rules	108

Before you begin	109
Spend some time planning	109
General guidelines	110
Considering legal implications	110
8 Managing SMTP e-mail	111
General e-mail (SMTP) configuration	111
E-mail (SMTP) delivery	111
Domain relays	112
DNS servers	112
Fallback relays	112
Creating a postmaster	113
Anti-relay	113
Local domains	114
Deny domains	114
Permit domains	114
Anti-relay response	115
How the appliance uses the anti-relay lists	115
Permit and deny settings	116
Permit Sender	117
Deny Sender	118
RBL servers	118
Responding to unwanted e-mail messages	119
Connection settings (Advanced)	120
Intercept ports	120
Listen ports	120
Listeners, connections and memory	120
Retryer settings	121
Policy-based e-mail configuration	121
Before you begin	121
Content Policies	122
Alert settings	123
Anti-spam settings	124
Anti-virus settings	124
Content scanner	124
Corrupt content	124
Disclaimer Text	124
Encrypted content	125

File filtering	125
HTML settings	126
Mail settings	127
Mail size filtering	127
Protected content	128
Scanner control (denial-of-service attacks)	128
Signed content (digital signatures)	129
Connection policies	131
Anti-Relay (routing characters)	131
Time-outs	132
Transport logging	132
Protocol policies	133
Data command options	133
Denial-of-service prevention	134
E-mail address configuration	134
Message processing	135
Transparency options	136
How e-mail messages are processed	137
Multi-policies for e-mail messages	138
Protocol policies	138
Connection policies	138
Content Policies	139
When a scanner triggers	141
Primary actions	142
Secondary actions	143
Setting up secondary actions in transparent modes	145

9 Scanning for Spam 147

What is SPAM?	147
What is SpamKiller for WebShield appliances?	147
Understanding spam scores	150
Disabling rules	151
Tips for avoiding spam	151
Updating your anti-spam software	152

10 E-mail (SMTP) content scanning 153

Content Rules and Rule Groups	153
Importing and exporting content rules	153
Scanning for content	154
Creating content rules	154
Giving a name and description to the rule	155
Specifying where the rule applies	155
Specifying the action to take when the rule is triggered	155
Specifying the word or phrase you want to detect	156
Adding optional advanced features	157
Understanding complex content rules for e-mail messages	158
Understanding limitations in content scanning	159
Examples of content rules	160
Keeping information confidential	160
Reducing network load	161
Blocking offensive words	161
Stopping nuisance e-mail messages	161
Reducing distractions	162

11 Virus-scanning 163

What is heuristic analysis?	163
Anti-virus software	164
Why update?	164
When is the best time to update?	164
Scheduling the updates	165
Local Updates and EXTRA DAT files	165
Scanning for viruses	165
Setting the action against viruses	165
Setting the level of scanning and type of protection	166
Customizing anti-virus settings	166
Blocking specific threats	168

12 Managing HTTP 169

General HTTP configuration	169
Connection settings (Advanced)	169
Intercept ports	169
Listen ports	169
Listeners, connections and memory	170
Policy-based HTTP configuration	170
Content policies	171
Alert settings	171
Anti-virus settings	171
HTML settings	172
Scanner control (denial-of-service attacks)	172
HTTP actions	172
Connection policies	173
Client (check for client)	173
Time-outs	173
Protocol policies	174
Client alert messages	174
Client download status messages	174
Denial-of-service prevention	174
Download status and data trickling	175
FTP over HTTP	176
Handoff host	176
Header blocking and modifications	177
Protocol details	177
Scanning	179
Streaming media	179
URL blocking and request permissions	180

13 Managing FTP 183

General FTP configuration	183
Connection settings (Advanced)	183
Intercept ports	183
Listen ports	183
Listeners, connections and memory	184
Policy-based FTP configuration	184
Content policies	184
Anti-virus settings	185

Scanner control (denial-of-service attacks)	185
FTP actions	185
Connection policies	185
Time-outs	186
Protocol policies	186
Data processing	186
Download status and data trickling	187
FTP handoff host	188
Upload status and data trickling	189
14 Managing POP3	191
General POP3 configuration	191
Connection settings (Advanced)	191
Intercept ports	191
Listen ports	191
Dedicated Ports	192
Listeners, connections and memory	192
Policy-based POP3 configuration	193
Before you begin	193
Content Policies	194
Alert settings	194
Anti-virus settings	195
Corrupt content	195
Encrypted content	195
Mail settings	195
Mail size filtering	196
Protected content	196
Scanner control (denial-of-service attacks)	197
Signed content (digital signatures)	198
POP3 Actions	199
Connection policies	200
Time-outs	201
Protocol policies	201
Protocol settings	201

15 Monitoring the appliance 203

Monitoring the appliance	204
System Status	204
Performance	206
Logging and Alerting Reports	208
Viewing the log	208
Displaying data graphically	209
Report using .TSV format	209
Logging charts	210
System Updates	211
System Resources	211
Configuring Logging and Alerting	211

16 Load Sharing 215

Load sharing basic concepts	215
Configuration scenarios	216
Scenario 1 — single appliance in non-sharing mode	216
Scenario 2 — single appliance in sharing mode	216
Scenario 3 — controlling appliance off-loads some workload	216
Scenario 4 — controlling appliance off-loads all workload	217
Configuring load sharing	217
Configuring the controlling appliance	217
Configuring the load sharing appliance	218
Load sharing examples	219
Example 1 — single appliance in non-sharing mode	219
Example 2 — single appliance in sharing mode	219
Example 3 — controlling appliance off-loads some workload	220
Example 4 — controlling appliance off-loads all workload	220
Viewing the load sharing status	221

17 Maintaining the Appliance 223

Changing the password	223
Turning off or rebooting the appliance	223
Setting the system date and time	223
Setting the operational language	224
Installing service packs and HotFixes	224
Evaluating SpamKiller	224

Activating SpamKiller	224
Enabling the ePolicy Orchestrator Agent	225
Uninstalling the ePolicy Orchestrator Agent	225
Saving the logs	225
Off-box logging with Syslog	226
Saving system configuration	226
Restoring system settings	226
Restoring default settings	227
Accessing the MIB definition file	227
Copying configuration	227
Removing old files	228
Quarantine maintenance	228
Deferred e-mail maintenance	228
Restricting the number of log files	228
A Word Separators	229
Type key	230
Character list	231
Index	243

Preface

This guide introduces McAfee® WebShield® appliance software version 3.0, and provides the following information:

- Overview of the product.
- Basic concepts about how the appliance works, and how you can integrate it into your existing network.

Audience

This information is designed for system and network administrators who are responsible for their company's anti-virus and security program.

Conventions

This guide uses the following conventions:

Bold All words from the user interface, including options, menus, buttons, and dialog box names.

Example

Type the **User name** and **Password** of the desired account.

Courier The path of a folder or program; a web address (URL); text that represents something the user types exactly (for example, a command at the system prompt).

Examples

The default location for the program is:

`C:\Program Files\Network Associates\VirusScan`

Visit the Network Associates web site at:

`http://www.networkassociates.com`

Run this command on the client computer:

`C:\SETUP.EXE`

Italic For emphasis or when introducing a new term; for names of product manuals and topics (headings) within the manuals.

Example

Refer to the *VirusScan Enterprise Product Guide* for more information.

<TERM> Angle brackets enclose a generic term.

Example

In the console tree under **ePolicy Orchestrator**, right-click **<SERVER>**.

NOTE Supplemental information; for example, an alternate method of executing the same command.

WARNING Important advice to protect a user, computer system, enterprise, software installation, or data.

Getting information

Package Guide *^	Lists the components included in the product box. <i>What is in the box?</i>
Read this first Guide **^	Important information about the product, license agreements, HotFixes and Service Packs, that should be read before reading other guides. <i>WebShield Appliance 3.0 Read this first</i>
Installation Guide *^	Instructions for installing the appliance. <i>WebShield 3100 version 3.0 Installation Guide</i> <i>WebShield 3200 version 3.0 Installation Guide</i> <i>WebShield 3300 version 3.0 Installation Guide</i>
Upgrade Guide *	Guide describing how to upgrade from WebShield appliance software version 2.7 to software version 3.0. <i>Upgrading WebShield appliance from version 2.7 to 3.0</i>
Procedures Guide *	Product introduction and features, detailed instructions for configuring the software, information on deployment, recurring tasks, and operating procedures. <i>WebShield Appliance 3.0 Product Guide</i>
Concepts Guide *	High level conceptual information about the appliance and how to integrate it into the network. <i>WebShield Appliance 3.0 Concepts Guide</i>
Help §	High-level information about each page in the user interface. <i>Quick Help panels</i>
Configuration Guide *	<i>For use with ePolicy Orchestrator™</i> . Procedures for configuring, deploying, and managing your McAfee Security product through ePolicy Orchestrator management software. <i>WebShield appliance version 3.0 for use with ePolicy Orchestrator 3.5</i>
Release Notes **^	Product information, resolved issues, any known issues, and last-minute additions or changes to the product or its documentation. <i>WebShield appliance version 3.0 Release Notes</i>

* An Adobe Acrobat .PDF file on the product CD or the McAfee Security download site.

^ A printed manual that accompanies the product CD. Note: Some language manuals may be available only as a .PDF file.

§ Help accessed from the software application, that provides high-level page-based help.

Contacting McAfee Security & Network Associates

Technical Support

Home Page	http://www.networkassociates.com/us/support/
KnowledgeBase Search	https://knowledgemap.nai.com/phpclient/homepage.aspx
PrimeSupport Service Portal *	https://mysupport.nai.com

McAfee Security Beta Program <http://www.networkassociates.com/us/downloads/beta/>

Security Headquarters — AVERT (Anti-Virus Emergency Response Team)

Home Page	http://www.networkassociates.com/us/security/home.asp
Virus Information Library	http://vil.nai.com
Submit a Sample — AVERT WebImmune	https://www.webimmune.net/default.asp
AVERT DAT Notification Service	http://vil.nai.com/vil/join-DAT-list.asp

Download Site

Home Page	http://www.networkassociates.com/us/downloads/
DAT File and Engine Updates	ftp://ftp.nai.com/virusdefs/4.x/
Anti-spam rules and Engine Updates	ftp://ftp.nai.com/spamdefs/1.x/
Product Upgrades *	https://secure.nai.com/us/forms/downloads/upgrades/login.asp

Training

McAfee Security University	http://www.networkassociates.com/us/services/education/mcafee/university.htm
----------------------------	---

Network Associates Customer Service

E-mail	services_corporate_division@nai.com
Web	http://www.networkassociates.com/us/index.asp
US, Canada, and Latin America toll-free:	
Phone	+1-888-VIRUS NO or +1-888-847-8766 Monday – Friday, 8 a.m. – 8 p.m., Central Time

For additional information on contacting Network Associates and McAfee Security— including toll-free numbers for other geographic areas — see the Contact file that accompanies this product release.

* Logon credentials required.

This chapter describes some basic concepts that will help you integrate your appliance into your existing network. It provides an overview of each concept and refers you to other parts of the guide where more detailed guidelines are provided.

What is the WebShield appliance?

The WebShield appliance is a purpose-built system that may be installed at key points on your network. It is typically installed at the Internet gateway, where it protects the points of entry to the network and minimizes the risk to your business-critical systems.

The WebShield appliance is available in three hardware platforms:

- WebShield 3100
- WebShield 3200
- WebShield 3300

NOTE

This guide provides information about all of these platforms, and clearly labels any platform specific information where necessary.

WebShield features

The appliance supports a wide range of features that help you combat threats to your organization. See *Introducing the WebShield appliance* in the *Product Guide* for a comprehensive list of supported features, including a list of features that are new for this release.

What are its default settings?

The user name is **webshield**, and it cannot be changed.

The default password is **webshieldchangeme**.

The default IP address for LAN1 port is **10.1.1.108**, and for LAN2 it is **10.1.2.108**.

The default system name depends on your appliance and is either **webshield3100**, **webshield3200**, or **webshield3300**.

NOTE

To improve security and deter hackers, you should always change the password, default system name, and IP addresses of the appliance.

All other defaults can be found by clicking on the relevant link in the **Resource Information** page. The **Resource Information** page can be accessed by clicking on the **Resources** link in the links bar at the top of the appliance application or web page.

How secure is the appliance?

The appliance is a secure device that can only be accessed through a secure HTTPs link.

If you are using a web browser, when entering the URL for the appliance, you must use HTTPS rather than HTTP.

The appliance's operating system prevents unauthorized access to its internal filing system.

The appliance is password protected.

NOTE

To make sure that only you can configure the appliance, you should change the default password for the appliance.

Important considerations

When setting up the appliance for the first time, you should consider:

- Which operational mode you require. See [Choosing the right operational mode on page 22](#).
- Which protocols you want the appliance to handle. See [Protocol support on page 22](#).
- How to integrate the appliance within your existing network. See [Recommended network topologies on page 23](#).
- Which access method you want to use. See [Accessing the appliance on page 24](#).
- How you can make the appliance more secure. See [Making the appliance more secure on page 25](#).
- Which networks and domains you want to include in the appliance's Inside Networks and Outside Networks list. See [Specifying your inside and outside networks on page 25](#).
- How you can make your network more secure. See [Using anti-virus scanning on page 26](#).
- How you can control SMTP e-mail access. See [Controlling SMTP e-mail access on page 28](#).
- How you can control the content of SMTP e-mail messages entering and leaving your network. [Handling unwanted content in SMTP e-mail messages on page 28](#).
- How you can protect your SMTP e-mail users from unwanted e-mail messages known as *spam*. See [Scanning SMTP e-mail messages for spam on page 28](#).
- Which policies you want to set up. See [Using policies on page 29](#).
- If you want to use load sharing to share the scanning workload. See [Sharing scanning resources on page 29](#).
- How you want to monitor the appliance. See [Monitoring the appliance on page 30](#).
- Which maintenance procedures you need to follow. For example, you might want to back up the system configuration so that it can be easily restored in the event of a problem with the appliance. See [Maintaining the appliance on page 30](#).
- What to do if you have a problem with the appliance, or want to submit a sample for analysis. See [Troubleshooting the appliance on page 30](#).

Choosing the right operational mode

The appliance operates in one of three modes:

- **Explicit Proxy** mode.
- **Transparent Router** mode.
- **Transparent Bridge** mode.

Selecting the right operational mode for the appliance is an important choice as it impacts how you integrate your appliance into your existing network and how the appliance handles traffic. After you select the right mode for the appliance, you should not need to change its mode until you next restructure your network. See [Which Operational Mode? on page 31](#) for details on operational modes.

Protocol support

The appliance can handle and scan traffic for the following protocols:

- Simple Mail Transfer Protocol (SMTP) e-mail messages.
- File Transfer Protocol (FTP) exchanges.
- Hypertext Transfer Protocol (HTTP) web browsing.
- Post Office Protocol version 3 (POP3) Internet e-mail messages.

NOTE

HTTPS, and FTP over HTTP (download only) are also supported as part of HTTP support. All other protocols will either be refused or not scanned, depending on the Operational Mode of the appliance. For example, the appliance does not scan Real Player traffic. For information on the handling of streaming media, see [Streaming media on page 179](#).

The appliance does not use the *Content Vectoring Protocol* (CVP).

There are appliance settings that are protocol-specific and settings that are similar for all protocols. Managing the protocol settings is described in more detail in:

- [Managing SMTP e-mail on page 111](#).
- [Managing HTTP on page 169](#).
- [Managing FTP on page 183](#).
- [Managing POP3 on page 191](#).

Handling protocol traffic

The appliance allows you to enable or disable each protocol (SMTP, FTP, HTTP and POP3).

If the appliance is in **Explicit Proxy** mode and a protocol is disabled, traffic that is directed to the appliance for that protocol will be refused, effectively blocking that protocol.

If the appliance is in either **Transparent Router** or **Transparent Bridge** mode, and the protocol is disabled, traffic for that protocol will pass through the appliance, but will not be scanned.

When operating in **Explicit Proxy** mode, only SMTP, FTP, HTTP and POP3 traffic should be sent to the appliance. All other traffic will be refused.

Recommended network topologies

The appliance can be used in almost any network topology, although there are some restrictions on how it can be used (See [Restrictions on page 24](#), for more information).

There are too many possible topologies to describe them all, but some typical topologies for each operational mode are provided in:

- [Explicit Proxy Mode on page 45](#)
- [Transparent Router Mode on page 67](#)
- [Transparent Bridge Mode on page 81](#)

NOTE

To scan a supported protocol, you must make sure that traffic for that protocol passes through the appliance. Any traffic that can bypass the appliance will not be scanned, leaving your network vulnerable to virus attacks.

For security reasons, you must use the appliance inside your organization and behind a correctly configured firewall.

If you are in any doubt about your network's topology and how you should integrate the appliance, consult your network expert.

Restrictions

The appliance cannot be used as:

- A firewall. You must use the appliance within your organization, and behind your existing firewall.
- A mail server. You might need to configure your firewall, mail server, web cache and other equipment to route protocol traffic through the appliance.
- A *Mail Transfer Agent (MTA)*. The appliance does not keep copies of the e-mail messages that pass through it, unless they are quarantined or deferred.
- A general purpose web server for storing web pages.
- A general purpose server for storing extra software and files. Do not install any software on the appliance or add any extra files to it unless specifically instructed by the appliance's documentation or a McAfee support representative.

The appliance is *not* able to handle *all* types of traffic. If you have selected Explicit Proxy mode, the appliance only handles SMTP, FTP, HTTP, and POP3 traffic. Do *not* attempt to route other traffic through the appliance.

Accessing the appliance

When you have installed or upgrade the appliance, you can access it using:

- The WebShield client application.
- A web browser.

NOTE

We recommend that you use the WebShield client application. If you use a web browser, be aware that the web browser's back button will take you to the appliance's log on screen and you will lose any unsaved changes.

Using a web browser

If you are using a web browser, when entering the URL for the appliance, you must use HTTPS rather than HTTP.

Using the WebShield client application

You can point a web browser at the appliance you want to manage, and from the log on page click on the link to install the WebShield client application. Once you have installed the WebShield client application, you can close the web browser, and launch the application from the desktop icon.

Making the appliance more secure

To improve security and deter hackers, you should always change the default:

- password
- appliance name
- IP addresses

See [Initial Configuration on page 95](#) for more information.

Specifying your inside and outside networks

The appliance has an *Inside Networks* list and an *Outside Networks* list that it uses to identify whether traffic passing through it has come from an internal or external network source.

WARNING

As traffic can be scanned according to its direction, it is important that you enter the correct information for your *Inside Networks* and *Outside Networks* lists.

It is up to you to decide which networks and domains should be set up as internal networks and which networks and domains should be treated as external networks. Internal networks are typically trusted networks.

By default, all domains and networks are treated as external networks, unless they are specifically set up as internal networks.

NOTE

For security reasons, your firewall should be listed as an external network.

To scan traffic passing between the outside world and your network, you must add the networks and domains within your organization, and protected by your firewall, to the appliance's *Inside Networks* list.

For more information about setting up Internal and External networks refer to [Initial Configuration on page 95](#).

Using anti-virus scanning

The appliance uses the McAfee virus-scanning engine and virus definition (DAT) files to scan and clean network traffic. The scanners can detect both known viruses and many new viruses and variants.

Each protocol has separate inbound and outbound virus scanning.

Traffic for a specific protocol will only be scanned if that protocol is enabled and scanning has been enabled in at least one direction.

By default, all protocols are enabled and the traffic is scanned in both directions.

The appliance divides its resources between the different protocols, virus scanning each protocol's inbound and outbound traffic ([Figure 1-1 on page 27](#)).

If you disable virus scanning for either direction, traffic passes through the appliance unscanned in that direction ([Figure 1-2 on page 27](#)).

NOTE

Do not disable virus scanning for any enabled protocol *unless* you are scanning its traffic at another point in your network. Allowing unscanned traffic to enter and leave your organization leaves it open to virus infection.

Ensure that you configure your other network equipment to route the protocols through the appliance, so nothing can bypass the appliance. Only traffic that passes through the appliance, or that is routed to the appliance in the case of **Explicit Proxy** mode, can be scanned.

Scanning is described in more detail in [Content rules on page 108](#).

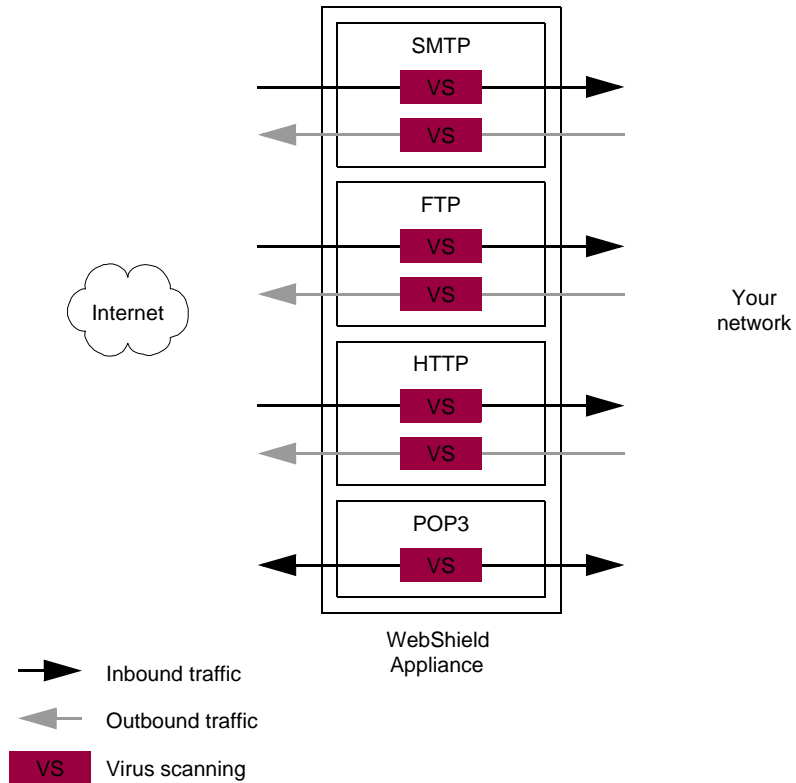


Figure 1-1. Inbound and outbound virus scanning




Figure 1-2. Outbound FTP virus scanning disabled

Controlling SMTP e-mail access

You can use the **Permit Sender** and **Deny Sender** options to specify who will be permitted or denied access to your organization via e-mail. See [Permit and deny settings on page 116](#) for more information about these options.

You can use the appliance's anti-relay features to prevent unscrupulous third parties using the appliance or the mail servers that it protects to deliver mail for them. See [Anti-relay on page 113](#).

See [E-mail \(SMTP\) delivery on page 111](#) for information about how the appliance attempts to deliver e-mail messages.

Handling unwanted content in SMTP e-mail messages

The appliance can use content rules to scan SMTP e-mail messages for undesirable content.

You create the content rules, stating what is not permitted in any message, and the appliance uses the rules to prevent such messages reaching their intended recipients.

Content scanning e-mail messages (SMTP) is described in more detail in [E-mail \(SMTP\) content scanning on page 153](#).

Scanning SMTP e-mail messages for spam

The appliance can use DNS blocking lists to block unwanted e-mail messages from particular sources. See [RBL servers on page 118](#) for more information about using DNS blocking lists.

The optional *McAfee SpamKiller for WebShield appliances* software provides additional protection from spam.

For information about evaluating SpamKiller for WebShield appliances, see [Evaluating SpamKiller on page 224](#).

If you have acquired the activation CD, needed to activate the software, see [Activating SpamKiller on page 224](#). See our web site for more information about SpamKiller for WebShield appliances and how to obtain an activation CD.

Using policies

A policy is a collection of settings and content rules that allow you to combat a specific threat to your network. You can tell the appliance what to do when a threat becomes a reality. See [Policies on page 103](#) for general information about policies.

For information about which policies you can create for each protocol, see:

- [Managing SMTP e-mail on page 111.](#)
- [Managing HTTP on page 169.](#)
- [Managing FTP on page 183.](#)
- [Managing POP3 on page 191.](#)

Sharing scanning resources

The ability to share the anti-virus and anti-spam scanning workload between appliances is called *load sharing*.

An appliance can be set up so that when it receives traffic from supported protocols it off-loads some or all of its scanning workload to other appliances.

An appliance that off-loads some or all of its scanning workload is known as a *controlling appliance*.

Appliances that receive scanning work from a controlling appliance are known as *load sharing appliances*.

The controlling appliance's virus-scanning settings are used to control virus-scanning on the load sharing appliances. That is, the virus-scanning settings on the controlling appliance override any virus-scanning settings on a load sharing appliance whenever that appliance receives traffic to scan from the controlling appliance.

If the controlling appliance is load sharing with five or more appliances, we recommend that it off-loads all of its scanning workload, so that more of its own resources can be dedicated to managing incoming traffic.

For information on installing load sharing appliances in your network see:

- [page 66](#) for appliances in **Explicit Proxy** mode.
- [page 79](#) for appliances in **Transparent Router** mode.
- [page 94](#) for appliances in **Transparent Bridge** mode.

For information on configuring load sharing, see [Load Sharing on page 215](#).

NOTE

The appliance does not use the Content Vectoring Protocol (CVP).

Monitoring the appliance

You can monitor the appliance in a number of ways. You can use:

- The Status page — summarizes the health of the appliance and the status of a number of different parameters.
- The log — the log records information that can be presented as charts and reports.
- Alerts — the appliance can be configured to generate alerts that can be used by the different network management systems monitoring the appliance. For example, the appliance can be remotely monitored by your SNMP platform, and by McAfee ePolicy Orchestrator.
- Notifications — the appliance can be configured to send e-mail and other alert messages to users and network administrators to tell them that a certain event has occurred.

For more information, see [Monitoring the appliance on page 203](#), and the protocol-specific sections in this guide.

Maintaining the appliance

The appliance allows you to save its configuration, so that it can be restored at a later time, if necessary.

It is important to perform regular maintenance on the appliance to ensure optimal performance. In most cases you can set up features that automate the maintenance tasks. See [Maintaining the Appliance on page 223](#).

Troubleshooting the appliance

If you are experiencing problems, refer the *Troubleshooting* chapter in the *Product Guide*. The *Troubleshooting* section describes the diagnostic tools you can use to identify potential problems, and answers some frequently asked questions (FAQs).

You can find our contact information in [Contacting McAfee Security & Network Associates on page 18](#).

The *Links Bar* at the top of the WebShield appliance application/web page provides some useful information about useful links to:

- Contacting support.
- Submitting a sample.
- The McAfee Virus Information library.
- Additional resources (which includes a link to list of default settings and a link to the MIB definitions).

Which Operational Mode?

2

This chapter describes the different operational modes for the appliance, and explains why it is important to choose the right mode.

The appliance can operate in one of three modes:

- **Explicit Proxy** mode.
- **Transparent Router** mode.
- **Transparent Bridge** mode.

The main differences between the operational modes are described in this chapter. They are:

- Whether communicating devices are aware of the existence of the appliance, that is, if the appliance is operating in one of the transparent modes.
- How you physically connect the appliance to your network.
- The amount of configuration you must perform to incorporate the appliance into your existing network.
- Where in the network the configuration takes place.

NOTE

The appliance can use only one of these modes. You must decide which mode to use before you install and configure the appliance.

The mode you select influences how you connect the appliance to your network and the management settings you use.

Choose a mode that suits your network configuration. If you are unsure about which mode to use, read the description of each mode in this chapter, refer to the mode-specific chapters within this Guide, or consult a network expert.

Explicit or transparent?

In **Explicit Proxy** mode some network devices must be set up to *explicitly* send traffic to the appliance. The appliance then works as a *proxy*, processing the traffic on behalf of these network devices.

For example, if the appliance is in **Explicit Proxy** mode you must explicitly configure your internal mail server to send e-mail traffic to the appliance. The appliance can then scan the e-mail traffic before forwarding it, on behalf of the sender, to the external mail server. The external mail server then forwards the e-mail to the recipient.

In a similar way, the network must be configured so that incoming e-mail messages from the Internet are delivered to the appliance rather than the internal mail server. The appliance can then scan the traffic before forwarding it, on behalf of the sender, to the internal mail server for delivery ([Figure 2-1](#))




Figure 2-1. Explicit communication

NOTE

In [Figure 2-1](#) the external mail server is in direct communication with the appliance, although traffic may pass through several network devices before reaching the appliance.

In either **Transparent Router** mode or **Transparent Bridge** mode the communicating devices are unaware of the intervention of the appliance — the appliance's operation is *transparent* to those devices.

For example, the internal mail server sends the e-mail to the external mail server. The internal mail server is unaware that the appliance has intercepted the e-mail and scanned it before forwarding it onto the external mail server for delivery.

In a similar way, the external mail server sends e-mail traffic to the internal mail server. The external mail server is totally unaware that the e-mail traffic it sent has been intercepted and scanned by the appliance ([Figure 2-2](#)).

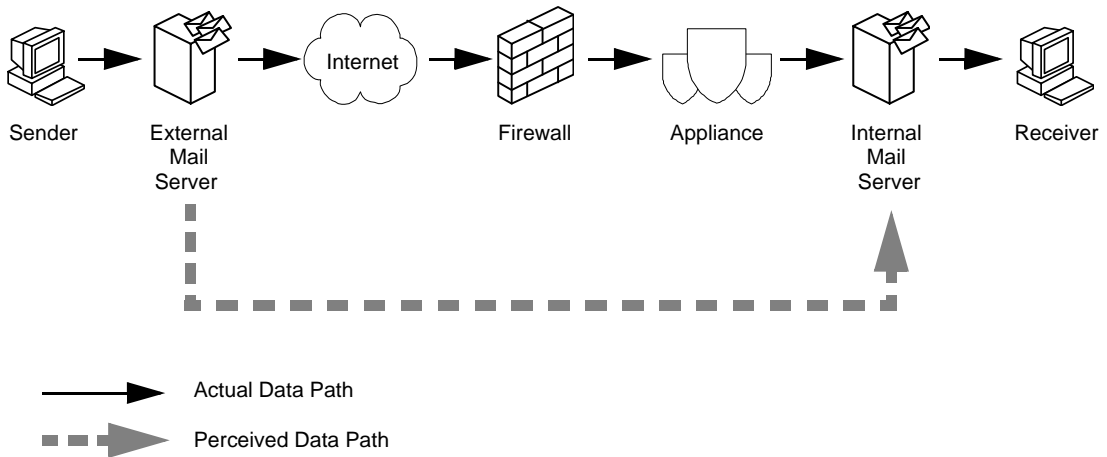


Figure 2-2. Transparent communication

NOTE

In [Figure 2-2](#) the external mail server believes it is in direct communication with the internal mail server. In reality, traffic may pass through several network devices and be intercepted and scanned by the appliance before it reaches the internal mail server.

The following sections describe each of the operational modes in more detail, and explain the consequence of changing to **Transparent Bridge** mode from any of the other modes.

Explicit Proxy mode

In **Explicit Proxy** mode, the appliance is connected to the network through the LAN1 port. LAN2 port is only used for changing the management settings of the appliance. You must explicitly configure your clients and other network devices to send SMTP, FTP, HTTP and POP3 traffic to the appliance, that is, to the IP address of LAN1 port.

NOTE

When operating in **Explicit Proxy** mode, only SMTP, FTP, HTTP and POP3 traffic should be sent to the appliance. All other traffic will be refused.

Explicit Proxy mode might not be the best option if there are a lot of network devices that must be reconfigured to send traffic to the appliance.

Explicit Proxy mode is best suited to networks where the client devices connect to the appliance through a single upstream and downstream device. For example, you can configure your network to have your web cache connect to the appliance on one side of the appliance and a firewall on the other side of the appliance, with both physically connected through the LAN1 port. The advantage of this scenario is that you only need to reconfigure the web cache and firewall. You do not need to reconfigure all of the clients.

The disadvantage of using **Explicit Proxy** mode is that it invalidates any firewall rules that you have set up for client access to the Internet. The firewall only sees the IP address information for the appliance rather than the IP addresses of the clients, which means that the firewall cannot apply its Internet access rules to the clients.

NOTE

Explicit Proxy mode does not support NetBEUI, IPX, or Multicast IP traffic.

Positioning the appliance

For security reasons, you must use the appliance **inside** your organization, **behind** a correctly configured firewall.

The location of the appliance in **Explicit Proxy** mode is not as important as the fact that all of the network devices are configured so that traffic to be scanned is sent to the appliance.




Figure 2-3. Existing network

In [Figure 2-3](#), the appliance could be connected to any of the three switches or to the router. In this case, all of the users would need to be reconfigured to direct traffic from the appropriate applications to the appliance.

The router must allow all users to connect to the appliance.

Typically the firewall would be configured to block supported protocol traffic that does not come directly from the appliance.

Configuration example

This section provides an example of the type of configuration that you would need to perform if:

- The appliance is operating in **Explicit Proxy** mode.
- You are using e-mail (SMTP).
- The appliance is connected to your network as shown in [Figure 2-4 on page 37](#).

For this configuration, you would need to:

- Configure the external Domain Name System (DNS) servers so that the external mail server knows that it must deliver mail to the appliance rather than the internal mail server.
- Configure the internal mail servers to send e-mail to the appliance. That is, the internal mail servers must use the appliance as a *smart host*.
- Make sure that your client devices can deliver e-mail to the mail servers within your organization.
- Make sure that your firewall rules are updated. You will need to make sure that the firewall will accept traffic from the appliance, but will not accept traffic that comes directly from the client devices. Set up rules to prevent unwanted traffic entering your organization.




Figure 2-4. Explicit Proxy mode configuration

For more information about how **Explicit Proxy** mode affects traffic handling, see [Explicit Proxy Mode on page 45](#).

Transparent Router mode

In **Transparent Router** mode, the appliance connects to your network using both the LAN1 and LAN2 ports. The appliance scans the traffic it receives and forwards it to the next network device. In this mode the appliance acts as a router (routing the traffic between the different networks based on the information held in its routing tables).

Unlike **Explicit Proxy** mode, you do not need to explicitly reconfigure all your network devices to send traffic to the appliance. All you need to do is configure the routing table for the appliance, and modify some of the routing information for the network devices either side of it (the devices connected to its LAN1 and LAN2 ports). For example, you might need to make the appliance your default gateway.

Transparent Router mode is suitable for networks that have firewall rules, as the firewall still sees the IP addresses of the clients and can therefore apply the Internet access rules to client traffic.

NOTE

Transparent Router mode does not support NetBEUI, IPX, or Multicast IP traffic.

Positioning the appliance

For security reasons, you must use the appliance **inside** your organization, **behind** a correctly configured firewall.

In **Transparent Router** mode we recommend that you position the appliance between the firewall and your router (Figure 2-5).




Figure 2-5. Recommended location in Transparent Router mode

In **Transparent Router** mode the appliance must join two networks ([Figure 2-6 on page 40](#)).

The appliance can be set up to route traffic to other networks. The networks routing traffic through the appliance are unaware that the appliance is intercepting and scanning the traffic, and that is why the appliance is said to be operating as a *transparent router*.




Figure 2-6. Routing between networks

Configuration example

This section provides an example of the type of configuration that you would need to perform if:

- The appliance is operating in **Transparent Router** mode.
- You are using e-mail (SMTP).
- The appliance is connected to your network ([Figure 2-5 on page 39](#)).

You would need to:

- Configure your client devices to point to the default gateway. In [Figure 2-5 on page 39](#) the client's default gateway is the router. See the user documentation that accompanied the client for further details about setting up proxy servers.
- Configure the appliance to use the Internet gateway as its default gateway. In [Figure 2-5 on page 39](#) the Internet gateway is the firewall, and the router's default gateway is the appliance.
- Make sure that your client devices can deliver e-mail to the mail servers within your organization.

For more information about how **Transparent Router** mode affects traffic handling, see [Transparent Router Mode on page 67](#).

Transparent Bridge mode

In **Transparent Bridge** mode the appliance connects to your network using both the LAN1 and LAN2 ports. The appliance scans the traffic it receives, and acts as a bridge that is connecting two separate networks, but treating them as a single network.

Transparent Bridge mode requires even less configuration than the **Transparent Router** mode. You do not need to reconfigure all of your clients or default gateway to send traffic to the appliance. As the appliance is not a router in this mode you will not need to update a routing table. You will still need to define the internal and external networks, but this is true of all operational modes.

Positioning the appliance

For security reasons, you must use the appliance **inside** your organization, **behind** a correctly configured firewall.

In **Transparent Bridge** mode we recommend that you position the appliance between the firewall and your router ([Figure 2-7 on page 42](#)).




Figure 2-7. Transparent Bridge mode

In **Transparent Bridge** mode you physically connect two network segments to the appliance and the appliance treats them as one logical network. As the devices are on the same logical network, they must all have compatible addresses for that network ([Figure 2-8 on page 43](#)).

Devices on one side of the bridge that communicate with devices on the other side of the bridge are unaware of the bridge's existence. They are unaware that the traffic is intercepted and scanned, and that is why the appliance is said to be operating as a *Transparent Bridge*.

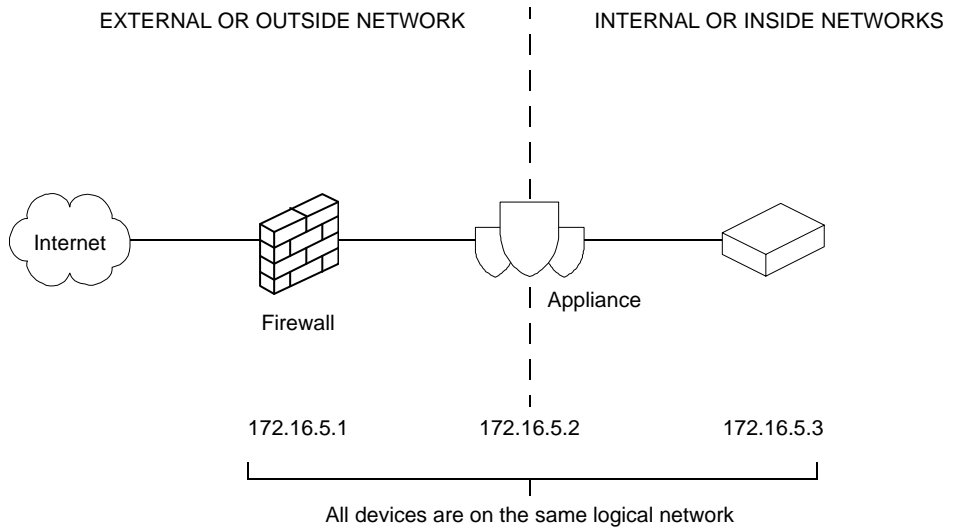


Figure 2-8. Single logical network

For more information about how **Transparent Bridge** mode affects traffic handling, see [Transparent Bridge Mode on page 81](#).

Changing operational modes

NOTE

Once you have selected an operational mode, we recommend that you do not change it unless you are moving the appliance or restructuring your network.

In **Explicit Proxy** mode and **Transparent Router** mode you can set up your appliance so that it sits on more than one network. You do this by setting up multiple IP addresses for LAN1 and LAN2 ports. If you change from **Explicit Proxy** or **Transparent Router** mode to **Transparent Bridge** mode, only the enabled IP addresses for each port (LAN1, LAN2) will be carried over into **Transparent Bridge** mode.

This chapter provides a few scenarios for the protocols to illustrate how you can integrate the appliance with your existing network when you want to use it in **Explicit Proxy** mode.

NOTE

To scan a supported protocol, you must configure your other network equipment or client computers to route that protocol through the appliance, so that nothing can bypass the appliance.

In **Explicit Proxy** mode you should only route SMTP, FTP, HTTP and POP3 traffic through the appliance. All other traffic will be refused.

You must use the appliance **inside** your organization and **behind** a correctly configured firewall.

If you are unsure about your network's topology and how you should integrate the appliance, consult your network expert.

This appliance offers even greater flexibility when used with other appliances to:

- Separate inbound and outbound scanning
- You can use two appliances to independently handle inbound and outbound scanning, providing a directional fail-safe configuration. See [Directional scanning on page 47](#).
- Increase throughput
- You can increase the available scanning throughput by adding more appliances as shown in the multiple appliance configurations.
- Share the scanning workload
- You can also share the scanning workload between appliances. See [Load sharing on page 66](#).

SMTP Scenarios

This section describes some SMTP configurations for when your appliance is in **Explicit Proxy** mode. This section describes:

- [One site.](#)
- [Directional scanning on page 47.](#)
- [Multiple sites on page 49.](#)
- [Demilitarized zone \(DMZ\) on page 50.](#)
- [International organization on page 52.](#)
- [Handling network failure on page 53.](#)

One site

Figure 3-1 shows a very common configuration for small companies with a few hundred users.




Figure 3-1. One site

NOTE

For this configuration, you must ensure that users cannot directly send outbound messages to the firewall or appliance. This can be enforced with a firewall policy.

Configure the firewall, appliance and mail server to relay inbound and outbound messages as shown in [Figure 3-2 on page 47](#).




Figure 3-2. Flow of e-mail messages

Figure 3-2 shows a fail-safe configuration:

- If the appliance's network connection fails between the appliance and the mail server, the outbound messages remain at the mail server, and the inbound messages are stored at the appliance. They will not be delivered unscanned to your mail server. This configuration ensures that unscanned e-mail messages cannot enter or leave your mail system.
- If the appliance's network connection fails between the appliance and the firewall, outbound mail is held at the appliance and inbound mail is held at the external mail server that is attempting to deliver the mail to you.

Directional scanning

Figure 3-3 and Figure 3-4 on page 48 show how you can configure your site for directional scanning. One appliance is dedicated to inbound scanning, and the other appliance is dedicated to outbound scanning.

Using two appliances doubles the available scanning throughput.

NOTE

For this configuration, you must ensure that users cannot directly send outbound messages to the firewall or appliance. This can be enforced with a firewall policy.




Figure 3-3. Directional scanning

Figure 3-4 shows a directional fail-safe configuration.




Figure 3-4. Directional flow of e-mail messages

Both appliances operate independently, as follows:

- **Appliance 1**

If the appliance's network connection fails between the appliance and the mail server, inbound messages are stored at the appliance. They will not be delivered unscanned to your mail server. This configuration ensures that unscanned e-mail messages cannot enter your mail system.

If the appliance's network connection fails between the firewall and the appliance, inbound mail is held at the external mail server that is attempting to deliver the mail to you.

■ Appliance 2

If the appliance's network connection fails between the mail server and the appliance, the outbound messages remain at the mail server. This configuration ensures that unscanned e-mail messages cannot leave your mail system.

If the appliance's network connection fails between the appliance and the firewall, outbound mail is held at the appliance.

Configure the firewall, appliances and mail servers to relay inbound and outbound messages as shown in [Figure 3-4 on page 48](#).

Multiple sites

[Figure 3-5](#) shows a typical configuration for a company that has multiple sites connected by a WAN (Wide Area Network) link. Each site has a mail server that handles the site's e-mail messages, but only the main site has a link to the Internet.




Figure 3-5. Multiple sites

Configure the firewall, appliance and mail servers to relay inbound and outbound messages, as shown in [Figure 3-6 on page 50](#).

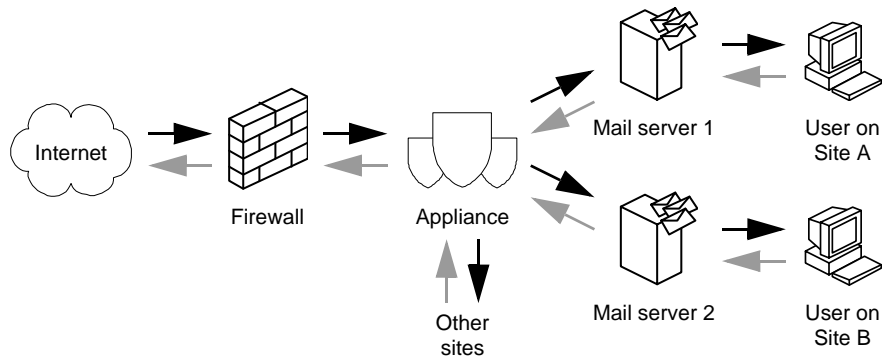


Figure 3-6. Multiple site traffic flow

Figure 3-6 shows a fail-safe configuration, ensuring that unscanned e-mail messages cannot enter or leave your mail system if the appliance's network connection fails.

You can introduce a second appliance to double the mail throughput capacity and create a directional fail-safe configuration. Directional scanning is explained in [Directional scanning on page 47](#). In directional scanning both appliances operate independently, being dedicated to either inbound or outbound scanning.

For a directional fail-safe configuration, configure the firewall, appliances and mail servers to relay inbound and outbound messages as shown in [Figure 3-4 on page 48](#). This provides the potential for greater throughput and reliability.

Demilitarized zone (DMZ)

Figure 3-7 and [Figure 3-8 on page 51](#) show how you can increase network security. The configuration is suitable for large organizations with a DMZ security model, but should not be attempted by inexperienced users, as it involves some complex configuration at the firewall.

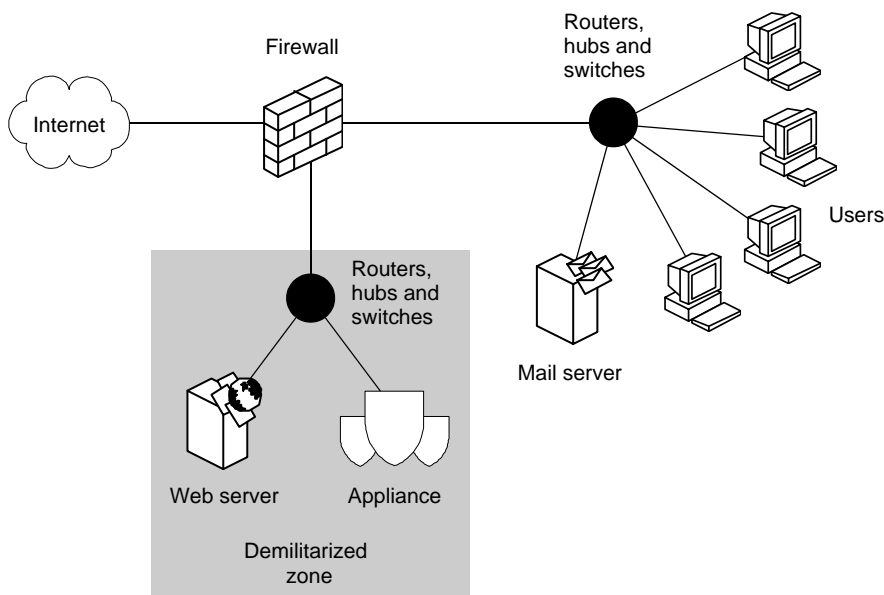


Figure 3-7. Demilitarized zone

Configure the firewall, appliance and mail server to relay inbound and outbound messages as shown in [Figure 3-8](#).




Figure 3-8. Flow of e-mail messages for demilitarized zone

The firewall routes the messages to and from the appliance, so each message travels through the firewall twice. The firewall is configured to accept no e-mail connections other than those from the mail server relays.

International organization

Figure 3-9 shows sites in different countries connected by a WAN (Wide Area Network) link. There are multiple appliances and multiple links, adding fault tolerance to the mail setup and ensuring that e-mail messages can still be scanned and delivered if one of the connections to an appliance or the Internet fails. Messages for any of the network domains can be received through any Internet link. The main mail server, called the *bridgehead* mail server, holds the main directory or alias list for the entire organization, and routes the e-mail messages to the correct internal mail servers.

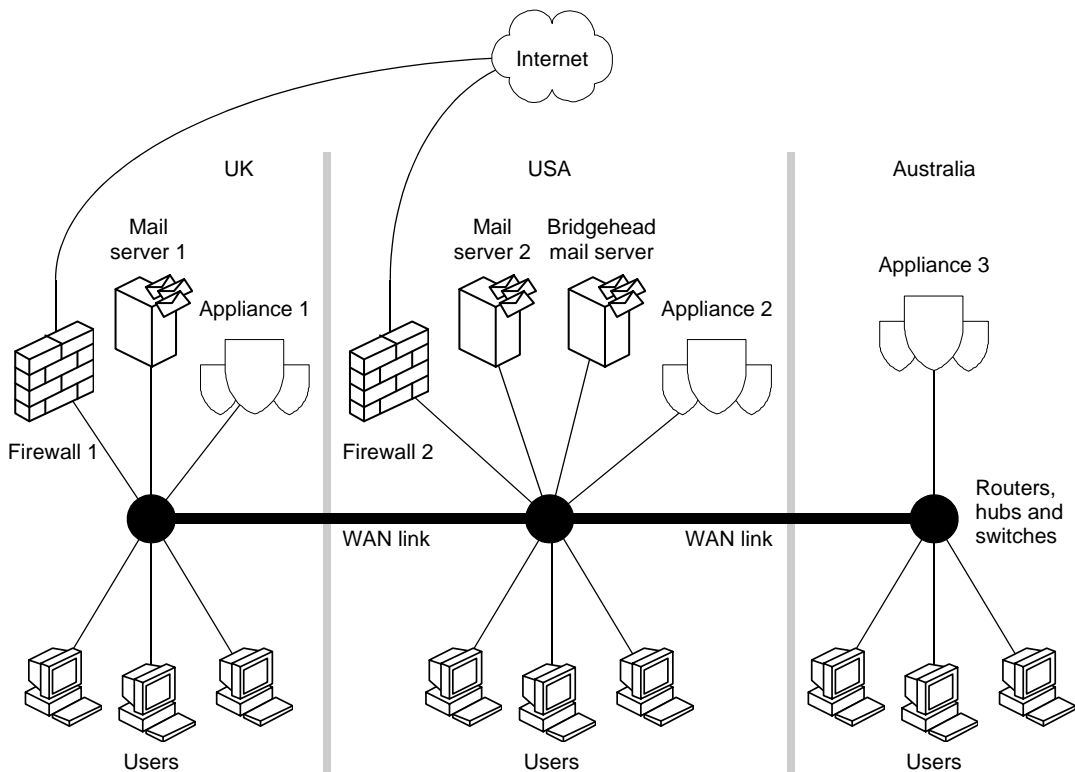


Figure 3-9. International organization

Handling network failure

This section describes three different ways to handle SMTP traffic when your network fails to deliver SMTP traffic to the appliance, or the appliance cannot process that traffic.

The appliance might not be able to process the supported protocol traffic if:

- There are internal network problems, such as a slow network or broken connection.
- There is a high quantity of network traffic, temporarily filling the appliance's buffers.
- The appliance is turned off or has been restarted. For example, you could be in the process of restoring its software or applying IP configuration changes.

The different ways to handle supported traffic are described in:

- [Fail-over](#).
- [Fail-closed on page 54](#).
- [Fail-open on page 55](#).

NOTE

All methods require modification to the records in your Domain Name System (DNS) server. See your DNS server software documentation to find out how to modify DNS records.

Fail-over

To provide fault tolerance, a fail-over configuration requires two or more appliances. If the first appliance cannot be reached, traffic is forwarded to the second appliance ([Figure 3-10 on page 54](#)). The second appliance scans the traffic. We strongly recommend that you use this type of configuration if you require continual anti-virus scanning and traffic flow.




Figure 3-10. Fail-over

For this configuration, you must modify your DNS server's MX (Mail Exchanger) records as follows:

- Add the first appliance, providing it with the highest priority so that it receives traffic first.
- Add the second appliance, providing it with a lower priority so that it receives traffic when the first appliance cannot be reached.
- Remove records for your servers and other equipment (except your firewall), so that nothing can bypass the appliances.

Fail-closed

The fail-closed configuration (Figure 3-11) is also known as *fail-safe*, and it ensures that unscanned traffic cannot enter or leave your organization if the appliance cannot be reached.



Figure 3-11. Fail-closed

We recommend that you use this type of configuration if you have just one appliance because it maintains your organization's anti-virus protection.

For this configuration, you must modify your DNS server's records as follows:

- Add the appliance, so that it receives traffic.
- Remove records for your servers and other equipment (except your firewall), so that nothing can bypass the appliance.

Fail-open

The fail-open configuration (Figure 3-12) ensures that traffic can enter or leave your organization if the appliance cannot be reached.




Figure 3-12. Fail-open

WARNING

We do not recommend that you use this type of configuration because it can allow infected material to enter or leave your organization unscanned. If possible, add more appliances to form a fail-over configuration.

Ensure that you have McAfee anti-virus software protecting your mail servers, file servers, clients (users) and other equipment.

For this configuration, you must modify your DNS server's records as follows:

- Add the appliances, providing them with the highest priorities so that they receive traffic first.
- Add your servers, firewall and other equipment, providing them with lower priorities so that they receive traffic when the appliances cannot be reached.

FTP scenarios

This section presents the following FTP scenarios:

- *Outbound FTP communication.*
- *Inbound FTP with one server on page 57.*
- *Inbound FTP with multiple servers on page 58.*

NOTE

We recommend that you protect your incoming FTP information (PUT commands to your FTP server) with the appliance.

For bi-directional (inbound and outbound) FTP scanning, you must use either true FTP clients, that use software such as CuteFTP, or command-line FTP clients.

Some web browsers send FTP requests over HTTP connections. In this case, you can only perform FTP downloads over HTTP. That is, only FTP GET requests can be made over an HTTP connection.

To perform an FTP download over an HTTP connection, you must configure your client browsers' FTP proxies to use the same server and port number that the appliance uses for the HTTP proxy (port 80 by default).

Outbound FTP communication

Figure 3-13 shows the communication that occurs when an internal user accesses an internal FTP server (within the organization) or external FTP server (outside the organization).




Figure 3-13. Outbound FTP

In **Explicit Proxy** mode, the appliance's FTP proxy is of the type "user@host" and users must connect to the FTP server using both their user name and the host name, separated by the @ symbol, for example `user@ftp.example.com`.

NOTE

If your firewall is also of the type "user@host", you can configure the appliance's handoff host as the firewall, so that the appliance appears and acts (to the connecting user) exactly as the firewall would. The appliance still scans the FTP traffic for viruses, using its FTP scan settings.

For this configuration, you must configure the FTP client to use the appliance as a proxy server.

Inbound FTP with one server

Figure 3-14 shows the communication that occurs when an external user accesses an internal FTP server.




Figure 3-14. Inbound FTP with one server

The user's FTP client is sent to the firewall by an external DNS server (on the Internet). The firewall forwards all traffic to the appliance, which uses its handoff host to forward traffic to the FTP server. When the traffic returns, the appliance scans it for viruses, using its *inbound* FTP scan settings.

The communication is transparent to the user. The user enters an FTP command, for example:

```
ftp ftpserver1.example.com
```

The user's FTP client then attaches to the FTP server using its domain name. For this configuration, you must perform the following actions:

- Create records for the external DNS server that match the FTP server to the firewall.
- Configure the firewall to only send and receive internal FTP traffic to and from the appliance, so that nothing can bypass it.
- Configure the appliance with the FTP server as the inbound handoff host.

Inbound FTP with multiple servers

If your organization has multiple FTP servers, you can use one appliance for each FTP server. Figure 3-15 shows the communication that occurs when an external user accesses the internal FTP servers through their dedicated appliances.




Figure 3-15. Inbound communication with multiple servers

An external DNS server (on the Internet) sends the user's FTP client to the firewall. The firewall forwards all traffic to the appropriate appliance, which uses its handoff host to forward traffic to the FTP server. When the traffic returns, the appliance scans it for viruses, using its *inbound* FTP scan settings.

The communication is transparent to the user. The user enters an FTP command, for example:

```
ftp ftpserver1.example.com
```

The user's FTP client then attaches to the FTP server using its domain name.

For this configuration, you must perform the following actions:

- Create records for the external DNS server that match the FTP servers to the firewall.
- Configure the firewall to only send and receive internal FTP traffic to and from the appliance, so that nothing can bypass it.
- Configure each appliance with the appropriate FTP server as the inbound handoff host.

HTTP scenarios

This section presents the following HTTP scenarios:

- *Outbound HTTP with internal web cache.*
- *Outbound HTTP with external web cache on page 60.*
- *Outbound HTTP without web cache on page 62.*
- *Inbound HTTP on page 62.*

NOTE

By default, the appliance only handles and scans HTTP traffic on port 80.

Outbound HTTP with internal web cache

Figure 3-16 shows the communication that occurs when an internal user accesses an external web server and there is a web cache between the user and the appliance.

The benefit of using this configuration is that, once cached, web pages do not need to be rescanned each time they are used. The amount of traffic passing through the appliance is reduced, which improves the throughput for other traffic being scanned.




Figure 3-16. Outbound HTTP with internal web cache

The user makes an HTTP request to the web cache. For example, the user requests a web page. If the page has already been cached, the web cache passes the page to the user. If the page has not already been cached, the web cache requests that page from the appliance. The appliance scans the request and it passes to the firewall. The firewall checks the request against the network security policy, and if it complies, passes the request to the external web server.

The response, in this case, the web page, is also intercepted and scanned by the appliance before it is stored in the web cache, and passed to the user.

The next time someone requests the same web page, the cached and clean copy of that page can be downloaded without having to pass through the appliance, and therefore without having to be rescanned.

For this configuration, you must perform the following actions:

- Configure the users' web browsers to send requests to the web cache.
- Configure the firewall to only allow HTTP traffic from the appliance, and to only send HTTP traffic to the appliance, so that HTTP traffic cannot bypass the appliance and firewall.

Outbound HTTP with external web cache

Figure 3-17 shows the communication that occurs when an internal user accesses an external non-transparent web cache. This configuration is common in organizations that use an *Internet Service Provider (ISP)* to provide web caching facilities.

Cached web pages will still be scanned by the appliance, as they still pass through the appliance.

NOTE

If you have an external web cache configuration, we recommend that you install an internal web cache between the users and the appliance to reduce the need for scanning.

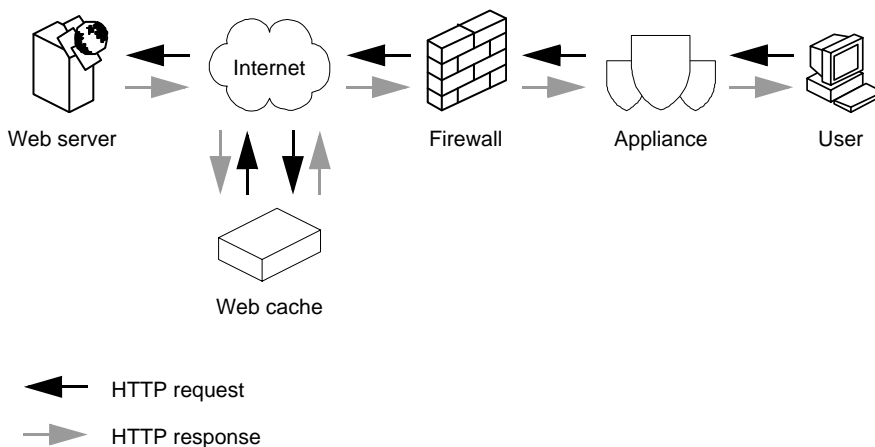


Figure 3-17. Outbound HTTP with external web cache

The user makes an HTTP request. For example, the user requests a web page. The request is scanned by the appliance and passes to the handoff host via the firewall. In this case the appliance's handoff host is the web cache.

The firewall checks the request against the network security policy, and if it complies, passes the request to the external web cache.

NOTE

Setting up an HTTP handoff host is described in [Handoff host on page 176](#).

If the web page has already been cached it is sent to the user via the firewall. The firewall will check its network security policy and if the connection is allowed the web page is passed to the appliance where it will be scanned before being passed to the user.

If the page has not been cached, the request is passed to the web server. The web server passes the page to the web cache. The web cache stores the web page before passing it back to the appliance via the firewall.

The firewall checks its network security policy and if the connection is allowed, the web page is sent to the appliance where it is scanned before being passed to the user.

For this configuration, you must perform the following actions:

- Configure the users' web browsers to send requests to the appliance.
- Configure the firewall to only allow HTTP traffic from the appliance, and to only send HTTP traffic to the appliance, so that HTTP traffic cannot bypass the appliance and firewall.
- Configure the appliance to use the web cache as its outbound handoff host. If you do not, the appliance tries the default route when connecting to external web servers, and in this case, bypasses the web cache.

NOTE

See the documentation that accompanies your web browsers for details of how to send traffic to the appliance.

For information about configuring firewall rules, see the firewall's user documentation.

Setting up an HTTP handoff host is described in [Handoff host on page 176](#).

If the external web cache is a *transparent* web cache configure your network as described in [Outbound HTTP without web cache on page 62](#).

Outbound HTTP without web cache

Figure 3-18 shows the communication that occurs when an internal user accesses an external web server and web pages are not cached.




Figure 3-18. Outbound HTTP without web cache

The user makes an HTTP request. For example, the user requests a web page. The appliance scans the request and the request passes to the firewall.

The firewall checks its network security policy and if the connection is allowed, the request is passed to the external web server.

The response, in this case the web page, passes back to the appliance via the firewall. The firewall checks its network security policy and if the connection is allowed, the web page passes to the appliance. The appliance scans the web page before passing it to the user.

For this configuration, you must perform the following actions:

- Configure the users' browsers to send HTTP requests to the appliance.
- Configure the firewall to only allow HTTP requests from the appliance, and to only send HTTP responses to the appliance, so that HTTP traffic cannot bypass the firewall and appliance.

NOTE

See the documentation that accompanies your web browsers for details of how to send traffic to the appliance.

For information about configuring firewall rules, see the firewall's user documentation.

Inbound HTTP

We do not recommend setting up inbound HTTP if your appliance is in **Explicit Proxy** mode. If you need to set up inbound HTTP, we recommend that you use an appliance in Transparent Router mode or **Transparent Bridge** mode.

POP3 scenarios

This section presents the following POP3 scenarios:

- [Using a generic connection.](#)
- [Using a dedicated connection on page 64.](#)
- [POP3 with multiple servers on page 65.](#)

NOTE

The appliance scans mail download (POP3) traffic without distinguishing its direction (inbound or outbound).

Using a generic connection

Figure 3-19 shows the communication that occurs when an internal user accesses an external POP3 server. For this type of configuration you would typically use a *generic* POP3 connection. A generic connection is a connection that by default uses port 110 to transmit and receive POP3 traffic.




Figure 3-19. Generic POP3 connection

The internal user makes a POP3 request to download mail. The user's POP3 client asks the appliance to log on to the external POP3 server and download mail on behalf of the user.

The request contains:

- A special form of the user name that is used to access the user's mail account on the POP3 server. An example of the special user name format is:
`username#pop3.example.com.`
- The name or address of the specific POP3 server.
- The password for the user's mail account on that POP3 server.

The appliance scans the request, and uses information in it to work out where to send it. In this case the request is sent to the external POP3 server, via the firewall.

The firewall, checks the request against the networks' security policy and if the request complies, the request passes to the POP3 server.

The appliance can then log on to the POP3 server and begin to download the mail on behalf of the internal user.

The response, in this case, mail, passes back to the appliance via the firewall, and is scanned before it passes to the user.

For this configuration, you must configure the users' POP3 clients to send POP3 traffic to the appliance.

Using a dedicated connection

Figure 3-20 shows the communication that occurs when an external user accesses an internal POP3 server. For this type of configuration you would typically use a *dedicated* POP3 connection. A dedicated POP3 connection is a connection that has been specially set up by the network administrator to force users to use a specific POP3 server.

NOTE

The dedicated connection should not use the generic POP3 port number or any other port number that is already in use on your appliance.




Figure 3-20. Dedicated POP3 connection

The external user makes a POP3 request to download e-mail. When the request reaches the firewall, the firewall applies a rule. The rule states that when POP3 requests are received, the requests are always diverted to the appliance.

The appliance scans the request and passes the request to the POP3 server.

The response passes to the appliance, where it is scanned before passing to the external user via the firewall.

For this configuration, you must perform the following actions:

- Configure the firewall to divert POP3 traffic to the appliance, so that nothing can bypass it.
- Configure the appliance with a dedicated connection for the specific POP3 server.

NOTE

For information about configuring firewall rules, see the firewall's user documentation.

POP3 with multiple servers

If your organization has multiple POP3 servers, you can configure them using one of the following methods:

- Configure the appliance with dedicated connections for each POP3 server, as described in [Using a dedicated connection on page 64](#).
 - ◆ Each dedicated connection must use a different port number. You might need to enable the use of these ports on the firewall, and tell external users which port number to use.
 - ◆ You can also set up firewall rules so that inbound traffic from a specific range of IP addresses will always be forwarded to the relevant appliance.
- Use one appliance for each POP3 server, as shown in [Figure 3-21](#). For this configuration, you must perform the following actions:
 - ◆ Set up your firewall rules to allow incoming connections to be diverted to the relevant appliance.
 - ◆ Configure each appliance with a dedicated connection for the appropriate POP3 server.




Figure 3-21. POP3 with dedicated appliances

Load sharing

If you have more than one appliance you can set them up to share the scanning workload. If the appliances are in **Explicit Proxy** mode they should be configured as shown in [Figure 3-22](#).




Figure 3-22. Load sharing Explicit Proxy configuration

NOTE

All of the appliances are connected through LAN1 port to the switch or router. The IP addresses are for example only.

For more information on configuring the appliances for load sharing, see [Load Sharing on page 215](#).

This chapter provides a few scenarios for the protocols to illustrate how you can integrate the appliance with your existing network when you want to use it in **Transparent Router** mode.

You can also share the scanning workload between appliances, by setting up the appliances as described in [Load sharing on page 79](#).

NOTE

To scan a supported protocol, you must ensure that traffic for that protocol passes through the appliance.

For security reasons, you must use the appliance **inside** your organization and **behind** a correctly configured firewall.

If you are unsure about your network's topology and how you should integrate the appliance, consult your network expert.

SMTP scenarios

This section describes some SMTP configurations for when your appliance is in **Transparent Router** mode. This section describes:

- [One site on page 68](#).
- [Multiple sites on page 69](#).
- [Demilitarized zone \(DMZ\) on page 70](#).
- [International organization on page 71](#).

One site

Figure 4-1 shows a very common configuration for small companies with a few hundred users.




Figure 4-1. One site

Configure the firewall, appliance, and mail server to relay inbound and outbound messages through the appliance as shown in Figure 4-2.




Figure 4-2. Flow of e-mail messages

Figure 4-2 on page 68 shows a fail-safe configuration:

- If the appliance's network connection fails between the appliance and the mail server, the outbound messages remain at the mail server, and the inbound messages are stored at the appliance. They will not be delivered unscanned to your mail server. This configuration ensures that unscanned e-mail messages cannot enter or leave your mail system.
- If the appliance's network connection fails between the appliance and the firewall, outbound mail is held at the appliance and inbound mail is held at the external mail server which is attempting to deliver the mail to you.

Multiple sites

Figure 4-3 shows a typical configuration for a company that has multiple sites connected by a WAN (Wide Area Network) link. Each site has a mail server that handles the site's e-mail messages, but only the main site has a link to the Internet.




Figure 4-3. Multiple sites

NOTE

Traffic between Site A and Site B will not be scanned, because traffic will never be routed to the appliance.

If you want traffic between internal sites to be scanned, we recommend that you install another appliance between the sites.

Configure the firewall, appliance and mail servers to relay inbound and outbound messages, as shown in [Figure 4-4](#).




Figure 4-4. Multiple sites traffic flow

This is a fail-safe configuration, ensuring that unscanned e-mail messages cannot enter or leave your mail system if the appliance's network connection fails.

Demilitarized zone (DMZ)

We do not recommend that you try to set up a demilitarized zone security model using appliances in **Transparent Router** mode. You should only set up a DMZ security model using appliances in **Explicit Proxy** mode. See [Demilitarized zone \(DMZ\)](#) on page 50.

International organization

Figure 4-5 shows sites in different countries connected by a WAN (Wide Area Network) link. There are multiple appliances and multiple links, adding fault tolerance to the mail setup and ensuring that e-mail messages can still be scanned and delivered if one of the connections to an appliance or the Internet fails. Messages for any of the network domains can be received through any Internet link. The main mail server, called the *bridgehead* mail server, holds the main directory or alias list for the entire organization, and routes the e-mail messages to the correct internal mail servers.

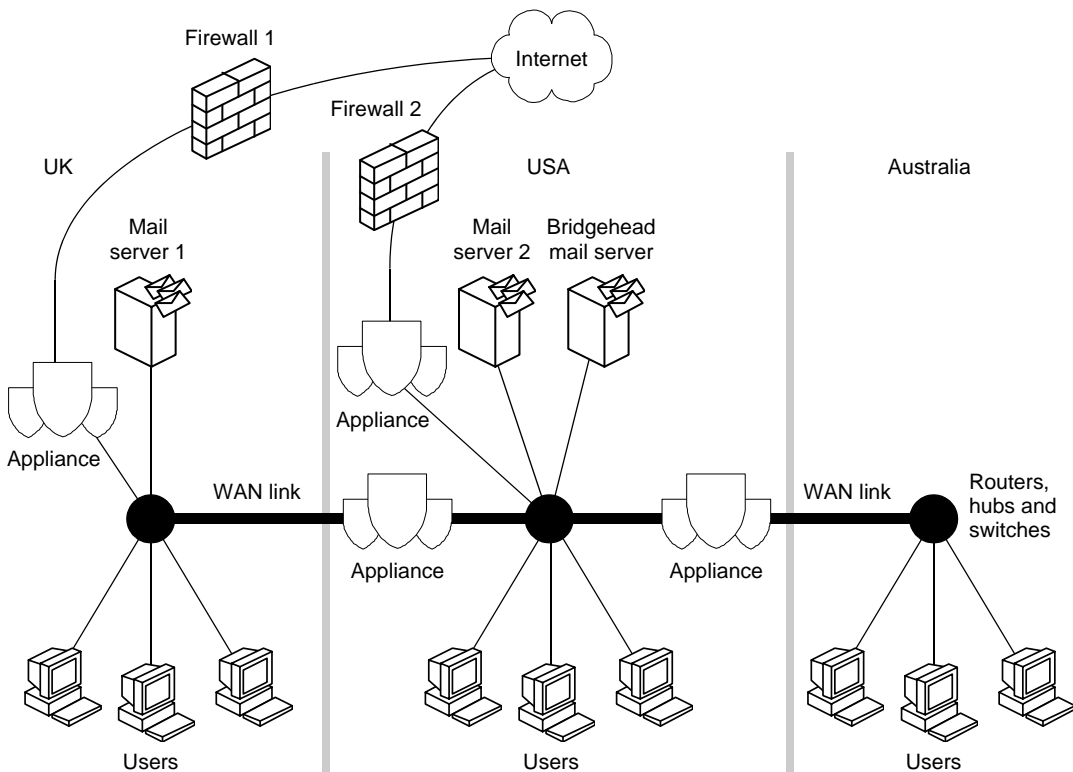


Figure 4-5. International organization

FTP scenarios

This section presents the following FTP scenarios:

- *Outbound FTP.*
- *Inbound FTP on page 73.*

NOTE

We recommend that you protect your incoming FTP information (PUT commands to your FTP server) with the appliance.

Outbound FTP

Figure 4-6 shows the communication that occurs when an internal user initiates an FTP connection with an external FTP server.




Figure 4-6. Outbound FTP

The internal user's FTP request is sent to the external FTP server. The appliance intercepts and scans the request. When the firewall receives this request, if it obeys its security policy, the request passes through the firewall to the Internet and external FTP server.

The FTP response is also intercepted and scanned by the appliance.

Inbound FTP

Figure 4-7 shows the communication that occurs when an external user accesses an internal FTP server.




Figure 4-7. Inbound FTP with one server

The external user's FTP request is sent to the FTP server. When the firewall receives this request, if it obeys its security policy, the request passes through the firewall. Before the request reaches the FTP server the appliance intercepts and scans the request.

The FTP response is also intercepted and scanned by the appliance.

HTTP scenarios

This section presents the following HTTP scenarios:

- *Outbound HTTP with internal web cache on page 74.*
- *Outbound HTTP with external web cache on page 75.*
- *Outbound HTTP without web cache on page 76*
- *Inbound HTTP on page 77.*

NOTE

By default, the appliance only handles and scans HTTP traffic on port 80.

Outbound HTTP with internal web cache

Figure 4-8 shows the communication that occurs when an internal user accesses an external web server and there is a web cache between the user and the appliance.

The benefit of using this configuration is that, once cached, web pages do not need to be rescanned each time they are used. The amount of traffic passing through the appliance is reduced, which improves the throughput for other traffic being scanned.




Figure 4-8. Outbound HTTP with internal web cache

The user makes an HTTP request to the web cache. For example, the user requests a web page.

If the page has already been cached, the web cache passes the page to the user.

If the page has not already been cached, the web cache requests that page from the web server.

The request is intercepted and scanned by the appliance.

If the request meets the network's security policy it will also pass through the firewall to the Internet and external web server.

The response, in this case, the web page, is also intercepted and scanned by the appliance before it is stored in the web cache, and passed to the user.

The next time someone requests the same web page, the cached and clean copy of that page can be downloaded without having to pass through the appliance, and therefore without having to be rescanned.

Outbound HTTP with external web cache

Figure 4-9 shows the communication that occurs when an internal user accesses an external web cache. This configuration is common in organizations that use an *Internet Service Provider (ISP)* to provide web caching facilities.

Cached web pages will still be scanned by the appliance, as they still pass through the appliance.

NOTE

If you have an external web cache configuration, we recommend that you install an internal web cache between the users and the appliance to reduce the need for scanning.

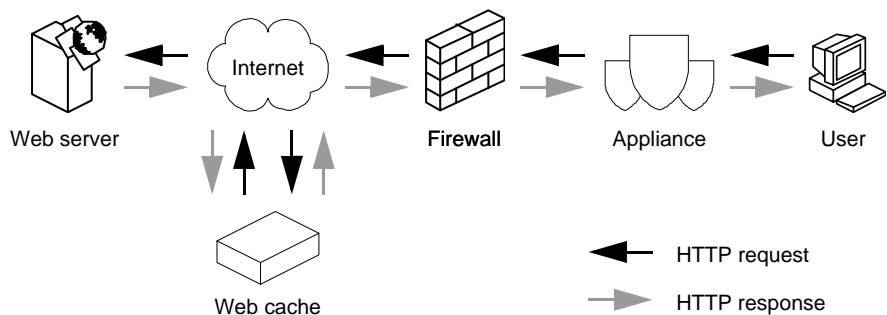


Figure 4-9. Outbound HTTP with external web cache

The user makes an HTTP request to the web cache. For example, the user requests a web page. The request is intercepted and scanned by the appliance.

If the request meets the network's security policy it will pass through the firewall to the web cache.

If the requested page has already been cached, the web cache will pass the page to the user. The page will be intercepted and scanned by the appliance before it passes to the user.

If the requested page has not already been cached, the web cache will pass the request to the web server.

The response, in this case, the web page, is passed to the web cache and stored before being passed to the user.

Before the page reaches the user, the firewall checks that the connection meets the network's security policy, and the appliance intercepts and scans the page.

If the external web cache is a *transparent* web cache, you do not need to perform any other configuration. If the web cache is not transparent, you should set up the appliance to use the web cache as its *handoff host*, or set up the users' browsers to send requests to the web cache.

NOTE

Setting up an HTTP handoff host is described in [Handoff host on page 176](#).

See the documentation that accompanies your web browsers for details of how to send traffic to the web cache.

Outbound HTTP without web cache

[Figure 4-10](#) shows the communication that occurs when an internal user accesses an external web server and web pages are not cached.




Figure 4-10. Outbound HTTP

The user's HTTP request is sent to the external web server, and is intercepted and scanned by the appliance.

The response is also intercepted and scanned by the appliance.

For this configuration, you ensure that the HTTP traffic you want scanned is routed through the appliance to the Internet.

Inbound HTTP

Figure 4-11 shows the communication that occurs when an external user accesses an internal web server.




Figure 4-11. Inbound HTTP

The external user sends an HTTP request to the internal web server. For example, the user requests a web page.

If the request meets the network's security policy it will pass through the firewall and be intercepted and scanned by the appliance before reaching the web server.

The response, in this case, the web page, is also intercepted and scanned by the appliance.

If you only have one web server that you want to be accessed by external users, set up the appliance to use the web server as its inbound *handoff host*. Setting up a handoff host provides added security, as external users can only access the specific web server and cannot access unauthorized web servers on your network.

NOTE

Setting up an HTTP handoff host is described in [Handoff host on page 176](#).

If you have more than one web server that you want to be accessed by external users, you should not set up the web server as the handoff host for the appliance. Instead, configure your firewall to only allow HTTP connections to the appropriate web servers, and route that traffic through the appliance.

POP3 scenarios

The appliance scans mail download (POP3) traffic without distinguishing its direction (inbound or outbound).

Figure 4-12 shows the communication that occurs when an internal user accesses an external POP3 server.




Figure 4-12. POP3 flow of traffic

The user makes a POP3 request to download mail from the external POP3 server. The request is intercepted and scanned by the appliance.

If the request meets the network's security policy it will pass through the firewall to the Internet and external POP3 server.

The response, in this case, mail messages, is also intercepted and scanned by the appliance.

Load sharing

If you have more than one appliance you can share the scanning workload between the appliances. If the controlling appliance is in **Transparent Router** mode, the appliances should be configured as shown in [Figure 4-13](#).

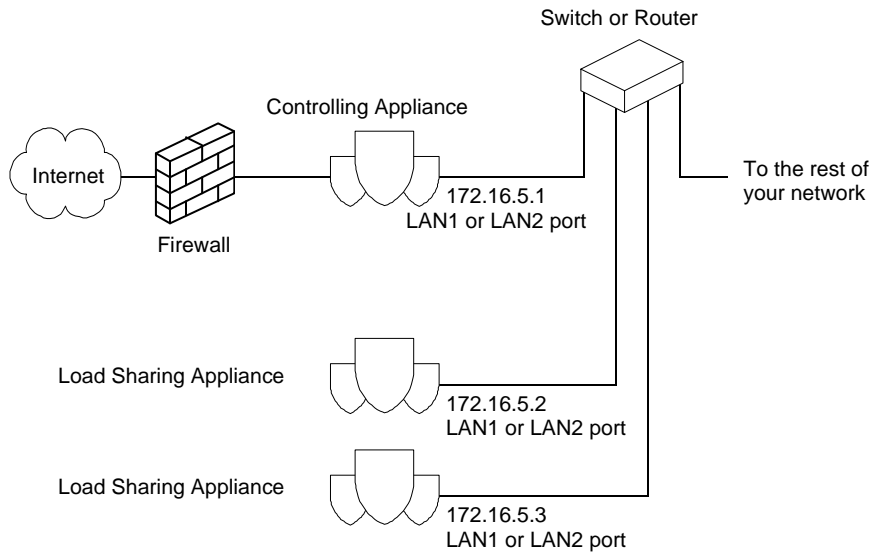


Figure 4-13. Load sharing Transparent Router configuration

NOTE

The IP addresses shown in [Figure 4-13](#) are for example only.

For more information on configuring the appliances for load sharing, see [Load Sharing on page 215](#).

This chapter provides a few scenarios for the protocols to illustrate how you can integrate the appliance with your existing network when you want to use it in **Transparent Bridge** mode.

You can also share the scanning workload between appliances by setting up the appliances as described in [Load sharing on page 94](#).

NOTE

You must use the appliance **inside** your organization and **behind** a correctly configured firewall.

When using appliances in **Transparent Bridge** mode, there must only be one appliance in operation on each subnet.

If you are in any doubt about your network's topology and how you should integrate the appliance, consult your network expert.

SMTP scenarios

This section describes some SMTP configurations for when your appliance is in **Transparent Bridge** mode. This section describes:

- [One site on page 82](#).
- [Dedicated appliances on page 83](#).
- [Multiple sites on page 84](#).
- [Demilitarized zone \(DMZ\) on page 85](#).
- [International organization on page 86](#).

One site

The one site topology shown in [Figure 5-1](#) is very common for small companies with a few hundred users.




Figure 5-1. One site: Transparent Bridge

NOTE

You must ensure that users cannot directly send outbound messages to the firewall or appliance. This can be enforced with a firewall policy.

Configure the flow of e-mail as shown in [Figure 5-2](#).




Figure 5-2. One site: flow of e-mail messages

Figure 5-2 on page 82 shows a fail-safe configuration:

- If the appliance's network connection fails between the appliance and the mail server, the outbound messages remain at the mail server, and the inbound messages are stored at the appliance. They will not be delivered unscanned to your mail server. This configuration ensures that unscanned e-mail messages cannot enter or leave your mail system.
- If the appliance's network connection fails between the appliance and the firewall, outbound mail is held at the appliance and inbound mail is held at the external mail server that is attempting to deliver the mail to you.

Dedicated appliances

Figure 5-3 shows how you could set up a network with a dedicated appliance for each of the supported protocols. In reality, you would probably only need to set up a dedicated appliance if there were sufficient traffic for that supported protocol to warrant its own appliance.

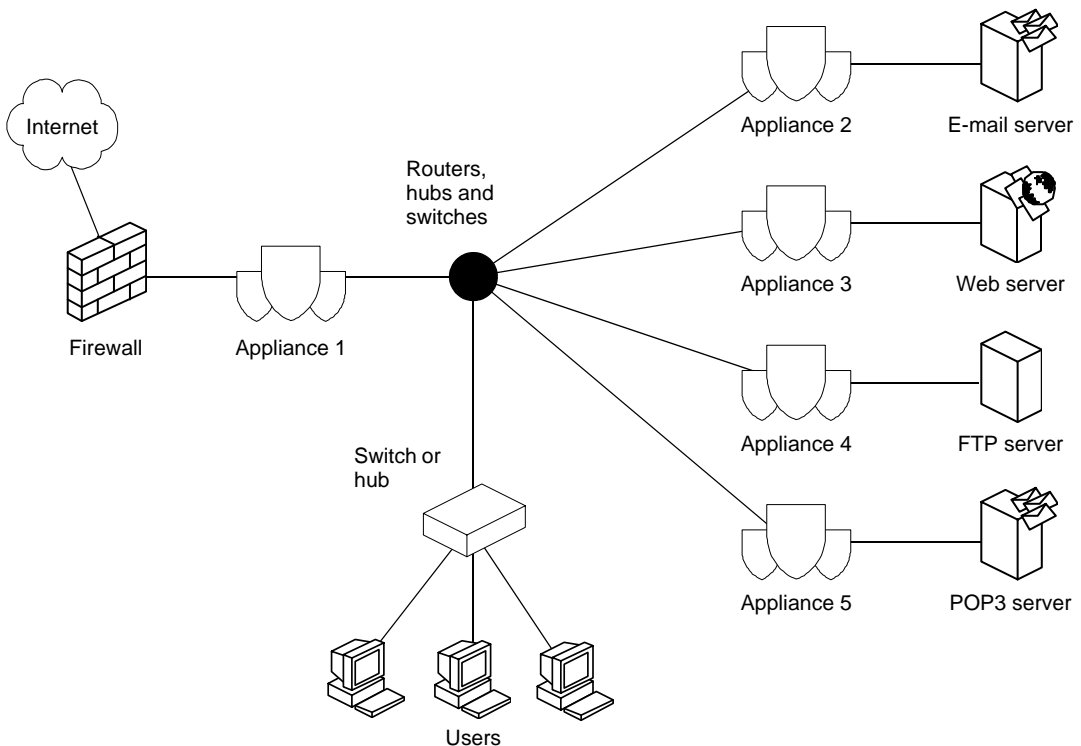


Figure 5-3. Dedicated appliances

For this configuration ([Figure 5-3 on page 83](#)):

- Appliance 1 has been set up to scan all support protocols.
- Appliance 2 has been set up to scan only SMTP traffic.
- Appliance 3 has been set up to scan only HTTP traffic.
- Appliance 4 has been set up to scans only FTP traffic.
- Appliance 5 has been set up to scan only POP3 traffic.

NOTE

Each appliance must be on a different subnet.

Multiple sites

[Figure 5-4](#) shows a typical configuration for a company that has multiple sites connected by a Wide Area Network (WAN) link. Each site has a mail server that handles the site's e-mail messages, but only the main site has a link to the Internet.




Figure 5-4. Multiple sites

NOTE

Traffic between Site A and Site B will not be scanned, because traffic will never be routed to the appliance.

If you want traffic between internal sites to be scanned, we recommend that you install another appliance between the sites.

Configure the flow of e-mail as shown in [Figure 5-5](#).

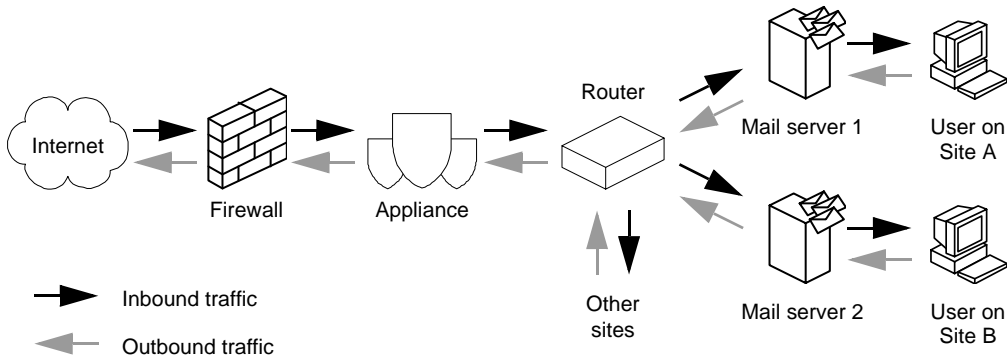


Figure 5-5. Multiple sites: flow of e-mail messages

This is a *fail-safe* configuration, ensuring that unscanned e-mail messages cannot enter or leave your mail system if the appliance's network connection fails.

Demilitarized zone (DMZ)

We do not recommend trying to set up a demilitarized zone security model using appliances in **Transparent Bridge** mode. You should only set up a DMZ security model using appliances in **Explicit Proxy** mode, as described in [Demilitarized zone \(DMZ\) on page 50](#).

International organization

Figure 5-6 shows sites in different countries connected by a Wide Area Network (WAN) link. There are multiple appliances and multiple links, adding fault tolerance to the mail setup, and ensuring that e-mail messages can still be scanned and delivered if one of the connections to an appliance or the Internet fails. Messages for any of the network domains can be received through any Internet link. The main mail server (called the *bridgehead* mail server) holds the main directory or alias list for the entire organization, and routes the e-mail messages to the correct internal mail servers.

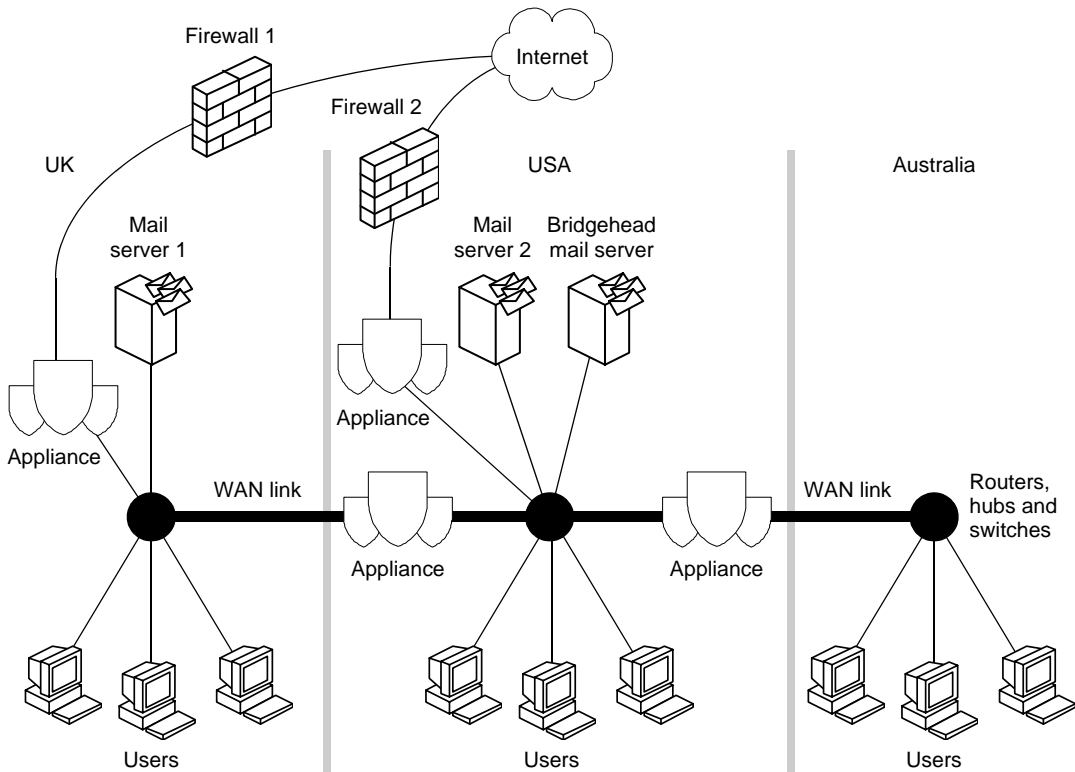


Figure 5-6. International organization

FTP scenarios

This section presents the following FTP scenarios:

- *Outbound FTP.*
- *Inbound FTP on page 88.*

NOTE

We recommend that you protect your incoming FTP information (PUT commands to your FTP server) with the appliance.

Outbound FTP

Figure 5-7 shows the communication that occurs when an internal user initiates an FTP connection with an external FTP server.




Figure 5-7. Outbound FTP

The internal user's FTP request is sent to the external FTP server. The appliance intercepts and scans the request. When the firewall receives this request, if it obeys its security policy, the request passes through the firewall to the Internet and external FTP server.

The FTP response is also intercepted and scanned by the appliance.

Inbound FTP

Figure 5-8 shows the communication that occurs when an external user accesses an internal FTP server.




Figure 5-8. Inbound FTP with one server

The external user's FTP request is sent to the FTP server. When the firewall receives this request, if it obeys its security policy, the request passes through the firewall. Before the request reaches the FTP server the appliance intercepts and scans the request.

The FTP response is also intercepted and scanned by the appliance.

HTTP scenarios

This section presents the following HTTP scenarios:

- *Outbound HTTP with internal web cache on page 89.*
- *Outbound HTTP with external web cache on page 90.*
- *Outbound HTTP without web cache on page 91.*
- *Inbound HTTP on page 92.*

NOTE

By default, the appliance only handles and scans HTTP traffic on port 80.

Outbound HTTP with internal web cache

Figure 5-9 shows the communication that occurs when an internal user accesses an external web server and there is a web cache between the user and the appliance.

The benefit of using this configuration is that, once cached, web pages do not need to be rescanned each time they are used. The amount of traffic passing through the appliance is reduced, which improves the throughput of other traffic being scanned.




Figure 5-9. Outbound HTTP with internal web cache

The user makes an HTTP request to the web cache. For example, the user requests a web page. If the page has already been cached, the web cache passes the page to the user. If the page has not already been cached, the web cache requests that page from the web server. The request is intercepted and scanned by the appliance. If the request meets the network's security policy it will also pass through the firewall to the Internet and external web server.

The response, in this case, the web page, is also intercepted and scanned by the appliance before it is stored in the web cache, and passed to the user.

The next time someone requests the same web page, the cached and clean copy of that page can be downloaded without having to pass through the appliance, and therefore without having to be rescanned.

Outbound HTTP with external web cache

Figure 5-10 shows the communication that occurs when an internal user accesses an external web cache. This configuration is common in organizations that use an Internet Service Provider (ISP) to provide web caching facilities.

Cached web pages will be scanned by the appliance, as they still pass through the appliance.

NOTE

If you have an external web cache configuration, we recommend that you install an internal web cache between the users and the appliance to reduce the need for scanning.




Figure 5-10. Outbound HTTP with external web cache

The user makes an HTTP request to the web cache. For example, the user requests a web page. The request is intercepted and scanned by the appliance.

If the request meets the network's security policy it will pass through the firewall to the web cache.

If the requested page has already been cached, the web cache will pass the page to the user. The page will be intercepted and scanned by the appliance before it passes to the user.

If the requested page has not already been cached, the web cache will pass the request to the web server.

The response, in this case, the web page, is passed to the web cache and stored before being passed to the user.

Before the page reaches the user, the firewall checks that the connection meets the network's security policy, and the appliance intercepts and scans the page.

If the external web cache is a *transparent* web cache you do not need to perform any other configuration. If the web cache is not transparent, you should set up the appliance to use the web cache as its *handoff host*, or set up the users' browsers to send requests to the web cache.

NOTE

Setting up an HTTP handoff host is described in [Handoff host on page 176](#).

See the documentation that accompanies your web browsers for details of how to send traffic to the web cache.

Outbound HTTP without web cache

Figure 5-11 shows the communication that occurs when an internal user accesses an external web server and web pages are not cached.




Figure 5-11. Outbound HTTP

The user's HTTP request is sent to the external web server, and is intercepted and scanned by the appliance.

The response is also intercepted and scanned by the appliance.

Inbound HTTP

Figure 5-12 shows the communication that occurs when an external user accesses an internal web server.




Figure 5-12. Inbound HTTP

The user sends an HTTP request to the internal web server. For example, the user requests a web page.

If the request meets the network's security policy it will pass through the firewall and be intercepted and scanned by the appliance before reaching the web server.

The response, in this case, the web page, is also intercepted and scanned by the appliance.

If you only have one web server that you want to be accessed by external users, set up the appliance to use the web server as its inbound *handoff host*. Setting up a handoff host provides added security. External users can only access the specific web server, they cannot get unauthorized access to other web servers on your network.

NOTE

Setting up an HTTP handoff host is described in [Handoff host on page 176](#).

If you have more than one web server that you want to be accessed by external users, you should not set up the web server as the handoff host for the appliance. Instead, configure your firewall to only allow HTTP connections to the appropriate web servers, and to ensure that traffic passes through the appliance.

POP3 scenarios

The appliance scans mail download (POP3) traffic without distinguishing its direction (inbound or outbound).

Figure 5-13 shows the communication that occurs when an internal user accesses an external POP3 server.




Figure 5-13. POP3 flow of traffic

The user makes a POP3 request to download mail from the external POP3 server. The request is intercepted and scanned by the appliance.

If the request meets the network's security policy it will pass through the firewall to the Internet and external POP3 server.

The response, in this case, mail messages, is also intercepted and scanned by the appliance.

Load sharing

If you have more than one appliance you can share the scanning workload between the appliances. If the controlling appliance is in **Transparent Bridge** mode, the appliances should be configured as shown in [Figure 5-14](#).

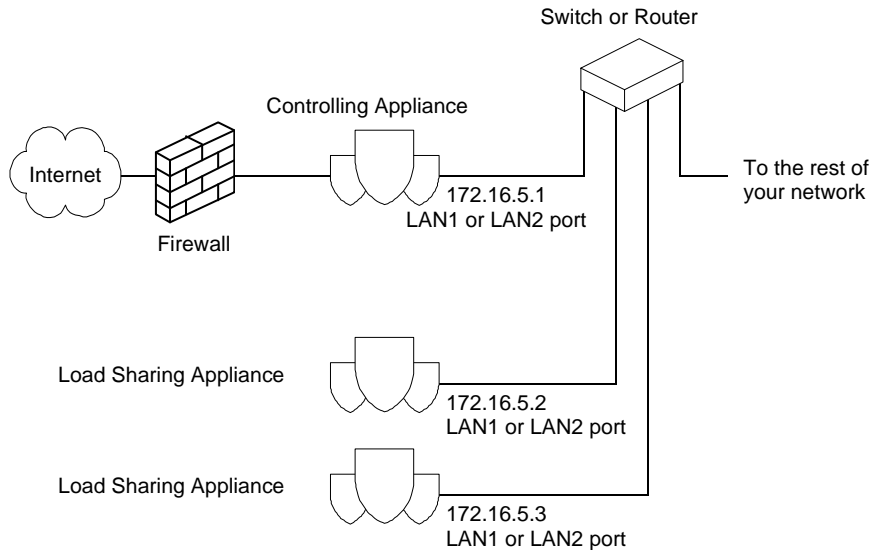


Figure 5-14. Load sharing Transparent Bridge configuration

NOTE

The IP addresses shown in [Figure 5-14](#) are for example only.

For more information on configuring the appliances for load sharing, see [Load Sharing](#) on page 215.

This section provides extra information that might help you perform the initial configuration of the appliance.

Accessing the appliance

When you have installed or upgraded the appliance, you can use the following methods to manage the appliance:

- The stand-alone WebShield client application
- A supported web browser

NOTE

We recommend that you use the WebShield client application to avoid issues with the use of the back button in certain web browsers.

User name and default password

When you log on for the first time the username is:

webshield

The default password is:

webshieldchangeme

For security reasons you should change the password.

Initial configuration

The initial configuration of the appliance is through the **Setup Wizard**.

Network settings can then be reconfigured using the **Network -> Settings** option, or by relaunching the **Setup Wizard** using **Network -> Setup Wizard**.

The following sections provide some more information about the initial configuration settings. For more detailed information about using the **Setup Wizard** or **Network -> Settings** options, see the *Product Guide*.

Appliance name

This is the name used to identify the appliance on the network.

NOTE

Appliance names must not contain spaces.

If you are using McAfee ePolicy Orchestrator, this is also the name used by ePolicy Orchestrator to identify the appliance.

We recommend that you change the default name to prevent the appliance being a target for hackers. The name must be unique and no longer than 15 characters.

If you are using other appliances in the same network, you must ensure that all the appliances have unique names.

NOTE

You must make sure that the appliance name is registered in the Domain Name System (DNS) servers. DNS servers are described in [DNS servers on page 101](#).

Domain name

On the Internet, computers and networks are generally grouped according to their organization type or geography. These network groupings are known as *domains*.

You must supply the domain or subdomain in which the appliance is located.

The domain or subdomain must be the full name of the system and not just the host name.

NOTE

The full name is often known as the *Fully Qualified Domain Name (FQDN)*. An example of an FQDN is *sample.nai.com*.

The domain name should be entered into the DNS servers that the appliance will use, so the appliance can be identified.

Default gateway

The default gateway is the next hop out of your network or subnet. Enter the IP address of the default gateway/router, for example: 111.111.111.1

Operational mode

The appliance can operate in one of three modes:

- **Explicit Proxy** mode.
- **Transparent Router** mode.
- **Transparent Bridge** mode.

Selecting the right operational mode for the appliance is an important choice as it impacts how you integrate the appliance into your existing network and how the appliance handles traffic. After you select the right mode for the appliance, you should not need to change its mode until you next restructure your network. Operational modes are described in more detail in [Which Operational Mode? on page 31](#).

Bridge Priority

If the appliance is operating in **Transparent Bridge** mode, and you are running the *Spanning Tree Protocol (STP)* on your network, you need to make sure that the appliance is given the right *bridge priority* according to STP rules.

By default the MAC address of the appliance is used to determine bridge priority.

If you do not want to use MAC addresses to determine priority, you must assign a unique number, in the range 0 through 65535, to each **Transparent Bridge** appliance. The lower the number the higher the bridge priority.

Protocols

The appliance can handle and scan traffic for the following protocols:

- Simple Mail Transfer Protocol (SMTP) e-mail messages.
- File Transfer Protocol (FTP) exchanges.
- Hypertext Transfer Protocol (HTTP) web browsing.
- Post Office Protocol version 3 (POP3) Internet e-mail messages.

You can enable or disable each of the supported protocols. By default the protocols are enabled, and traffic is scanned in both directions.

Protocol specific configuration is described in:

- [Managing SMTP e-mail on page 111](#).
- [Managing HTTP on page 169](#).
- [Managing FTP on page 183](#).
- [Managing POP3 on page 191](#).

Interface addresses

You must supply TCP/IP network address information for the appliance so that it can communicate with the network to which it is connected.

To prevent duplicate IP addresses on your network and to deter hackers, you should change the default IP addresses for the appliance to unique IP addresses. The IP address must be suitable for your network.

You can move IP addresses up and down the list.

An IP address at the top a list is known as a *primary IP address*, and any IP addresses below it in the list are known as *aliases*.

You cannot delete or disable a primary IP address (with the exception of an appliance in **Explicit Proxy** mode, where the primary IP address on the LAN2 interface can be disabled, but only as part of disabling the whole LAN2 interface).

You can use the **Advanced** button to configure the network adaptor settings.

For more information about interface addresses, refer to *Getting Started* in the *Product Guide*.

Inside and Outside Networks

The appliance has an *Inside Networks* list and an *Outside Networks* list that it uses to identify whether traffic passing through it has come from an internal or external network source.

NOTE

As traffic can be scanned according to its direction, it is important that you enter the correct information in your *Inside Networks* and *Outside Networks* lists.

Network devices that are not listed in either the *Inside Networks* list or the *Outside Networks* list will be refused access to the appliance.

What should be in the Inside Networks list?

The following information must be listed in the *Inside Networks* list:

- The domains or IP addresses for the internal networks (those inside your organization and behind a correctly configured firewall) with which the appliance communicates.
- Load sharing appliances should list their controlling appliances.

What should be in the Outside Networks list?

The appliance uses a * domain to identify its external networks as everything that has not been specified as an internal network. You can either accept this default or list individual networks and domains.

If you want to list individual networks and domains, you must include:

- The firewall.
- Any internal subdomains that you want to be treated as external networks, such as `testing.example.com`.

NOTE

There should be at least one entry in the *Outside Networks* list, and if you only have one entry then it must be your firewall.

To allow incoming connections from the Internet, we strongly recommend that you keep the * domain entry in your *Outside Networks* list. This is particularly required for inbound mail. Hosts that are not listed as internal or external will be blocked from using the appliance.

All other entries should be added above the * domain entry.

Adding domains and networks

The process of listing domain and network information is similar for both *Inside Networks* and *Outside Networks*.

To set up information that is used to identify internal domains and networks, use the **Inside Networks** tab.

To set up information that is used to identify external domains and networks, use the **Outside Networks** tab.

If you already have the domain and network information in a file with .CSV (comma-separated values) format, you can import that information.

Each entry in the file must be on a single line using the following format:

Type	Format	Example
Domain	D, <domain>	D, www.nai.com
Network Address	N, <IP address>, <IP subnet mask>	N, 192.168.1.99, 255.255.255.0

Alternatively, you can manually specify the domain and network information. You can enter the domain name or the IP address. You can use domain names because the appliance can perform reverse DNS server lookups to check the hosts.

You can identify as many domains and networks as you want.

NOTE

If you are configuring the **Inside Networks** for the first time, the inside networks information is used as the *local domains*. The appliance uses the local domains information to prevent the unwanted relaying of e-mail messages through your organization. See [Anti-relay on page 113](#).

Resolving conflicts

If a network device is effectively listed in both the *Inside Networks* and *Outside Networks* lists, the appliance must decide whether to treat that network device as part of the internal network or part of the external network.

The appliance compares the entries for the device and applies the following rules, in order, until it can determine how to treat the device:

- A more precise domain name takes precedence over a less well-defined domain name. For example:

host.sales.example.com takes precedence over **.sales.example.com*

sales.example.com takes precedence over **.example.com*

- A longer subnet mask will take precedence over a shorter subnet mask. For example:

255.255.255.0 takes precedence over 255.255.0.0

255.255.0.0 takes precedence over 255.0.0.0

- If the entries are identical for both lists, the network device will be considered part of the internal network.

NOTE

Although you can use a mixture of both domain names and IP addresses in the *Inside Networks* and *Outside Networks* lists, we recommend that you use either domain names **or** IP addresses. Using the same method throughout will help you identify potential conflicts that could otherwise lead to direction-related virus-scanning problems.

DNS servers

Domain Name System (DNS) servers translate the user-friendly names of network devices and domain names, such as *www.aol.com*, into IP addresses.

As DNS is a distributed service, you can specify a list of name servers that will be used to resolve the host name to IP address mapping and IP address to host name mapping.

The servers will be contacted in top-down order. For example, if the request cannot be resolved by the first name server in the list, the second name server in the list will be contacted.

If the request cannot be resolved by the name servers in the list, the request will be forwarded to the DNS root name servers on the Internet. You can prevent requests being sent to the root name servers by selecting the

Use forward only mode option.

The appliance uses the local DNS server for the following activities:

- Attempting to deliver scanned SMTP e-mail messages, if the DNS delivery method is selected.
- Verifying web browsing (HTTP) requests and determining which URLs to block, if URL blocking is configured.

Static routes

If you are using the appliance in **Explicit Proxy** mode, there are circumstances where it needs to use more than one gateway or router to communicate with the entire network. In these circumstances, specify the default gateway or router in the **Initial network settings** page of the wizard (in **Default gateway**), and specify the other gateways or routers under static routes.

If you need static routes, specify the routing information that will be stored in the appliance's routing table. Every network is different. If you do not know the routing information for your network, consult your network expert.

Dynamic routes

Dynamic routing is a system that allows your network equipment, including the appliance, to listen for the routing information that is broadcast by the routers on your network. The equipment can then use that information to configure its own routing information.

NOTE

The appliance only supports the RIP (Routing Information Protocol) and OSPF (Open Shortest Path First) routing protocols.

Every network is different. If you do not know whether your network uses dynamic routing, consult your network expert.

If necessary, enable dynamic routing for the appliance.

Routing information is stored in the *Routing Table*. You can use the **Troubleshoot -> Diagnostics -> Display Routing Information** option to view the appliance's routing table, to ensure that it is receiving routing information.

NOTE

If you disable dynamic routing, any dynamic routing information already acquired by the appliance can only be removed by rebooting the appliance.

Load sharing

It is possible to share the task of scanning traffic between several appliances on the same network.

Appliances can give some of their scanning work-load to other appliances.

They can also accept or refuse requests to scan traffic on behalf of other appliances. For more information about load sharing, see [Load Sharing on page 215](#).

Password settings

When you first access the appliance you should change the password used to access the appliance's web interface.

NOTE

The default password for all appliances is now *webshieldchangeme*.

Date and time settings

By default the appliance is set to use UTC (Temps universel coordonné) or Coordinated Universal Time in English.

If you want the appliance to use a different time zone, select it from the list of time zones.

Operational language

Select the language in which you want the appliance to operate. This affects the reports and messages generated by the appliance, excluding any text strings that you can specify or modify through the interface. This is separate from the language used for viewing the interface (chosen when logging on to the appliance).

This section provides a general overview of policies and how they can be used. For more detailed information about how to set up policies, see the *Product Guide*.

What is a policy?

A policy is a collection of settings and content rules that allow you to combat a specific threat to your network. You can tell the appliance what to do when a threat becomes a reality. See [Policy actions](#) for more information about the type of actions the appliance can perform. You can:

- Apply ready-made policies to your entire organization. These policies are known as *global policies*. See [Global policies on page 104](#).
- Create new policies that can be tailored to suit the specific needs of any part of your organization. These policies are based on the global policies, and are called *non-global policies* in this user documentation. See [Non-global policies on page 104](#).
- Set up groups to which you can assign policies. These groups, which could represent the different departments or functions within your organization, are known as *Policy Groups*. Policy groups are described in more detail in [Policy Groups on page 108](#).
- Create the content rules, that define what content will trigger a response from the appliance. Content rules can be created independently of a policy and applied to policies at a later date. Content rules are described in [Content rules on page 108](#).

Policy actions

A policy specifies how the appliance must act when a threat becomes a reality.

You can specify the action that the appliance takes when any part of the policy is violated. For example, if a virus is detected, you can choose to clean, quarantine, or delete the infected item.

If the appliance finds an undesirable phrase that you specified in a content rule, you can choose to block the item or allow it through.

If a large file is detected, you can choose to block it, or you can allow it through and issue an alert.

You can also issue alerts and notifications.

Issuing alerts and notifications

Each item in the policy has an action associated with it. When a rule or setting is triggered, both the user and an administrator can be informed.

The appliance will replace the offending message or attachment with text that you prepare. Any users who later read the message will see the replacement text instead. You can also request the appliance to send a message to an administrator, and record the event in a log.

The appliance uses substitution variables (also known as *tokens*) to customize alert messages. For example, a message of the form:

```
Bad content detected at %LOCALTIME%.
```

might become:

```
Bad content detected at 12:34 on Monday.
```

To find out which substitution variables can be used and where, see the *Substitution Variables* appendix in the *Product Guide*.

Global policies

When it is first installed, the appliance has one or more global policies, which describe how items will be scanned for viruses, file-filtering rules, and various other settings. Initially, such policies apply to the whole organization. From a global policy, you can create further policies as necessary to apply to groups of users or domains.

As you create further policies, each one records whether any of its current settings are inherited from its global policy. A change to the global policy — such as an increased level of anti-virus protection or a new file-filtering rule — is propagated instantly to the other policies. The global policy also includes an indication of how many other policies have inherited its settings.

Non-global policies

A non-global policy allows you to create exceptions to the general rules encapsulated in a global policy.

For example, as a general rule, you might not want anyone in your organization to be able to send or receive attachments in e-mail, with the exception of the marketing department.

In this case the global policy would be set up to disallow attachments and you would create a non-global policy that allows the marketing group to receive and send e-mail messages with attachments.

Inheriting global settings

Non-global policies can *inherit* settings from the global policy.

In the following example, the policy on the left is called a global policy. It is well suited to most departments in the organization. However, it is not ideal for the sales department because they often handle much larger files for customers. Therefore the sales department needs a different policy. You can create their policy by creating a new non-global policy based on the global policy, then modifying some parts of the non-global policy to better suit that department.

Table 7-1. Creating policies

Global policy	Non-global policy for the Sales Department
Apply medium-level scanning for viruses.	Apply medium-level scanning for viruses.
Do not accept files that are larger than 10MB.	Do not accept files that are larger than 50MB.

The sales department has *inherited* one item from the global policy (for medium-level scanning) and modified one item (for the size of files).

For each non-global setting, you can specify if the actions should be inherited or modified (*disinherited*).

Ordering non-global policies

In some cases an item, such as a file, e-mail message, or user's document, might be covered by two or more policies.

For example, a person might work for two departments, where each department has its own policy. One department might be allowed to access a certain web site, while the other department is banned from accessing that web site.

When the user who is in both policy groups tries to access the web site the appliance must know which of the two conflicting policies to apply to that user.

The appliance will only apply the topmost non-global policy. For this reason, it is important to make sure that non-global policies are correctly ordered.

Adding time-specific settings to non-global policies

The settings within global policies are not time-specific. For example, you can either enable or disable the **Anti-Virus** setting. You cannot configure the global **Anti-Virus** setting to only scan e-mail messages between the hours of 2am and 4am on a Friday morning.

You can apply time restriction settings to non-global content policies.

For example, you might want to specify that users can only send e-mail messages with large attachments at certain times of the day; typically when the network is less congested.

To do this you would set up a global policy that prevents the sending of e-mail messages that have large attachments. You would then set up a non-global policy that allows the sending of large attachments at a specific time each day.

The appliance will always give precedence to the more precise (granular) policy. In this example, specifying that the non-global policy should only apply at a set time is more granular than the global policy that specifies that the policy applies at “any time” of the day. At the set time, the time-specific non-global policy will override the less-specific global policy.

Multiple instances of rules and settings

You can have multiple instances of the same rule or some other settings within one policy. For example, a sales department has the following schedule:

Time	Activity
Monday morning	Prepare a list of sales targets for the week.
Monday afternoon	Publish the report of sales targets.
Tuesday	Sell!
Wednesday	Sell!
Thursday	Sell!
Friday morning	Sell!
Friday afternoon	Publish the report of sales achieved.

The two sales reports are sent or published on Monday afternoon and Friday afternoon only. The reports must not be circulated or modified at any other times. To control this, you can create two rules:

- Ban any document that contains the phrase “Sales Report” on Monday morning.
- Ban any document that contains the phrase “Sales Report” between Tuesday morning and Friday morning.

The rules are identical but their times are different.

Handling overlapping time restrictions

If you have specified time restrictions that overlap each other, the appliance handles these in the following way:

- Only one instance of a specific rule can be active at any one time.
- Only one instance of the anti-virus settings is active at any one time.
- If a rule or anti-virus setting is set to operate all the time, it has the lowest priority; it becomes active only when no time-restricted rule or anti-virus settings is in operation.
- If a rule or anti-virus setting is set to operate at various times, the most recently activated instance is active. For example, you have these settings:
 - 1 Apply a low-level of anti-virus scanning all the time.
 - 2 Apply a medium-level of anti-virus scanning on weekdays only.
 - 3 Apply a high-level of anti-virus scanning Thursday mornings only.

The first set of anti-virus settings will apply, unless it is now a weekday. This setting is ignored if now is a Thursday morning.

In general, avoid creating complex time restrictions. If a rule or anti-virus setting has several time restrictions applied to it, the restriction that is in force has these qualities:

- Latest start time.
- Earliest finish time.
- Latest first day (using Sunday as first day of the week).
- Earliest last day (using Sunday as first day of the week).
- Least number of days.

Specifying time restrictions

You can specify the time restrictions precisely, even to the minute. You can give a name to each period such as **Morning**, then define it, for example, as 9:05 to 11:59.

Policy Groups

You might want to apply different policies to different groups within your organization. Before you can apply a policy to these different groups, you must first define the groups. The groups to which policies are applied are called *Policy groups*.

To create a policy group you need to:

- Give the policy group a name.
- Define the membership of the policy group.
- Specify which conditions must be met before the appliance can apply a policy to this policy group.

NOTE

If you want to use information from your Lightweight Directory Access Protocol (LDAP) servers to create e-mail groups, you must have imported the directory information from those servers before you create a policy group. See the *Product Guide* for more information about using LDAP servers.

Content rules

A content rule is used to define unacceptable content. For example, you could set up a content rule to ban the use of a particular swear word in e-mail messages entering or leaving your organization.

That content rule can then be applied to any content policy, and the appliance configured to perform certain actions when that content rule is triggered. For example, when the appliance detects the swear word in an e-mail message it could refuse the e-mail message and send a warning to a network administrator that an offensive e-mail message has been detected.

As you could create a large number of content rules, content rules are organized into *Rule Groups*. Each rule group has one or more *Content Rules*.

The appliance comes pre-configured with a standard set of rule groups. You can add content rules to these rule groups or create new rule groups of your own. You create the rule groups first and then add content rules to them.

You can assign the whole rule group to a policy or just assign selected content rules. Assigning selected content rules allows you to set up policy-specific settings for those content rules.

For more information about content rules, see the *Content Rules* section of the *Product Guide*.

Before you begin

This section describes some of the issues you should consider before creating policies.

Spend some time planning

Policies are very powerful tools, and if you want to make full use them, you need to spend some time planning. Poorly configured policies can cause serious security and connectivity issues for your network.

You should:

- Familiarize yourself with the concepts described in this section.
- Spend some time thinking about how to organize users and computers into different policy groups.

NOTE

This is particularly important if you are setting up SMTP e-mail policies. The number of non-global policies you create can affect the number of scans the appliance has to perform. This in turn can affect the appliance's performance. See [Multi-policies for e-mail messages on page 138](#).

- Consider which policies you want to assign to those policy groups. The policies that you can assign are described in:
 - ◆ [Policy-based e-mail configuration on page 121](#).
 - ◆ [Policy-based HTTP configuration on page 170](#).
 - ◆ [Policy-based FTP configuration on page 184](#).
 - ◆ [Policy-based POP3 configuration on page 193](#).
- Consider the legal implications of setting certain policies. See [Considering legal implications on page 110](#).
- Follow the general guidelines. See [General guidelines on page 110](#).

General guidelines

There are some general guidelines that you should follow when setting up policies:

- Set up the global policies to cover most scenarios.

See [Global policies on page 104](#) for more information about global policies.

- Only set up a non-global policy if you need to create exceptions to the way that the global policy handles the item. For example, if you want to create exceptions to the way that connections or traffic is normally handled by the appliance.

See [Non-global policies on page 104](#) for more information about non-global policies.

- When the appliance is in Transparent Router or Transparent Bridge mode, the priority assigned to non-global content policies is important.

See [Ordering non-global policies on page 105](#) for more information about prioritizing non-global policies.

Caution

Incorrect configuration of advanced policy settings can cause serious security and connectivity issues for your network. For this reason, we recommend that you do not change advanced settings unless instructed to do so by Technical Support or your network consultant.

Considering legal implications

Before applying any restrictions on employees' e-mail and Internet access, check any requirements in your local laws. In some instances, such restrictions might be illegal. You should at least consider informing employees that restrictions are in force. It might be useful to display a statement when they start up their computers, or you can attach a disclaimer to each e-mail message. We advise you to discuss the implications with your legal department.

This section provides more information about SMTP e-mail configuration. It includes the following topics:

- [General e-mail \(SMTP\) configuration.](#)
- [Policy-based e-mail configuration on page 121.](#)
- [How e-mail messages are processed on page 137.](#)

General e-mail (SMTP) configuration

The features listed under the **Configure** -> **SMTP** menu option are not based on specific policies.

This section describes these non-policy based options, and it includes the following topics:

- [E-mail \(SMTP\) delivery.](#)
- [Anti-relay on page 113.](#)
- [Permit and deny settings on page 116.](#)
- [Connection settings \(Advanced\) on page 120.](#)
- [Retryer settings on page 121.](#)

E-mail (SMTP) delivery

The appliance attempts one or more of the following methods in the order shown, when delivering scanned SMTP e-mail message:

- Domain relays
- DNS
- Fallback relays.

Domain relays

You can specify mail relays to route e-mail messages destined for specific domains to specific mail servers ([Figure 8-1 on page 112](#)). You can create as many relays as you want.

NOTE

Identify the most common relays first because the appliance tries the relays in order.

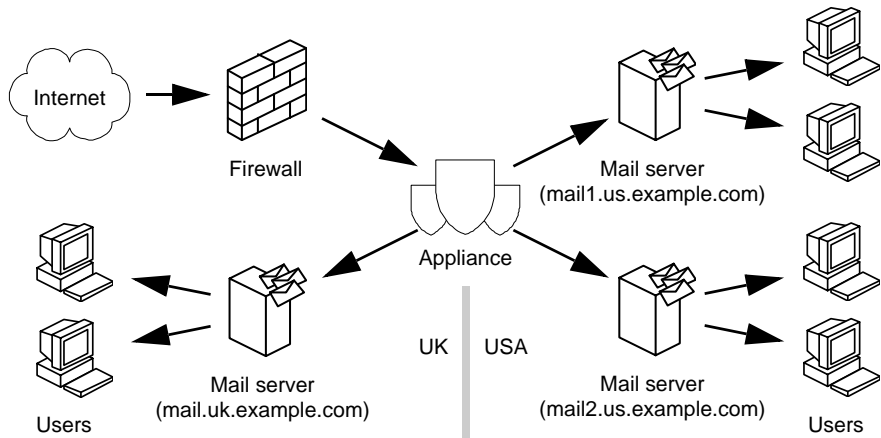


Figure 8-1. Mail relays to multiple mail servers or gateways

DNS servers

If you enable DNS, the appliance uses the DNS servers that you specify when supplying the TCP/IP network address information; see [DNS servers on page 101](#).

Fallback relays

You can specify fallback relays to route e-mail messages that cannot be delivered by the other delivery methods. You can create as many relays as you want.

NOTE

Identify the most common relays first because the appliance tries the relays in order.

Creating a postmaster

The appliance can assign an administrator (or *postmaster*) to handle any queries from senders about e-mail messages that were returned because they triggered one or more of the appliance's scanners.

Choose a user who will read e-mail messages regularly. Examples:

User@example.com

mail_master@example.net

Besides individual names, you can use distribution lists, such as IT_Managers@example.com. In e-mail products such as Microsoft Outlook, your recipients can also use rules to forward their e-mail messages in case they themselves are not available. In this way, you can be sure that any user problems can be resolved quickly.

Anti-relay

The appliance can prevent third-party relaying.

Mail relaying is often used for malicious purposes, such as *mail bombing* or *spamming*. The anti-relay feature prevents unscrupulous third parties using the appliance or the mail servers that it protects to deliver mail for them.

Consider the possible consequences of your business clients receiving relayed messages, maybe of a distasteful nature, that appear to be sent from your company. Using the anti-relay feature, you can avoid such embarrassment, and protect your organization's professional image.

The appliance's anti-relay feature prevents the forwarding of e-mail (SMTP) messages that do not originate in one of the domains that the appliance has been configured to accept. Sites that forward all messages are commonly called *open relays* and are currently considered undesirable. The operators of such sites might become banned by organizations such as MAPS which attempt to control the propagation of spam mail.

NOTE

Messages with special routing characters such as "%" can also be permitted access or denied access. See [Anti-Relay \(routing characters\) on page 131](#) for more information about routing characters.

Local domains

Only domains and networks specified in the **Configure -> SMTP -> Anti-Relay -> Local Domains** list can use the appliance to relay e-mail messages. When the appliance receives an e-mail message with an address that matches a local domain name, it handles the e-mail message as described in [How the appliance uses the anti-relay lists on page 115](#).

There are three ways to set up entries for your local domains:

- If you have already specified your *Inside Networks*, using the initial **System Configuration** page, that information will be used as your local domains. See [Inside and Outside Networks on page 98](#) for more details.
- If you have a .CSV (comma-separated values) file that lists your Inside Networks you can import that file. Importing a file overwrites the existing entries.
- Alternatively, you can manually specify the local domains and networks for your organization. You can enter the domain name or the IP address.

You must define at least one local domain.

If you use a .CSV file, each entry in the file must be on a single line and have the following format:

Type	Format	Example
Domain	D, <domain>	D, www.nai.com
Network Address	N, <IP address>, <IP subnet mask>	N, 192.168.1.99, 255.255.255.0

Deny domains

Domains and networks specified in the **Configure -> SMTP -> Anti-Relay -> Deny Domains** list cannot use the appliance to relay e-mail messages. When the appliance receives an e-mail message with an address that matches a deny domain name, it handles the e-mail message as described in [How the appliance uses the anti-relay lists on page 115](#).

Permit domains

Domains and networks specified in the **Configure -> SMTP -> Anti-Relay -> Permit Domains** list can use the appliance to relay e-mail messages. You can use this option to permit sub-domains when the parent domain is listed in the **Deny domains** list, as entries in the **Permit domains** override those in the **Deny domains** list.

When the appliance receives an e-mail message with an address that matches an entry in the **Permit domains** list, it handles the e-mail message as described in [How the appliance uses the anti-relay lists on page 115](#).

Anti-relay response

You can choose how the appliance responds when refusing to relay an e-mail message. The appliance can:

- **Reject the recipient** — the appliance can refuse the e-mail message.

The appliance sends a rejection code (SMTP 550 Fail).

We recommend this option, because the sender is normally informed that the message was not relayed.

- **Accept and ignore the recipient** — The appliance accepts the e-mail message, but does not deliver the e-mail message to that recipient.

The appliance sends an acceptance code (SMTP 250 OK).

We do not recommend this option, because it suggests to the sender that the message was received as intended.

How the appliance uses the anti-relay lists

If the **Local Domains** list is left blank, the appliance is an open relay, accepting and passing on all e-mail messages it receives. This occurs regardless of whether there are entries in the **Deny Domains** or **Permit Domains** lists.

NOTE

We recommend that you do not leave the appliance as an open relay, because it could be exploited by senders of spam to pass on their spam messages.

If the **Local Domains** list is not blank, the appliance checks the e-mail messages that pass through it. For each e-mail message, the appliance checks the e-mail message source IP address and recipients against the entries in the **Local Domains**, **Deny Domains** and **Permit Domains** lists:

Scenario	Local Domains	Deny Domains	Permit Domains	Outcome
1	?	?	Yes	Allowed
2	?	Yes	No	Rejected
3	Yes	No	No	Allowed

Scenario 1 — if the message matches an entry in the **Permit Domains** list, it is allowed through, regardless of whether it matches entries in the other lists.

Scenario 2 — if the messages does not match an entry in the **Permit Domains** list but does match an entry in the **Deny Domains** list, it is rejected, regardless of whether it matches entries in the **Local Domains** list.

Scenario 3 — if the message only matches an entry in the **Local Domains** list, it is allowed through.

In all three lists, you can specify a number of domains and IP address ranges. When checking each message to determine a match, the appliance interprets the list entries like this:

- **Domain entry** — the appliance checks the message's destination e-mail address (the recipient) against the domain to see if there is a match. If the Domain entry has A records on the DNS server, this address will also be checked against the message's source IP address (the sending server).
- **Domain range entry** — the appliance checks the message's destination e-mail address (the recipient) against the domain range to see if there is a match.
- **IP address or IP address range entry** — the appliance checks the message's source IP address (the sending server) against the IP address or IP address range to see if there is a match.

In the special circumstances that the appliance receives an e-mail message addressed to a specific IP address — such as user@[192.164.4.170] — it interprets the list entries as follows:

- **Domain entry** — the appliance accesses the A records on the DNS server to retrieve the domain's corresponding IP address. It then checks both the message's source IP address (the sending server) and the destination e-mail address (the recipient) against the IP address to see if there is a match.

This activity is not performed for wildcard domain ranges because the IP addresses to which they refer cannot be determined.

- **IP address or IP address range entry** — the appliance checks both the message's source IP address (the sending server) and destination e-mail address (the recipient) against the IP address or IP address range to see if there is a match.

Permit and deny settings

You can:

- Use the **Permit sender** option to specify which e-mail messages can bypass the anti-spam checks.
- Use the **Deny sender** option to block e-mail messages from unwanted sources.
- Use **RBL servers** to specify which blackhole list servers will be used during the anti-spam checks.
- Use the **Response** option to specify how the appliance will respond to e-mail messages blocked by the **RBL Servers** check or **Deny Sender** check

Permit Sender

The **Permit Sender** option allows you to specifically permit e-mail access from an individual or organization that could be sending unwanted e-mail messages.

For example, you might want to receive information from other organizations in your business sector in order to keep up-to-date with their products and promotions. This type of product and promotional material would probably be identified as spam if the e-mail message containing it is scanned for spam. To avoid this situation, you need to add the sender of the promotional material to the **Permit Sender** list, so that e-mail messages from this sender bypass the spam-scanning process.

There are two ways to enter the details of a sender you want to deny access to:

- Import a .CSV (comma-separated values) file.
- Manually enter details.

Manually entering sender details

You can enter the names of senders in any of the following formats:

Type	Format
Domain	example.com, example*.com
E-mail Address	john_smith@example.com, john*@example.com
Network Address	10.1.1.0:255.255.255.0

Importing a .CSV file

If you have a .CSV (comma-separated values) file that lists senders you want to deny, you can import details from that file. Each entry must be on a single line and have the following format:

Type	Format	Example
Domain	D, <domain>	D, www.example.com
E-mail Address	E, <e-mail address>	E, user@example.com
Network Address	N, <IP address>, <IP subnet mask>	N, 192.168.1.99, 255.255.255.0

NOTE

Importing a .CSV file overwrites any existing entries.

Deny Sender

The **Deny Sender** option allows you to specifically deny e-mail access to an individual or organization that is sending unwanted e-mail messages.

For example, you might want to block e-mail messages from a specific recruitment company, or abusive ex-employee.

There are two ways to enter the details of a sender you want to deny access to:

- Import a .CSV (comma-separated values) file — the format required is the same as that described in [Importing a .CSV file on page 117](#).
- Manually enter details — the format required is the same as that described in [Manually entering sender details on page 117](#).

NOTE

The **Permit Sender** option overrides the **Deny Sender** option.

RBL servers

The appliance can block unwanted e-mail messages (often spam) from particular sources.

The appliance does this by comparing the IP address of an e-mail source against lists of potential sources of spam. These lists are known as blackhole lists and are maintained by such organizations as:

- <http://www.ordb.org>
- <http://www.mail-abuse.com/>
- <http://www.dsbl.org>
- <http://www.spamhaus.org>

NOTE

There may be other organizations not listed that provide similar services. While we take every effort to ensure the web site addresses are correct at the time of publication, they could change over time. We do not accept any responsibility for the availability or these services, or the accuracy of results obtained when using them.

You can specify which anti-spam checklists are used by entering the host names of the servers that maintain those lists.

Different lists can offer different services. For example, they might list known spammers, sites that allow open-relaying of e-mail messages, or lists of dial-up systems.

You can enter the names of the lists that you want to use. For example, you could enter `relays.ordb.org` or `blackholes.mail-abuse.org`. You can enter more than one list, but only one list at a time.

Responding to unwanted e-mail messages

You can choose how the appliance responds when it receives unwanted e-mail messages. The appliance can:

- **Reject the e-mail** — the appliance can deny the e-mail message and keep the connection open.

The appliance sends a rejection code (SMTP 550 Fail), and any further communication follows the RFC standard.

The sender is normally informed that the message was not accepted. The appliance sends a rejection code (SMTP 550 Fail). Normally, the sender is informed that the message was not relayed.

- **Reject the e-mail and close the connection** — the appliance can deny the e-mail message and keep the connection open.

Before closing the connection, the appliance sends a rejection code (SMTP 550 Fail).

We recommend this option, because it suggests to the user that their spam e-mail message did not reach the intended recipient.

- **Accept and drop the e-mail** — the appliance accepts the e-mail message and discards it.

The appliance issues an acceptance code (SMTP 250 OK).

We do not recommend this option, because it suggests to the sender that the message was received as intended.

NOTE

The response only applies to e-mail messages blocked by the **RBL Servers** check or **Deny Sender** check.

Connection settings (Advanced)

For TCP and UDP protocols, ports numbers are used to identify the ends of logical connections which carry specific services. Each service has an associated port number.

The default port number for SMTP is port 25.

NOTE

We recommend that you do not change the port numbers unless you understand port assignments and the implications of changing the port number.

Intercept ports

The **Intercept ports** list only applies to appliances operating in Transparent Router mode or Transparent Bridge mode. You can specify the ports on which the appliance will intercept SMTP e-mail traffic.

Listen ports

The **Listen ports** list only applies to appliances operating in Explicit Proxy mode. You can specify the ports on which the appliance will listen for SMTP traffic.

Listeners, connections and memory

You can set up scanning resources on a protocol-by-protocol basis.

For each protocol, you need to set up the number of processes listening for the protocol-specific traffic. These processes are known as *Listeners*. Each listener can handle a number of connections. The number of scans that can be performed simultaneously for each protocol is equal to the number of listeners multiplied by the number of connections.

As there is a finite amount of resource available for scanning, there will always be a resource trade-off between:

- The number of listeners for that protocol.
- The number of connections handled by the listeners for the protocol.
- The amount of memory required for scanning that protocol.
- The scanning resources assigned to other protocols.

When you first use the appliance, some default values are in place. You can restore these settings at any time if the modified values are unsuitable.

Retryer settings

The appliance can store e-mail messages and attempt to deliver them at a later time. You can specify:

- How long the appliance will wait between attempts to forward an e-mail message.
- How long the appliance should try to forward an e-mail message before that e-mail message is bounced.
- The maximum number of retryers that can be attempting to forward e-mail messages concurrently.

Policy-based e-mail configuration

You can use policies to determine how the appliance handles e-mail messages in certain circumstances. For example, how the appliance handles an e-mail message could be determined by:

- Whether the e-mail message is from hosts in your Inside Networks Lists or Outside Networks list.
- Who sent the e-mail message.
- Who is going to receive the e-mail message.
- The content of the e-mail message.

Before you begin

This section describes some of the issues you should consider before setting up policies for SMTP e-mail messages.

Before setting up policies, you should:

- Familiarize yourself with the concepts described in this section.
- Spend some time thinking about how to organize users and computers into different policy groups. See [Policy Groups on page 108](#) for more information about policy groups.
- Consider which policies you want to assign to those policy groups.

There are some general guidelines that you should follow when setting up policies for e-mail messages:

- Set up the global policies to cover most scenarios. See [Global policies on page 104](#) for more information about global policies.
- Only set up a non-global policy if you need to create exceptions to the way that most e-mail messages are handled, as specified by the global policy.

See [Non-global policies on page 104](#) for more information about non-global policies.

- When the appliance is in Transparent Router or Transparent Bridge mode, the priority assigned to non-global content policies is important. See [Ordering non-global policies on page 105](#) for more information about prioritizing non-global policies.

Caution

Incorrect configuration of advanced policy settings can cause serious security and connectivity issues for your network. For this reason, we recommend that you do not change advanced settings unless instructed to do so by Technical Support or your network expert.

Content Policies

The **Policy -> SMTP -> Content -> From Outside** menu options can be used to set up policies that tell the appliance how to handle e-mail messages from hosts in your **Outside Networks** list.

The **Policy -> SMTP -> Content -> From Inside** menu options can be used to set up policies that tell the appliance how to handle e-mail messages from hosts in your **Inside Networks** list.

You can specify the actions that the appliance will perform in certain circumstances. See [When a scanner triggers on page 141](#) for a list of possible actions.

You can configure the following SMTP content policies:

- [Alert settings on page 123.](#)
- [Anti-spam settings on page 124.](#)
- [Anti-virus settings on page 124.](#)
- [Content scanner on page 124.](#)
- [Corrupt content on page 124.](#)
- [Disclaimer Text on page 124.](#)
- [Encrypted content on page 125.](#)
- [File filtering on page 125.](#)
- [HTML settings on page 126.](#)
- [Mail settings on page 127.](#)
- [Mail size filtering on page 127.](#)
- [Protected content on page 128.](#)
- [Scanner control \(denial-of-service attacks\) on page 128.](#)
- [Signed content \(digital signatures\) on page 129.](#)

Alert settings

The appliance will send an HTML message to clients when a specific event occurs. This is known as an HTML alert.

You can:

- Change the text that appears at the start of the HTML alert, known as the *alert header*.
- Change the text that appears at the end of the HTML alert, known as the *alert footer*.

Anti-spam settings

The optional *McAfee SpamKiller for WebShield appliances* can help you control the amount of spam reaching end users.

For more information, see:

- [Evaluating SpamKiller on page 224.](#)
- [Activating SpamKiller on page 224.](#)
- [Scanning for Spam on page 147.](#)

Anti-virus settings

The appliance can be configured to detect viruses and other potentially harmful software. For more information about scanning for viruses, see [Virus-scanning on page 163.](#)

Content scanner

The appliance can be configured to detect unwanted content in e-mail messages. See [E-mail \(SMTP\) content scanning on page 153](#), for a detailed description of content scanning.

Corrupt content

If content is corrupt, the appliance might not be able to scan the file for viruses or banned content.

You can:

- Specify the action the appliance should take when corrupt content is detected.
- Specify which alert should be used, and if necessary, customize the alert text.

Disclaimer Text

A disclaimer is some text — an explanation, information, a legal statement, or warning — that the appliance can add to an e-mail message. The appliance enables you to add disclaimers to inbound and outbound e-mail messages, and to e-mail messages for specific groups of users.

By adding a disclaimer to outbound messages, you can limit the liability posed by statements that might be legally damaging, for example, those containing offensive remarks. Disclaimers are also useful for renouncing the contents of a message as the view of the author, not of the organization, to avoid any damaging publicity. For example:

The information contained in this message is confidential and may be legally privileged. Views or opinions expressed in this e-mail message are those of the author only.

By adding a disclaimer to inbound messages, you can make staff aware that all e-mail messages and attachments are being scanned for viruses and content. For example:

This e-mail message and any attachments were scanned for viruses when they left the organization. Communications will be monitored regularly to improve our service and for security and regulatory purposes. Thank you for your assistance.

By adding a disclaimer according to a specific department, you can protect your organization against costly misunderstandings. For example:

Prices quoted by our sales personnel in this e-mail message provide only a rough guide and do not represent any part of a formal contract.

You can:

- Enable or disable the use of disclaimers.
- Specify where the disclaimer should be inserted.
- Edit the disclaimer text.

Encrypted content

If content is encrypted, it cannot be scanned. You can:

- Specify how encrypted content is handled by the appliance.
- Specify which alert message to use, and if necessary, customize the alert message text.

If you allow it through, it must be scanned after it is decrypted, and this typically occurs at the client computer.

File filtering

You can configure the appliance to perform different actions on different file types. This is known as file filtering.

You can:

- Enable or disable file filtering.
- Create rules that specify which circumstances will trigger a reaction from the appliance, and what that reaction will be.
- Edit the alert text that is sent when a file is blocked by the appliance.

Any network contains many types and sizes of files, though not all are useful or desirable to your organization:

- Some graphic file formats such as bitmap (suffixed “.BMP”) use large amounts of computer memory and can affect network speed when transferred. You might prefer that users work with other more compact formats such as GIF or JPEG.

For example, if your organization produces computer software, you might see executable files (suffixed with the file name extension “.EXE”) moving around the network. Within any other organization, those files might be games or illegal copies of software. Similarly with movie files (suffixed “.MPEG”), unless your organization handles files of this type, they are probably for entertainment only.

- Much of your organization’s most valuable information — such as designs and lists of customers — is in databases or other special files, so it is important to control the movement of these files. However, it is possible to make any file masquerade as another. An employee with malicious intent might rename an important database file called CUSTOMERS.MDB to NOTES.TXT and attempt to transfer that file, believing that it cannot be detected. Fortunately, you can configure the appliance to examine each file based on its content or *file format*, and not on its file name extension alone.

The file-filtering rules provided enable you to examine any file in several ways:

- Name of the file, such as GOODGAME.EXE.
- Type of the file as indicated by its extension, such as *.EXE and *.JPG.
- Format of the file as indicated by its content such as spreadsheet or graphic content.
- Size of the file, whether above or below a specified size.

When you create settings to control the use of any file, remember that some departments within your organization might need fewer constraints. For example, a marketing department might use large high-quality graphic files for advertising purposes.

HTML settings

You can configure how the appliance handles certain elements and components embedded in HTML data.

Mail settings

You can specify how the appliance handles e-mail messages that use the MIME format. You can:

- Specify the action the appliance should take when a partial message is detected. A partial message is a message that has been divided into smaller parts for sending as several separate e-mail messages.
- Specify the action the appliance should take when a message contains a reference to an external resource and the scheme needed (usually FTP) to retrieve that resource. These messages are known as *external-body messages*.
- Specify which alert message should be used, and if necessary, customize the alert text.
- Add a prefix to the subject line of a message.
- Tell the appliance how to handle MIME messages that have corrupt header files.
- Tell the appliance where to position the alert and disclaimer attachments. The text can appear in the body text or be included as an attachment.
- Set the re-encoding options.
- Tell the appliance how to handle MIME header files that contain null characters.
- Tell the appliance how to handle inconsistent line endings.
- Specify the maximum number of MIME parts a message can have before the appliance considers it to be corrupt.
- Specify which MIME types should be treated as text attachments and which MIME types should be treated as binary attachments.

Mail size filtering

You can specify how the appliance handles e-mail messages that:

- Are larger than a pre-defined limit.
- Have attachments that are larger than a pre-defined limit.
- Have more attachments than are allowed.

An attachment, typically a graphic, a document, or a spreadsheet can greatly increase the size of a complete message — a typical memo of a few kilobytes can grow to many megabytes. Normally this flow of information is necessary for your organization to function, but problems arise when attachments are used excessively or when their use is abused.

For example, computer games are sometimes attached to e-mail messages. Each game typically consumes a few megabytes. Large audio or graphics files — whether for entertainment or business purposes — approach similar sizes. Popular items, when copied and forwarded many times over, can add a heavy load to your mail server. All users will suffer from the slower performance.

The appliance allows you to remove attachments from e-mail messages if they exceed a specified size or quantity. Discarded attachments can be replaced by a small text file, which informs the recipient that attachments were removed. You can also specify special actions against any e-mail message that exceeds a specified size overall.

Protected content

You can specify how the appliance handles e-mail messages that contain data that cannot be scanned because it is protected in some way. For example, it is password protected. You can:

- Specify the action the appliance should take when it detects protected content.
- Edit the replacement message.

Scanner control (denial-of-service attacks)

Large or complex files such as compressed files or .ZIP files can take some time to scan. Such files can be used to attack your network, deliberately slowing its performance. For these reasons, you can limit the size to which any file may be expanded and the depth of nesting.


You can also specify the amount of time that the appliance may spend scanning any file.

Depth of nesting in compressed files


To understand the effect of scanning to a depth of nesting, consider the next diagram. This shows a compressed file, which contains documents and a compressed file. That compressed file contains documents and another compressed file, and so on.



- A depth of 2 scans only the non-compressed files inside a compressed file (as shaded). The contents of any compressed files are not scanned.



- A depth of 3 scans the non-compressed files inside a compressed file, plus only the non-compressed files inside any compressed file that it contains (as shaded).



Signed content (digital signatures)

You can specify how the appliance handles signed e-mail messages.

You can:

- Specify the action the appliance should take when an e-mail message is signed.
- Edit the replacement message.

Whenever information is sent electronically, it runs the risk of being accidentally or wilfully altered. To overcome this, some e-mail software uses a digital signature — the electronic form of a handwritten signature. A digital signature is extra information added to a sender's message, which identifies and authenticates the sender and the information in the message. It acts like a unique summary of the data. Typically, a long string of letters and numbers appears at the end of a received e-mail message. The e-mail software then re-examines the information in the sender's message, and creates a digital signature. If that signature is identical to the original, you can be sure that the data has not been altered.

While this method is useful most of the time, it can cause problems if the message violates your policy. For example, if the message contains a virus, bad content, or is too large, the appliance might clean or remove some part of the message. The original digital signature is now 'broken'. In other words, although the message is still valid (and usually readable), its signature is invalidated. Now the recipient cannot rely on the contents of the message at all because the contents might also have been altered in other ways.

You need to consider carefully how you handle signed messages. You can choose one of the following actions:

- **Refuse the original data and return a rejection code** — the appliance rejects the e-mail message and sends an SMTP 550 command to the mail server.
- **Accept and then drop the data** — the appliance accepts the e-mail message and discards it. It sends an SMTP 250 command to the mail server.
- **Allow changes to break the signed E-mail** — the appliance modifies the message, even if this breaks the signature. The modified e-mail message is sent to the recipient.
- **Do not allow changes to break the signed E-mail** — the appliance only performs actions that will not break the signed e-mail message signature. It then attempts to deliver the e-mail message to the original recipients.
- **Allow through** — some e-mail software might not accept any changes to the signed message, and therefore you cannot allow the appliance to alter the content. The danger here is that if you choose to allow all signed messages through, an undesirable item can escape detection if it is inside a signed message. If you allow all signed messages through, you need to be sure that the messages come from a trusted source, or that they will be scanned at a later stage. Any scanner detections will be logged but not acted upon.

In all cases, you can record the arrival of signed messages and notify an administrator.

Connection policies

The **Policy -> SMTP -> Advanced Policies -> Connection -> From Outside** menu options can be used to set up policies for SMTP connections initiated by hosts in your **Outside Networks** list.

The **Policy -> SMTP -> Advanced Policies -> Connection -> From Inside** menu options can be used to set up policies for SMTP connections initiated by hosts in your **Inside Networks** list.

You can configure the following SMTP connection policy:

- [Anti-Relay \(routing characters\) on page 131.](#)
- [Time-outs on page 132.](#)
- [Transport logging on page 132.](#)

Anti-Relay (routing characters)

An e-mail address can contain routing characters, such as “%” and “!” which enable a message to be passed between computers. You can permit or block this form of relaying by specifying which routing characters you will permit or deny.

For example, if you want to block the relaying of addresses of the type “user@host”@relay.com, add *@* to the list of characters blocked.

Typically you would set up the routing characters as described below:

- Do not enter any patterns in **Permit routing characters**.
- Enter the standard patterns shown below in **Deny routing characters**:

Pattern	Pattern Description
%	Right-binding route character (%-exploit).
!	Local or mail gateway routing.

NOTE

An e-mail address consists of parts before and after the final @ sign. These patterns work on the part of the address before the final @ sign. The appliance examines the destination e-mail address for these patterns.

Entering these patterns in **Deny routing characters** prevents computers inside your network relaying e-mail messages (spam) on behalf of unauthorized users.

Time-outs

You can configure the following connection time-outs:

- Client to appliance time-outs.
- Appliance to server time-outs.
- Total appliance to server time-outs.

See the online *Quick Help* for a more detailed description of each of the time-out settings.

Transport logging

The appliance can log certain events that occur when handling SMTP e-mail messages. You can enable or disable transport logging.

An entry in the log will be created when:

- The sender is listed in the Blackhole Lists.
- The sender is listed in the Deny Sender list.
- The sender is listed in the Anti-Relay Deny Domains list.
- An SMTP 250 OK message is sent to the sender's mail server.
- The appliance receives an SMTP 250 OK message from the receiving mail server.
- An e-mail message is deferred.
- The appliance attempts to deliver a deferred e-mail message.
- The appliance has successfully delivered a deferred e-mail message to the next mail server, and received an SMTP 250 OK message from that mail server.
- The appliance has determined what eventually happened to the e-mail message. For example, it was delivered, refused, or accepted and dropped.

You can also configure the transport log to include additional entries when the following scanning events occur:

- A virus is detected.
- Banned content is detected. For example, banned words.
- Banned content type is detected. For example, banned file types.
- An e-mail is encrypted, signed, a partial e-mail message, or corrupt.
- Spam is detected.

Where relevant, the transport log will record information about:

- The sender.
- The recipients.
- The number of recipients.
- The e-mail message identification number.
- The size of the e-mail message.
- The mail relay used to forward the e-mail message (host name and IP address).
- Which scanner was triggered.

Protocol policies

The **Policy -> SMTP -> Advanced Policies -> Protocol -> From Outside** menu options can be used to set SMTP-specific policies that control the communication between the appliance and hosts listed in your **Outside Networks** list.

The **Policy -> SMTP -> Advanced Policies -> Protocol -> From Inside** menu options can be used to set SMTP-specific policies that control the communication between the appliance and hosts listed in your **Inside Networks** list.

You can configure the following SMTP protocol policies:

- [Data command options on page 133.](#)
- [Denial-of-service prevention on page 134.](#)
- [E-mail address configuration on page 134.](#)
- [Message processing on page 135.](#)
- [Transparency options on page 136.](#)

Data command options

You can set up how the appliance responds to e-mail messages that exceed any of the following parameters:

- Maximum amount of data that can be received during the DATA phase of the communication.
- Maximum number of characters per line.
- Maximum number of Hops allowed (that is, the maximum number of Received lines allowed in the e-mail header).

You can also specify the maximum number of megabytes that the e-mail message can have before the appliance closes the connection.

Denial-of-service prevention

To prevent denial-of-service attacks the appliance can close a connection if one or more of the following conditions occur:

- The average data throughput over a set interval is less than a pre-defined value.
- The number of trivial commands received before the appliance receives a successful DATA command is exceeded.
- The maximum command length permitted by the RFC standard is exceeded.
- The length of the SMTP conversation (defined as the time between the opening of the connection and receiving the final dot (.) command) exceeds a pre-set time.
- The AUTH phase of a communication exceeds a pre-defined limit (Transparent Bridge mode only).
- The maximum number of recipients allowed is exceeded. The appliance can:
 - ◆ Sends the SMTP failure response.
 - ◆ Delay the response by a set amount of time.

E-mail address configuration

The appliance can generate an e-mail message, if it receives an e-mail message that triggers a scanner detection, or if it is unable to deliver the original e-mail message.

The appliance can generate:

- A *notification* e-mail message — the appliance can be set up to notify a recipient that an e-mail message has triggered a detection setting. The notification e-mail message does not contain the original e-mail message. It only contains the *notification*, in the form of an HTML attachment.
- An *annotated* e-mail message — an annotated e-mail message contains the original e-mail message and some additional text, known as a *notification*. The notification text is sent as an HTML attachment.

You can:

- Specify which addresses should be used in the TO and FROM fields of the e-mail messages generated by the appliance.
- Specify the notification text that is used for annotated and notification e-mail messages generated by the appliance.
- Set up advanced address parsing options.

For more information about e-mail address configuration, refer to the *Product Guide*.

Message processing

You can set up:

- *A Welcome Message.*
- *Store and forward options.*
- *DNS data limits.*
- *Advanced SMTP options.*

A Welcome Message

When a host using SMTP connects to an appliance in Explicit Proxy mode, by default, the following welcome message is displayed:

<appliance name and domain>WebShield<product number>/SMTP Ready

You can use the welcome message option to replace this message with some other text of your choice.

Store and forward options

The appliance can be configured to store e-mail messages when the maximum size of the e-mail message is exceeded, or the number of recipients exceeds a pre-defined limit.

The appliance will attempt to deliver those e-mail messages at a later time.

DNS data limits

When the appliance tries to deliver an e-mail message by doing a DNS look-up, it examines the number of Mail Exchange records (MX records) and Address records (A records) returned by the DNS server.

You can limit the number of delivery attempts the appliance will make by limiting the number of MX and A records returned by the DNS server.

Advanced SMTP options

You can set up the following advanced SMTP options:

- Send SMTP traffic to a different port number.
- Specify the maximum number of policies that can be applied to an e-mail message.
- Add the WebShield IP address to the Received e-mail header.
- Force the HELO command to automatically perform a reset (RSET command).
- Force the use of the HELO or EHLO command in any SMTP communication.

Refer to the *Product Guide* for more information about these options.

Transparency options

These options only applies to appliances operating in Transparent Router or Transparent Bridge mode.

You can set up the appliance to:

- Use the Welcome Message from the mail server, or use its own Welcome Message.
- Add some text to the front of the Welcome Message provided by the mail server.
- Allow Extended Simple Mail Transfer Protocol (ESMTP) extensions. For example, Delivery Sender Notification (DSN), Authentication (AUTH), and 8BITMIME.
- Send keep-alive commands during the DATA phase to prevent the connection timing-out. The NOOP command is used as the keep-alive command.
- Generate additional scanning alerts to warn a network administrator or other users when specific events occur. For example, the appliance can warn users when viruses, spam, or banned content have been detected.
- Allow or prevent the use of multiple policies for e-mail messages with more than one recipient.
- Add a Received header to the e-mail.

How e-mail messages are processed

When the appliance receives an e-mail message it processes it in the following order:

- 1 Permit Sender checks.
- 2 Deny Sender checks.
- 3 Anti-relay checks are made in the following order:
 - a Permit domains
 - b Deny domains
 - c Local domains
- 4 Scanning checks — the scanning checks are made in the following order:
 - a If the optional *SpamKiller for WebShield appliances* is installed and activated, the anti-spam scan will be performed first.
 - b Anti-virus and content scanning are performed next, and are performed in parallel.

NOTE

If there are actions associated with anti-virus scanning and content scanning detections, the highest priority action will be performed. The priority the appliance gives to actions is pre-determined and cannot be configured.

- 5 Delivery checks — the delivery checks are made in the following order:
 - a Domain relays
 - b DNS
 - c Fallback relays

Multi-policies for e-mail messages

If an e-mail message is sent or received and it has more than one recipient, the appliance needs to know which policies to apply to that e-mail message.

If the recipients are in the same policy group, the appliance will apply the policies associated with that policy group.

If the recipients are in different policy groups, with perhaps conflicting policies, the appliance must decide how best to handle that e-mail message. How it handles the e-mail message depends on:

- Which policies need to be applied — Content policies, Connection policies, or Protocol policies.
- Its operational mode — Transparent Bridge mode, Transparent Router mode, or Explicit Proxy mode.
- If the handling of multiple policies is enabled (Transparent mode only).
- The maximum number of policies allowed to apply to a single e-mail message — as specified in **Policy -> SMTP -> From Inside/From Outside -> Advanced Policies -> Protocol -> Message Processing -> Advanced**.
- The priority assigned to a non-global policy by its order in the tree node.

The following sections explain how e-mail messages with multiple recipients are handled for each of the policy types.

Protocol policies

If the appliance receives an e-mail message with multiple recipients, and it needs to apply protocol policies to that e-mail message, it always applies the highest priority policy to *all* of the recipients.

The priority is determined by the order in which the non-global policies are within the tree node. See the *Product Guide* for more information about ordering non-global policies.

It does this regardless of which operational mode is being used by the appliance.

Connection policies

If the appliance receives an e-mail message with multiple recipients, and it needs to apply connection policies to that e-mail message, it always applies the highest priority policy to *all* of the recipients.

The priority is determined by the order in which the non-global policies are within the tree node. See the *Product Guide* for more information about ordering non-global policies.

It does this regardless of which operational mode is being used by the appliance.

Content Policies

If the appliance receives an e-mail message with multiple recipients, and it needs to apply content policies to that e-mail message, it handles the e-mail message as described in the following table:

Operational Mode	Multi-policies setting:	Maximum number of policies:	How the e-mail message is treated
Explicit Proxy mode	Not applicable	Not exceeded	<p>The e-mail message is effectively replicated according to the different policies that must be applied.</p> <p>Each replicated e-mail message passes through the scanners separately, and the relevant policies and actions are applied.</p> <p>Separate entries appear in the logs and reports for each replicated e-mail message.</p> <p>If the appliance is configured to generate alerts, separate alerts will be generated for each replicated e-mail message.</p>
Explicit Proxy mode	Not applicable	Exceeded	Only the highest priority policy is applied, and it is applied to all recipients.
Transparent Bridge or Transparent Router	Disabled	Not applicable	Only the highest priority policy is applied, and it is applied to all recipients.
Transparent Bridge or Transparent Router	Enabled	Exceeded	Only the highest priority policy is applied, and it is applied to all recipients.
Transparent Bridge or Transparent Router	Enabled	Not exceeded	<p>The e-mail message is effectively replicated according to the different policies that must be applied.</p> <p>Each replicated e-mail message passes through the scanners separately, and the relevant policies and actions are applied.</p> <p>Separate entries appear in the logs and reports for each replicated e-mail message.</p> <p>If the appliance is configured to generate alerts, separate alerts will be generated for each replicated e-mail message.</p> <p>The replicated e-mail messages are delivered using proxy delivery methods.</p>

If the highest priority method is applied, the e-mail message is not replicated and only passes through the scanners once.

Content policies and performance issues

When setting up content policies, it is particularly important to make sure that the majority of e-mail messages will be covered by the global policy.

Each time the appliance has to replicate an e-mail message in order to apply a different non-global policy to it, the e-mail message is rescanned.

The amount of scanning required has an impact on the general performance of the appliance.

Number of scans example

You could have set up two non-global policy groups, called *Directors* and *Managers*.

- director@example.com is a member of the *Directors* policy group.
- manager@example.com is a member of the *Managers* policy group.

You have assigned a **Mail Size Filtering** policy to each of these policy groups:

- Members of the *Directors* policy group can receive e-mail messages with attachments over 32000 kilobytes.
- Members of the *Users* policy group can only receive attachments if they are less than 32000 kilobytes. If an attachment is more than 32000 kilobytes, the *Users* policy group will receive an HTML alert instead of the attachment.

Scenario 1

The appliance receives an e-mail message containing a 35000 kilobyte attachment that is addressed to the following recipients:

- director@example.com
- manager@example.com

The appliance must scan the e-mail twice and perform 2 different actions, because the e-mail message is addressed to two different recipients. Each recipient is affected by a different policy with different actions:

- For the recipient director@example.com — the appliance will allow the attachment through.
- For the recipient manager@example.com — the appliance will replace the attachment with an HTML alert, telling manager@example.com that the attachment has been removed.

Scenario 2

The appliance receives an e-mail message containing a 35000 kilobyte attachment that is addressed to the following recipients:

- director@example.com
- manager@example.com
- user@example.com

The appliance must scan the e-mail three times and perform 3 different actions, because the 3 different recipients are members of 3 different policy groups.

The appliance will apply the actions for director@example.com and manager@example.com as described in scenario 1.

However, user@example.com is not a member of either of the policy groups affected by the *Directors* and *Managers* policies.

When deciding how to handle the attachment for user@example.com, the appliance must refer back to the default global policy and apply whatever action the global policy states.

For example, the global policy could be **From Outside**, and the **Mail Size Filtering** setting state that the appliance must **Refuse the original data and return a rejection code**.

When a scanner triggers

The appliance can be configured to perform certain actions when a scanner triggers. For example, when the anti-virus scanner detects a virus, you can tell the appliance to attempt to clean the e-mail message.

The actions that are available depend on the selected SMTP policy setting and on which scanner detected the issue.

This section summarizes all of the possible SMTP actions that could occur.

The actions can be divided into *Primary Actions* and *Secondary Actions*. Primary actions determine how the original e-mail message is handled. Secondary actions apply to additional, copied e-mail messages, and notifications.

Topics include:

- [Primary actions on page 142.](#)
- [Secondary actions on page 143.](#)
- [Setting up secondary actions in transparent modes on page 145.](#)

Primary actions

When you are configuring SMTP content policy settings, you can select one or more of the following primary actions:

- **Refuse the original data and return a rejection code**

The appliance rejects the e-mail message and sends an SMTP 550 command to the mail server.

- **Accept and then drop the data**

The appliance accepts the e-mail message and discards it. It sends an SMTP 250 command to the mail server.

- **Do not allow changes to break the signed E-mail**

The appliance only performs actions that will not break the signed e-mail message signature. It then attempts to deliver the e-mail message to the original recipients.

- **Allow changes to break the signed E-mail**

If a scanner is triggered due to the content of an e-mail message, the appliance modifies the message, even if this breaks the signature. The modified e-mail message is sent to the recipient.

- **Replace the content with an HTML alert**

If a scanner is triggered due to the content of an e-mail message, it replaces the content with an HTML alert. The modified e-mail message is then sent to the recipient.

- **Remove the content**

You can configure the appliance to limit the number and size of attachments it will scan. If an e-mail message exceeds these limits, the appliance removes the excess content, scans the remaining e-mail message, and sends the modified e-mail message to the recipient.

- **Clean the content**

If the appliance detects a virus within the e-mail message, it will attempt to clean that virus. If the e-mail message can be cleaned, the modified e-mail message is sent to the recipient.

You can also tell the appliance how to handle files that have zero bytes after they have been cleaned. The zero byte file options are accessed using the **Advanced** button in the anti-virus actions page.

You can:

- ◆ **Keep zero byte files** — the appliance will allow files that have been cleaned to have zero bytes.
 - ◆ **Remove zero byte files** — the appliance will remove any file that has zero bytes after it has been cleaned.
 - ◆ **Treat zero byte files as a failure to clean** — the appliance will treat the files as if it cannot clean them, and performs the action specified in **If cleaning fails, take the following action**.
- **Allow Through**

The appliance will let the e-mail message through. Any scanner detections will be logged but not acted upon. For example, you would select this action if you want to monitor the use of certain words in e-mail messages, without preventing their use.

Secondary actions

When you are configuring SMTP content policy settings, you can select one or more of the following primary actions:

NOTE

Some of these actions require that e-mail delivery methods are set up as described in **E-mail address configuration** in the *Product Guide*. You can also use the **E-mail Address Configuration** option to change the text that appears in the e-mail notifications and annotated e-mail messages described below.

- **Deliver modified e-mail message to sender**

The appliance returns a copy of the modified e-mail message to the sender.

- **Deliver a notification E-mail to the original recipient(s)**

If a scanner is triggered due to the content of an e-mail message, the appliance can send an e-mail message to the original recipients. This e-mail message can be used to tell the recipient that there has been a problem, and it is known as a *notification* e-mail. You can use the **E-mail Address Configuration** option to change the text that appears in the notification e-mail message.

- **Deliver an E-mail alert to a GUI defined recipient**

If a scanner is triggered due to the content of an e-mail message, the appliance can send an e-mail message to a user-defined recipient. This e-mail message can be used to tell the recipient that there has been a problem, and it is known as a *notification* e-mail.

- **Deliver a notification E-mail to the original sender**

If a scanner is triggered due to the content of an e-mail message, the appliance can send an e-mail message to the original sender. This e-mail message can be used to tell the sender that there has been a problem, and it is known as a *notification* e-mail. You can use the **E-mail Address Configuration** option to change the text that appears in the notification e-mail message.

- **Quarantine the original E-mail**

The appliance can be set to quarantine e-mail messages.

- **Quarantine the modified E-mail**

The appliance can be set to quarantine e-mail messages that have been modified by the appliance.

- **Forward the original E-mail to a GUI defined recipient(s)**

When a scanner triggers, the appliance can be configured to forward the offending e-mail message to a user-defined recipient. For example, the e-mail message could be forwarded to the e-mail administrator, or to a spam administration mailbox.

- **Forward the modified E-mail to a GUI defined recipient(s)**

The appliance can be configured to forward an e-mail message that it has modified to a user-defined recipient. For example, the e-mail message could be forwarded to the e-mail administrator, or to a spam administration mailbox.

- **Deliver an annotated original E-mail to a GUI defined recipient**

The appliance can be configured to send an annotated e-mail message to a user-defined recipient. An annotated e-mail message, is similar to a notification e-mail message, except that it contains the e-mail message as an attachment. If this action is selected, it will contain the original e-mail message.

- **Deliver an annotated modified E-mail to a GUI defined recipient**

The appliance can be configured to send an annotated e-mail message to a user-defined recipient. An annotated e-mail message, is similar to a notification e-mail message, except that it contains the e-mail message as an attachment. If this action is selected, it will contain the modified e-mail message.

Setting up secondary actions in transparent modes

By default, most of the secondary actions are not available when the appliance is operating in Transparent Router mode, or Transparent Bridge mode. Only the quarantine actions are available by default.

To enable secondary actions for transparent modes:

- 1 Set up e-mail delivery addresses as described in the *Product Guide*.
- 2 Enable **Allow WebShield to generate additional scanning alerts**, as described in *Generating additional scanning alerts* in the *Product Guide*.

This section describes spam and how you can use the appliance to control the amount of spam reaching users.

What is SPAM?

Any unsolicited and unwelcome e-mail messages can be considered spam. Spam includes commercial e-mail messages, the electronic equivalent of “junk mail,” and unwanted non-commercial e-mail messages, such as virus hoaxes, jokes, and chain letters.

Frequently, people who create spam, known as *spammers*, forge the headers of the e-mail messages to hide their true identity, often deflecting retaliation toward innocent parties.

What is SpamKiller for WebShield appliances?

McAfee SpamKiller for WebShield appliances is an optional add-on for certain WebShield appliances. When activated, it is fully integrated into the WebShield appliance and user interface, and provides additional protection against spam at the network perimeter.

SpamKiller for WebShield appliances uses the McAfee SpamAssassin anti-spam engine (called the anti-spam engine in this guide), and a set of anti-spam rules to scan e-mail (SMTP) traffic for potential spam.

We say “potential spam” because identifying spam is not an exact science. Anti-spam rules can only identify characteristics within the e-mail message that make it more likely that the e-mail message contains spam. For example, an anti-spam rule might look for certain words or phrases that typically appear in spam e-mail messages, such as “get rich quick.”

It is important to maintain a balance between blocking potential spam and allowing normal e-mail communication through:

- If your anti-spam measures are too stringent, normal e-mail communication could be wrongly identified as spam and blocked. Users will complain that they are not receiving the e-mail messages they were expecting.
- If your anti-spam measures are not stringent enough, too much spam will get through and interfere with normal communication.

Maintaining the right balance is quite difficult for a number of reasons:

- The nature of spam is always changing; as the people who write spam (known as *spammers*) change their tactics to avoid detection.
- The definition of spam changes according to the context. For example, a joke received at home from a friend might not be considered spam; the same joke sent to 3000 employees might be considered spam by the employer.
- There will probably be exceptions to specific rules. For example, you might want to block commercial spam, unless it comes from organizations within your business sector, in which case, you might need to receive it to keep up-to-date with their products and promotions.

NOTE

For these reasons we cannot guarantee that the anti-spam software will detect and block all e-mail messages that could contain spam.

SpamKiller for WebShield appliances works with the appliance's existing anti-spam features to help you maintain the best balance between blocking potential spam and letting normal e-mail communication through. In particular:

- To counter changing spammer tactics, McAfee Security regularly updates the anti-spam engine and anti-spam rules files used to detect spam. These files can be automatically downloaded using the WebShield appliance's update facility. You can even download special rules, known as *extra rules*, that have been designed to combat a sudden outbreak of a specific type of spam e-mail message.
- You can set up separate inbound and outbound anti-spam policies, and specify the level of spam detection that should be used for each policy.
- You can use the existing **Permit Sender** feature to allow e-mail messages from specific senders, networks, and domains to bypass anti-spam scanning.

The anti-spam software also allows you to decide who deals with the spam, once it has been identified. You can:

- Deal with spam at the appliance level, so that it never reaches the end users. For example, e-mail messages that contain potential spam can be refused, discarded, or forwarded to a special mailbox.
- Use the appliance to add a warning to e-mail messages containing potential spam, and let the end users choose how to deal with those e-mail messages when they receive them. For example, the mail administrators and end-users can set up their mail clients to automatically place spam into a special spam mail folder.

You can control the amount of spam that your organization receives by blocking all e-mail messages from known unwanted senders, marking the subject line of any suspicious e-mail messages, deleting messages, or moving messages to a safe area (or “quarantine”). Additionally, you can inform an administrator of the detection, or record the event in a log.

The appliance provides several techniques to guard against the nuisance of spam e-mail messages:

- Blacklists

The appliance matches every e-mail message against a *blacklist*. A blacklist is the list of e-mail addresses from which your company does not want to receive messages because those messages are always spam messages. Besides blacklisting “From” addresses, you can also blacklist “To” addresses. For example, if an e-mail address in your organization receives a large amount of spam, you can prevent that address forwarding any e-mail.

- Whitelists

The appliance matches every e-mail message against a *whitelist*. A whitelist is the list of e-mail addresses that you trust not to send unwanted messages. The list can contain addresses of business partners or organizations that sell essential products. Such messages are allowed through without scanning for spam phrases.

- Rules and scores

The appliance matches an extensive set of rules against every e-mail message. Each rule is associated with a score — positive or negative. Rules that match for spam-like characteristics give a positive score. Rules that match attributes of legitimate messages give a negative score. When added together, the scores give each message an overall spam score. Some rules are simple, and match only on popular phrases. Other rules are more complex and match on the header information and structure of e-mail messages.

Understanding spam scores

Spam often contains well-known phrases. For example, these phrases are good indicators:

Phrase	Spam score per phrase
Dear Friend	1.5
amazing offers	1.0
believe your eyes	1.2
incredibly low	0.8
best ever	0.8

(The values shown here are for example only. The actual values might be different in the product. This example is deliberately simple, and does not attempt to demonstrate any complex matching.)

Consider the following two messages. The phrases are highlighted for clarity.

Message	Total spam score
Dear John, Our computer suppliers have some amazing offers on PCs this year. I'll send you their catalogue and discuss my requirements with you on Tuesday. Looking forward to our best ever year on this project! Regards, Peter	 $1.0 + 0.8 = 1.8$
 Dear Friend, See our web site for amazing offers on PCs. You won't believe your eyes! These incredibly low prices are our best ever!	 $1.5 + 1.0 + 1.2 + 0.8 + 0.8 = 5.3$

The second message has a higher score, which indicates that it is *possibly* spam. It is possible for a legitimate message to attain a high score. Therefore, the detection of spam cannot be precise. You can determine how the appliance will respond to messages based on their spam scores:

- You can specify a level at which you regard a message as spam. Typically, a score of 5 indicates that a message is spam. You can inform the recipients that a message is likely to be spam by adding some text, such as **** SPAM ****, to the subject line of the message. Recipients can then easily identify a spam e-mail message, and decide how to handle the message. For example, some e-mail products such as Microsoft Outlook and Lotus Notes can redirect mail to specific folders based on rules or filters.

- You can specify a level at which the appliance will handle spam messages automatically. For example, the appliance can automatically block or quarantine messages that have high spam scores. In addition, you can inform an administrator or log the event.
- You can specify that the appliance adds a report to a message's Internet headers that tells its recipients of any rules that triggered and the message's spam score. You can choose whether to add the report, and whether such information is included in all messages or only those messages that the appliance identifies as spam.

The report includes a spam score and optionally a *spam score indicator*. For example, a spam score of 5.6 can have an indicator of five asterisks, and a spam score of 6.2 can have an indicator of six asterisks. The indicator is rounded to the integer and ignores any decimal fraction. The indicator provides a simple character string for filtering messages.

We recommend that you set this option for initial testing only, because it can impact your server's performance. When you have the information that you need, turn the option off.

Disabling rules

The appliance contains numerous anti-spam rules that it applies against e-mail messages. However some rules might not be appropriate for your organization, so you can disable them.

For example, advertisements for unproven slimming aids are common, so a rule that detects the phrase "weight loss" is useful for identifying a possible spam message. However, if your organization produces health products, you might not want this rule applied against your e-mail messages.

Tips for avoiding spam

We recommend the following tips to reduce unwanted e-mail messages. Make these tips available to users to help them reduce the amount of spam they receive:

- Use a different e-mail address or "public" e-mail address when participating in news groups, joining contests, or responding to any third-party requests online.
- Avoid using a Reply or Remove option. Some senders remove the address, but others record the e-mail address and later send more spam, or sell the address to other spammers.
- Limit Internet usage at work. When at work, do not access sites that are not business-related such as message boards, e-trade sites, Internet auctions, and e-commerce sites.

- Do not post e-mail addresses online. Know whether your e-mail address will be displayed or used before posting an e-mail address online. Read the privacy policy on the web site before posting your address and opt out, if possible.
- Beware of purchasing products that are advertised by spam. When you respond to this type of e-mail message, you often make more personal information such as your name, address, telephone number or credit-card numbers available to spammers, which can lead to increased spam. Furthermore, in order to provide themselves with an income, spammers must issue large numbers of e-mail messages in order to get enough responses. By not responding at all, you can discourage this advertising technique by making it unprofitable.

Updating your anti-spam software

The **Update** -> **Anti-Spam** option is only available if *McAfee SpamKiller for WebShield appliances* is being evaluated or is fully activated. It allows you to respond to the ever changing nature of spam, by regularly downloading the latest anti-spam files.

The anti-spam files help you maintain a balance between the e-mail messages you want to filter out because they probably contain spam, and those that you want to let through because they are unlikely to contain spam. You can download:

- **Anti-spam rules** — define what is spam. There are anti-spam rules that are updated on a regular basis, and rules that are only produced in response to an urgent need to combat a sudden outbreak of a specific type of spam (these rules are called *extra rules*).
- **Anti-spam engine** — uses anti-spam rules to scan e-mail messages for spam.

E-mail (SMTP) content scanning

10

This appliance uses content rules to prevent SMTP e-mail messages with unwanted content reaching their intended recipients.

Content Rules and Rule Groups

A content rule is used to define unacceptable content. For example, you could set up a content rule to ban the use of a particular swear word in e-mail messages entering or leaving your organization.

That content rule can then be assigned to an SMTP content policy, and the appliance configured to perform certain actions when that content rule is triggered. For example, when the appliance detects the swear word in an e-mail message it could refuse the e-mail message and send a warning to a network administrator that an offensive e-mail message has been detected.

As you could create a large number of content rules, content rules are organized into *Rule Groups*. Each rule group has one or more *Content Rules*.

For example, you can create a rule group, called *Offensive descriptions*, then within the group you can create a content rule that detects *cruel*, another content rule that detects *unkind*, and another content rule that detects *uncaring*, and so on.

You can assign the whole rule group to a policy or just assign selected content rules. Assigning selected content rules allows you to set up policy specific settings for those content rules.

The appliance comes pre-configured with a standard set of rule groups. You can add content rules to these rule groups, or create new rule groups of your own.

You create the rule groups first, and then add content rules to them.

Content rules are likely to grow in number and complexity over time, and it is important to think about how you will group content rules, and what each rule group and content rule should be called.

Importing and exporting content rules

Having created a content rule, you can share its rules and settings with other computers and with our other products, because rules can be imported and exported as text files in XML format.

Scanning for content

You can have a large number of content rules, and each content rule can specify words in various combinations. The content rules can be simple such as detecting the use of a single word or phrase. They can be more complex and include combinations of phrases that appear closely together. A complex content rule can allow the use of a word in one situation, but prevent its use in others.

Typically, you will want a content rule to scan for undesirable words in the content of each e-mail message. However, you can also scan the following items:

- Content in attachments.
- Names of files attached to e-mail messages.
- Name of sender.
- Name of recipient.
- Name of domain.

Creating content rules

To use content scanning, you create *rules* using the following steps:

- 1 *Giving a name and description to the rule on page 155.*
- 2 *Specifying where the rule applies on page 155.*
- 3 *Specifying the action to take when the rule is triggered on page 155.*
- 4 *Adding optional advanced features on page 157.*

These steps are next described in more detail.

Giving a name and description to the rule

Over time, you can create many rules, so each needs an accurate name and description.

Remember that when the rule is triggered, the *name* of the rule appears in the alert message that users see. Therefore, if you are trying to prevent the use of an insulting phrase, do not include that phrase in the name of the rule. Instead, name your rule as something like “Ban Insult 23.”

Each rule can also have a *description*. You can provide more information here about the purpose of the rule. The rule’s description does not appear in the alert message.

Specifying where the rule applies

A banned phrase might appear inside a variety of files or documents. You can specify precisely which types you want to scan.

A banned phrase might appear in the body of the message, its subject line or even inside a plain-text attachment. The appliance can scan the file name of any attachments too, so you can block attachments by exact name, such as *goodgame.exe*, or by file extension such as *.JPG files. A scan on each message sender can be used to block known unwanted senders, especially those sending nuisance mail.

See also the anti-spam features, described in [Scanning for Spam on page 147](#).

Specifying the action to take when the rule is triggered

You can take several actions against any item that triggers a rule. The available range of actions depend on your product’s configuration, but probably includes the following actions:

- **Quarantine** — the appliance places the item in a quarantine area, where you can examine the item and decide how to handle it.
- **Replacement** — the item is automatically replaced by a ready prepared document that explains why the original was replaced.
- **Allow through** — the item is not changed but the event might be logged or the administrator is alerted.

Specifying the word or phrase you want to detect

You can specify precisely how a word or phrase appears by specifying its case, using wild cards, and specifying its position:

- Ignoring case.

Normally the appliance scans for the word or phrase exactly as it is written. If you specify that case is to be ignored, the appliance matches the word or phrase regardless of its case. So, "abc" will match abc, Abc, ABC and aBc, or any combination of uppercase and lowercase letters contained in the phrase.

- Using wildcards.

With this feature, you can use the characters * and ? to represent missing characters:

- ◆ ? represents any single character.
For example, "??g" will match dig, dog and tug.
- ◆ * represents any number of characters including none at all.
For example, "s*ing" will match sing, singing and sting.

- Specifying characters at the start or end of words.

You can match characters that appear only at the start of a word.
For example, "hat" matches hat, hate, hats, and hatter.

You can match characters that appear only at the end of a word.
For example, "hat" matches hat, that and what.

You can match characters at the start and at the end of a word. This is sometimes called "exact word matching."
For example, "hat" matches hat but does not match hate, that, or what.

You can match characters anywhere in the word.
For example, "hat" matches hat, hate, that and what.

Some types of file use special formatting characters to specify the layout of text. For example, attachments can contain characters to denote word breaks, line breaks, tabs, cells, end of lines, and other format information. See the table, [Word separators](#) on [page 158](#) for details.

Some characters such as currency symbols and accented characters might be difficult to match because of variations in character sets. You might need to experiment to ensure that your rules can detect such characters.

Adding optional advanced features

You can further refine the conditions that trigger a rule by specifying how other words or phrases may appear in combination with the first word or phrase — their *context* and their *nearness*.

- *Words in context with other words.*
- *Words that are near other words.*

Words in context with other words

- A rule may trigger if **all** of the additional words or phrases are present.

For example, a rule is triggered when the name of a secret new product is used in the same e-mail message as the date for the product's launch.

- A rule may trigger if **any** of the additional words or phrases are present.

For example, a rule is triggered when any word appears that is on a list of offensive words, or a list of secret projects.

- A rule may trigger if **none** of the additional words or phrases are present.

A rule is triggered when an offensive word, for example “dog” is used *except* when it was used to specify a type of that animal, for example, a corgi or alsatian.

Words that are near other words

Normally, when you are searching content in any small document, the banned words are near each other. However in a longer document, the words might appear anywhere, and falsely trigger the rule. To avoid this, your rule can consider the nearness of the words.

As a simple example, a rule might trigger if two words such as *ugly* and *manager* appear together within a block of 50 characters. In the following example, the second paragraph will be detected, and the document can be blocked to prevent the insult.

The latest version of the product looks *ugly*. We need to consider several problems. I will discuss improvements with the *manager* of that department.

I attended the meeting about that new product today. The new *manager* is so *ugly*, nobody will ever want to work with him.

This feature is useful in blocking some offensive phrases. They often contain words that do not cause offence when used alone, but become offensive when grouped together.

Note that nearness is best suited to plain text. It cannot accurately interpret character counts in binary files or files that contain complex text formatting.

Definition of a word

A word is any number of characters bounded by a word separator, which is usually some form of punctuation. The following table shows some examples of word separators that are recognized by the appliance.

Table 10-1. Word separators

horizontal tab	line feed	line break
space	exclamation mark (!)	quotation mark (")
number sign (#)	percent sign (%)	ampersand (&)
apostrophe (')	left parenthesis '('	right parenthesis ')'
asterisk (*)	plus sign (+)	comma (,)
hyphen-minus (-)	full stop (.)	solidus (/)
colon (:	semicolon (;)	less-than sign (<)
equals sign (=)	greater-than sign (>)	question mark (?)
commercial at (@)	left square bracket ([)	reverse solidus (\)
right square bracket (])	low line (_)	left curly bracket ({)
vertical line ()	right curly bracket (})	tilde (~)

For a full list of word separators recognized by the appliance, see [Word Separators on page 229](#).

Understanding complex content rules for e-mail messages

E-mail messages typically have a different structure to documents, and this can affect the way that content rules apply.

For example, consider the following text in a document:

I think our manager is stupid and ugly.

To prevent the words “stupid” and “ugly” appearing together in a *document*, you can create a rule with a *complex phrase* — the rule triggers when these words appear together.

The same rule will work on the following simple e-mail message:

To: user1@example.com
From: user2@example.com
Subject: Our manager
I think he is stupid and ugly. What do you think?

Now consider a second example:

```
To: user1@example.com
From: user2@example.com
Subject: Our stupid manager
        I think he is ugly too. What do you think?
```

The complex rule you have already created will not trigger in this case. Most e-mail messages are based on the MIME format, and they comprise several parts. You can think of each part as a separate file — one for the “To” address, the “From” address, the subject line, and the message body. In this example, no part contains both words — “stupid” is in the subject line, while “ugly” is in the message body.

To trigger a content rule on the words “stupid” and “ugly” appearing together in an *e-mail message*, you must create a rule that combines two simple *conditions* — the rule triggers when the word “stupid” appears anywhere in an e-mail message *and* when the word “ugly” appears anywhere in an e-mail message.

Understanding limitations in content scanning

A rule can only apply to a single file, document or attachment at any time.

For example, you may have a rule that triggers on finding the word “ugly” in databases and in spreadsheets. When the appliance encounters any database, it searches for the word “ugly”. Similarly, when the appliance encounters any spreadsheet, it searches for the word.

You can make such a rule more complex. For example, you may make the rule search for both “ugly” and “stupid” in databases and in spreadsheets. When the appliance encounters any database, it searches for the word “ugly” and the word “stupid”. If both words are present, the rule triggers your defined action. When the appliance encounters any spreadsheet, the rule is also triggered.

It is possible to create combinations of rules that will not work. For example, you can create a rule which detects “ugly” in databases, and “stupid” in spreadsheets. If used separately, those rules will work. However a compressed file (such as a WinZip file) could contain a database with “ugly” and a spreadsheet with “stupid”. This combination of files will not be detected.

Examples of content rules

Content scanning enables you to create rules that detect the appearance of words and phrases in many situations and combinations, as in the following examples:

- *Keeping information confidential on page 160.*
- *Reducing network load on page 161.*
- *Blocking offensive words on page 161.*
- *Stopping nuisance e-mail messages on page 161.*
- *Reducing distractions on page 162.*

Each example described here can block e-mail messages — by destroying them, or by moving them to a quarantine area where they can be examined later. You need to be aware of local legislation that affects how e-mail may be treated.

Keeping information confidential

If your organization prefers that details of a new event, product or project are not discussed outside the company, you can prevent the name being discussed in outbound e-mail messages.

For example, your company plans to release a new product called SuperThing. To prevent anyone outside the organization knowing about the product, you need to detect the word inside each e-mail message.

You create a rule called “Confidential information about SuperThing” and apply this rule to a plain-text attachment, the body of the message, and the subject line of message. You specify SuperThing as the word on which to trigger the rule.

As a second example, your organization plans to launch the new product in January. The date must be kept secret. Messages like this must not leave the organization:

We are ready to launch SuperThing in January.

Before that date, less harmful e-mail messages will discuss the product’s details and preparations for its launch. Other products will also be launched, but their dates are less relevant. You do not want to block this message:

The agenda for tomorrow’s meeting:
1 Progress towards the launch of SuperThing
2 How to reduce our stationery costs
3 Launch of MegaBox in January

You can create a rule that triggers only when the two words — SuperThing and January — are close to each other, perhaps within 30 characters.

As a final example, your organization is planning to promote Mr. Jones to the position of CEO. Your rule must trigger on the combination of two words — CEO and Jones.

Reducing network load

The transfer of some file types, such as movie files (MPEGs) and bitmap graphics impact heavily on networks. By creating a list of unacceptable file extensions, you can discourage their use. Your trigger words might be “.BMP” or “.MPG” and you set them to apply to the names of attachments only.

Blocking offensive words

Insulting messages from your own staff or customers might damage the company's reputation. By creating a list of unacceptable words, you can prevent their use.

For example, imagine that it is very offensive to say “You are a dog” to another person. However when used in other contexts, such as discussing types of dog like corgi or alsatian, the word is not offensive. To prevent the word entering or leaving the company in its offensive context, create a new rule called “Offensive word — dog.” You set the rule to apply inside the body of the message, and you set an action to discard such messages. After entering the word dog, you can further refine its context. For example, this rule is to be triggered only if none of these words — alsatian, corgi, spaniel, and so on — appear in the message.

Stopping nuisance e-mail messages

Disgruntled ex-employees, virus hoaxers, and unscrupulous retailers who know the e-mail addresses of your staff can cause problems.

For example, John Smith has been annoying employees by sending unwanted e-mail messages. The content of his messages vary but he always uses one of two e-mail addresses. You create a rule called *Annoying Person*. As the trigger phrase, enter John Smith's two e-mail addresses, and apply the rule to the message's sender only.

The appliance incorporates an anti-spam feature that prevents known spam sources from attacking your network. However, you can also scan content for common or known phrases to limit this kind of attack further. Phrases such as “get rich quick” or “this virus will destroy your computer” can become the trigger phrase for another useful rule.

Reducing distractions

When frequent inappropriate messages are distracting your staff, the appliance can block these messages and deter their senders. For example, advertisements broadcast via e-mail might have “Car for sale” or “House for sale” as their subject line. The messages waste your e-mail resources and distract your staff. To block such an e-mail message, you create a rule called Distracting Advertisements. Specify the trigger phrase as “for sale” and apply the rule to the subject line of a message only.

Many games are sent by e-mail as computer programs (.EXE files). You can block these by creating a rule that triggers when attachments have “.EXE” in their name. This type of rule has an added advantage as games are a popular hiding place for viruses.

This section describes how you can use the anti-virus features provided by the appliance to help protect your network from viruses and other potentially harmful software. The appliance:

- Detects and cleans viruses.
- Automatically scans within compressed files.
- Automatically decompresses and scans files compressed in the packages that include: PKZip, .LHA, and .ARJ.
- Detects macro viruses.
- Upgrades easily to new anti-virus technology.
- Detects polymorphic viruses.
- Detects new viruses in executable files and OLE compound documents, using a technique called *heuristic analysis*.

What is heuristic analysis?

An anti-virus scanner uses two techniques to detect viruses: signatures and heuristic analysis. A *virus signature* is simply a binary pattern that is found in a virus-infected file. Using information in its virus definition files, the scanner searches for those patterns. This approach cannot detect a new virus because its signature is not yet known. Therefore another technique, known as *heuristic analysis*, is employed.

Programs that carry a virus often have distinctive features. They might attempt unprompted modification of files, invoke mail clients or use other means of self-propagation. The scanner analyzes the program code to detect these kinds of computer instructions. The scanner also searches for legitimate behavior, such as prompting the user before taking action, and thereby avoids raising false alarms.

In an attempt to avoid detection, some viruses are encrypted. Each computer instruction is simply a binary number, but the computer does not use all the possible numbers. By searching for unexpected numbers inside a program file, the scanner can detect an encrypted virus.

By using these techniques, the scanner can detect both known viruses and many new viruses and variants.

Anti-virus software

Every month, hundreds of new viruses appear. To ensure that your network is always protected, you must keep your anti-virus software up-to-date.

NOTE

All appliances need updating individually.

The anti-virus software has two parts:

- **Virus definition (DAT) files** — the *DAT files* contain descriptions of new viruses, enabling the engine to detect and clean them. New DAT files are normally released weekly, and occasionally more often.
- **A virus-scanning engine** — the *engine* contains the software to process the DAT files. As new types of virus appear, the engine may need to be improved to handle them. If an outdated engine attempts to use current DAT files, it cannot make full use of their new content. The engine file is upgraded less often than the DAT files, usually every few months.

Both parts are essential to provide full protection against computer viruses, and both must be kept updated. They can both be downloaded from the Network Associates FTP site or other authorized web site using automatic updating.

Why update?

To offer you the best protection possible, we continually update the virus definition (DAT) files that our software uses to detect viruses. Although the software uses heuristic analysis that enables it to detect some previously unknown strains of viruses or malicious code, many new virus types and other agents appear frequently.

Often, your existing software cannot detect these intruders because the virus definition (DAT) files that came with it are outdated. For maximum protection, we strongly recommend that you update your files on a regular basis — at least once per week.

When is the best time to update?

We strongly recommend that you update your virus definition (DAT) files on a regular basis, at least once per week. We produce new DAT files regularly each week, but we occasionally release DAT files (called *EXTRA DAT files*) to counter sudden appearances of new viruses.

You need to consider a time when the network is not too busy — possibly during the night, or during the day but outside normal business hours.

If your business operates on weekends, you might consider updating on Friday evenings and Sunday evenings.

Scheduling the updates

The appliance allows you to schedule regular updates at least once per hour from any one of these sources:

- Our FTP servers or other authorized provider.
- A computer acting as a proxy if your appliance cannot access those servers directly.
- A computer within your network to which the required files have already been downloaded.

The files are compressed to ensure fast downloading.

Use the **Monitor -> Status -> General Status** or **Monitor -> Updates** options to view version number and update information.

Use the **Update -> Anti-Virus** option to set up an update schedule or perform an immediate update.

Local Updates and EXTRA DAT files

The appliance also allows you to download the latest DAT, virus-scanning engine and EXTRA DAT files from a local computer. This facility is available from the **Update -> Anti-Virus** option.

Scanning for viruses

When you prepare settings for scanning viruses and other potentially unwanted software, you need to consider the following:

- Action to take when a virus is found.
- How to handle mass-mailer viruses. See [Blocking specific threats on page 168](#).
- The level of anti-virus protection that you need. See [Setting the level of scanning and type of protection on page 166](#) and [Customizing anti-virus settings on page 166](#).

Setting the action against viruses

You can choose to clean each virus that is detected. If this is not possible, you may delete the infected file or move it to a safe area (or “quarantine”). Additionally, you may inform an administrator of the detection, or record the event in a log.

Setting the level of scanning and type of protection

The appliance provides several levels of anti-virus protection, allowing you to choose high, medium, and low levels of scanning:

- **High** — Most secure. Scans all files, including compressed files.
- **Medium** — Scans executables, Microsoft Office files and compressed files.
- **Low** — Least secure. Scans executables and Microsoft Office files.

Be aware that a higher level of scanning provides good security but can affect performance. In some cases, high levels of scanning are unnecessary if data is being scanned for viruses elsewhere in your network.

In addition, you can customize the scanning by choosing exactly what to scan from a range of options. See [Customizing anti-virus settings on page 166](#).

You can also determine when scanning will occur. See [Content rules on page 108](#).

Customizing anti-virus settings

Besides giving you the preset levels of scanning described in [Setting the level of scanning and type of protection on page 166](#), the appliance also allows you to specify various options when scanning for viruses. Be aware that although more options can provide greater security, the scanning will take longer. The scanning options are described next.

- Detecting possible new viruses in programs and documents.

An anti-virus scanner typically detect viruses by looking for the *virus signature*, which is a binary pattern that is found in a virus-infected file. However, this approach cannot detect a new virus because its signature is not yet known, therefore the scanner use another technique — *heuristic analysis*.

Programs, documents or e-mail messages that carry a virus often have distinctive features. They might attempt unprompted modification of files, invoke mail clients, or use other means to replicate themselves. The scanner analyzes the program code to detect these kinds of computer instructions.

Documents that carry a virus often have distinctive features such as a common technique for replicating themselves. The scanner analyzes the document to detect these kinds of computer instructions.

Program file heuristics scans program files and identifies potential new file viruses. Macro heuristics scans for macros in the attachments (such as those used by Microsoft Word, Microsoft Excel, and Microsoft Office) and identifies potential new macro viruses.

- Scanning inside archive files.

By default, the scanner does not scan inside file archives such as .ZIP or .LZH files because any virus-infected file inside them cannot become active until it has been extracted.

- Scanning default file types.

Normally the scanner examines only the default file types — in other words, it concentrates its efforts on scanning those files that are susceptible to viruses. For example, many popular text and graphic formats are not affected by viruses. Currently the scanner examines over 100 types by default, which includes .EXE and .COM.

- Scanning all files.

Some operating systems such as Microsoft Windows use the extension name of a file to identify its type. For example, files with the extension .EXE are programs. However, if a virus-infected file is renamed with a harmless extension such as .TXT, it can escape detection. The operating system cannot run the file as a program, unless it is renamed later. This option ensures that every file is scanned.

- Scanning files according to file name extension.

Some operating systems such as Microsoft Windows use file name extensions to identify the type of file. For example, files with the extension .EXE are programs, files with the extension .TXT are simple text files. the appliance allows you to specify the types of files you wish to scan according to their file name extension.

- Treating all macros as viruses.

Macros inside documents are a popular target for virus writers. Therefore for added security, you might consider scanning all files for macro viruses, and optionally removing any macro that is found, regardless of whether it is infected.

- Scanning compressed program files.

Compressed files (such as those compressed with PKLITE). If you are scanning selected file extensions only, include the needed compressed file extensions in the list of file extensions to be scanned.

- Scanning suspicious programs.

These programs might be dangerous but they are not viruses. They include programs such as remote-access utilities and password crackers.

- Scanning joke programs.

These programs are not harmful. They play tricks on the user such as displaying a hoax message.

Blocking specific threats

Normally, the appliance handles all potentially harmful software in the same way, however you can specify that certain types are handled differently.

For example, you can configure the appliance to inform the sender, the recipient and an administrator with an alert message whenever a virus is detected in an e-mail message. This feature is useful in that it shows that the anti-virus scanner is working correctly, but it can become a nuisance when a mass-mailer virus is encountered.

Mass-mailer viruses such as Melissa and Bubbleboy propagate themselves rapidly using e-mail services. As a result, numerous alerts are generated, and these can be as irritating as the surge of infected e-mail messages that have been blocked.

A feature in the appliance allows you to handle any mass-mailer virus separately from other types of virus. You can choose to discard the infected document immediately, and thereby suppress any alert messages that would otherwise be generated.

This section provides more information about HTTP web browsing configuration. It includes the following topics:

- [General HTTP configuration](#)
- [Policy-based HTTP configuration on page 170.](#)

General HTTP configuration

The following feature, listed under the **Configure** -> **HTTP** menu option, is not based on specific policies.

Connection settings (Advanced)

For TCP and UDP protocols, ports numbers are used to identify the ends of logical connections which carry specific services. Each service has an associated port number.

The default port number for HTTP is port 80.

NOTE

We recommend that you do not change the port numbers unless you understand port assignments and the implications of changing the port number.

Intercept ports

The **Intercept ports** list only applies to appliances operating in Transparent Router mode or Transparent Bridge mode. You can specify the ports on which the appliance will intercept HTTP traffic.

Listen ports

The **Listen ports** list only applies to appliances operating in Explicit Proxy mode. You can specify the ports on which the appliance will listen for HTTP traffic.

Listeners, connections and memory

You can set up scanning resources on a protocol-by-protocol basis.

For each protocol, you need to set up the number of processes listening for the protocol-specific traffic. These processes are known as *Listeners*. Each listener can handle a number of connections. The number of scans that can be performed simultaneously for each protocol is equal to the number of listeners multiplied by the number of connections.

As there is a finite amount of resource available for scanning, there will always be a resource trade-off between:

- The number of listeners for that protocol.
- The number of connections handled by the listeners for the protocol.
- The amount of memory required for scanning that protocol.
- The scanning resources assigned to other protocols.

When you first use the appliance, some default values are in place. You can restore these settings at any time if the modified values are unsuitable.

Policy-based HTTP configuration

The section briefly describes the purpose of each of the policies you can apply to HTTP traffic and connections.

Caution

Incorrect configuration of settings labelled “advanced” can cause serious security and connectivity issues for your network. For this reason, we recommend that you do not change these settings unless instructed to do so by Technical Support or your network consultant.

Content policies

The **Policy** -> **HTTP** -> **Content** -> **From Outside** menu options can be used to set up policies that tell the appliance how to handle the HTTP requests and responses from hosts in your **Outside Networks** list.

The **Policy** -> **HTTP** -> **Content** -> **From Inside** menu options can be used to set up policies that tell the appliance how to handle the HTTP requests and responses from hosts in your **Inside Networks** list.

You can specify the actions that the appliance will perform in certain circumstances. See [HTTP actions on page 172](#) for a list of possible actions.

You can configure the following HTTP content policies:

- [Alert settings on page 171.](#)
- [Anti-virus settings on page 171.](#)
- [HTML settings on page 172.](#)
- [Scanner control \(denial-of-service attacks\) on page 172.](#)

Alert settings

The appliance will send an HTML message to clients when a specific event occurs. This is known as an *HTML alert*.

You can:

- Change the text that appears at the start of the HTML alert, known as the *alert header*.
- Change the text that appears at the end of the HTML alert, known as the *alert footer*.

Anti-virus settings

The appliance can be configured to detect viruses and other potentially harmful software. For more information about scanning for viruses, see [Virus-scanning on page 163](#).

HTML settings

You can configure how the appliance handles certain elements and components embedded in HTML data.

When users view a web page, their browsers can download ActiveX components, MacroMedia Flash objects, Java applets, and scripting languages such as VBScript and JavaScript. Such objects can sometimes contain unwanted software. Although the anti-virus scanner used by the appliance detects many unwanted objects, you can provide extra security by choosing to block all such objects.

Web pages can also contain metadata, comments, and links (URLs) to other pages or web sites. If you are concerned that these areas might harbor unwanted software or undesirable content, you can choose to scan them too.

Scanner control (denial-of-service attacks)

Large or complex files such as compressed files or .ZIP files can take some time to scan. Such files can be used to attack your network, deliberately slowing its performance.

You can:

- Specify the amount of time that the appliance can spend scanning any file.
- Specify the action the appliance takes when that time limit is exceeded.

HTTP actions

The appliance can be configured to perform specific actions when a scanner triggers. For example, when the anti-virus scanner detects a virus, you can tell the appliance to replace the content with an HTML alert. The HTTP actions that the appliance can perform are described below:

- **Replace the content with an HTML alert**

If a scanner is triggered due to the content of a file, it replaces the content of that file with an HTML alert. The modified file is then sent to the recipient.

- **Allow Through**

The appliance will let the file through. Any scanner detections will be logged but not acted upon. For example, if you are expecting some large files that would normally trigger the denial-of-service limits, you could temporarily set the appliance to allow these files through, rather than replace them with an HTML alert.

Connection policies

The **Policy -> HTTP -> Advanced Policies -> Connection -> From Outside** menu options can be used to set up policies for HTTP connections initiated by hosts in your **Outside Networks** list.

The **Policy -> HTTP -> Advanced Policies -> Connection -> From Inside** menu options can be used to set up policies for HTTP connections initiated by hosts in your **Inside Networks** list.

You can configure the following HTTP connection policy:

- *Client (check for client).*
- *Time-outs.*

Client (check for client)

The appliance can be set to check that the HTTP client is still present.

You can specify how frequently the appliance should check for the presence of the HTTP client.

Time-outs

You can configure the connection and data time-outs.

The **Connection time-out** is the maximum number of seconds that the appliance will wait while trying to establish a connection with the remote web server.

The **Data time-out** is the maximum number of seconds that the appliance will wait for activity during the data transfer phase of the communication.

Protocol policies

The **Policy** -> **HTTP** -> **Advanced Policies** -> **Protocol** -> **From Outside** menu options can be used to set HTTP-specific policies that control the communication between the appliance and hosts listed in your **Outside Networks** list.

The **Policy** -> **HTTP** -> **Advanced Policies** -> **Protocol** -> **From Inside** menu options can be used to set HTTP-specific policies that control the communication between the appliance and hosts listed in your **Inside Networks** list.

You can configure the following HTTP protocol policies:

- [Client alert messages on page 174.](#)
- [Client download status messages on page 174.](#)
- [Denial-of-service prevention on page 174.](#)
- [Download status and data trickling on page 175.](#)
- [FTP over HTTP on page 176.](#)
- [Handoff host on page 176.](#)
- [Header blocking and modifications on page 177.](#)
- [Protocol details on page 177.](#)
- [Scanning on page 179.](#)
- [Streaming media on page 179.](#)
- [URL blocking and request permissions on page 180.](#)

Client alert messages

You can edit the messages that are sent by the appliance to alert users about specific events that have occurred.

Client download status messages

You can edit the message that is generated by the appliance and shown to users who download large files over an HTTP connection. The message is known as the *Client Download Status Message*.

Denial-of-service prevention

To prevent denial-of-service attacks the appliance can refuse an HTTP request if:

- The header size exceeds a pre-defined limit.
- The header line count exceeds a pre-defined limit.

Download status and data trickling

You can use one of the following options to help overcome the problems associated with downloading large files:

- Download status pages
- Data trickling
- Keep-alive headers

Download status pages

When a user tries to download a file over HTTP, the appliance contacts the relevant server and starts downloading the requested file to a storage area on the appliance. The appliance will download the whole file and scan it before sending the file to the user.

It can take a long time to download and scan large files, and the user might think that the connection has timed-out if they do not see any activity from the appliance.

To prevent this problem, you can configure the appliance to display a page that reassures the user that the download is still in progress. This page is known as the *Download Status Page*.

Data trickling

You can use data trickling to solve the problems that occur when the user downloads large files over HTTP.

You can configure the appliance to send the large file to the client as a series of much smaller data chunks. The large file can be trickled to the client in small amounts, before the whole file has been received from the server and scanned by the appliance. This is known as *Data Trickling*.

Caution

Data trickling can leave your network vulnerable to viruses and other potentially harmful software.

Keep-alive headers

It can take a long time to download and scan large files, and the client's software can time-out the connection.

To prevent this problem, you can configure the appliance to repeatedly send HTTP header lines to the client software to prevent that software timing out the connection. These headers are known as *Keep-alive headers*.

Refer to the *Product Guide* for more information about these options.

FTP over HTTP

The appliance allows FTP transfers over an HTTP connection. For example, a user could click on a link in a web browser page that launches a dialog box allowing them to download some software or a user guide from the Internet.

You can configure the appliance to contact the remote FTP server using a “passive FTP” or “active FTP” data connection over HTTP.

You can also specify which format will be used when the FTP directory information is returned. The format can be:

- Short form
- Full form
- Full merged form

NOTE

Typically, FTP over HTTP is only used when the client browser is configured to use the appliance as a proxy.

Handoff host

The appliance can be configured to use a handoff host for HTTP traffic.

A handoff host is used to divert all client requests to a specific server. This server is then responsible for handling the client requests. For example, a handoff host is used to divert requests to another proxy, such as a web cache ([Figure 12-1](#)).

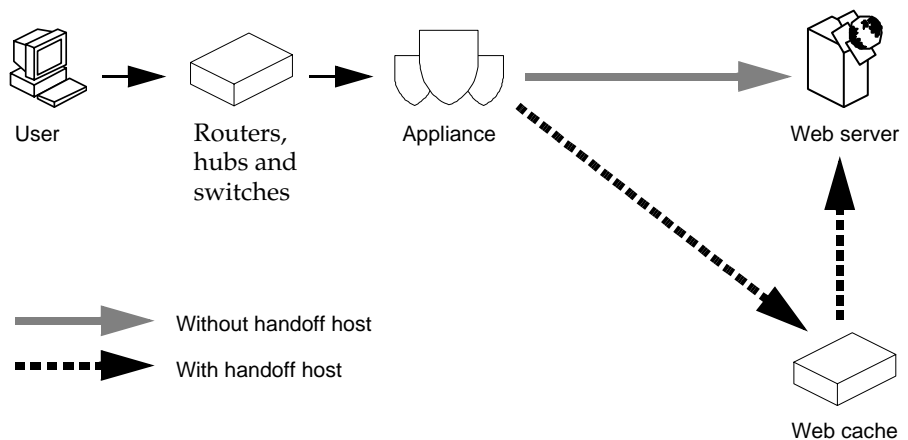


Figure 12-1. Using a handoff host

You can use this option to:

- Set up a handoff host for the appliance.
- Specify if the HTTP GET requests sent by the appliance should be formatted for use by a proxy or by a web server.

If a GET request is formatted for use by a proxy, it will contain the domain name of the server hosting the requested HTTP service and the name of the requested resource.

For example, if the user wants to view the default page (index.html) for the web site example.com, the GET request would look something like:

```
GET http://example.com/index.html HTTP/1.0
```

If the GET request is formatted for use by a web server, it will only contain the name of the requested resource (index.html in this example):

```
GET /index.html HTTP/1.0
```

Header blocking and modifications

The appliance can be configured to:

- Remove certain request and response headers. For example, the appliance can be set up to remove HTTP header records "Accept-Encoding" and "Accept-Ranges" records.
- Add VIA headers to HTTP request and responses.

Refer to the *Product Guide* for more information about these options.

Protocol details

You can configure the following options:

- NTLM failure pages.
- Non-compliant POST requests.
- Downgrade to HTTP 1.0.
- TRACE and OPTIONS requests.
- Server for internal information pages.

NTLM failure pages

Some clients and servers use the Microsoft Windows NT LAN Manager (NTLM) authentication protocol for the secure transmission of passwords.

In certain circumstances the NTLM authentication process will fail. For example, if a client using a web browser configured to operate in proxy mode tries to connect via the appliance to a server that requires NTLM authentication, the authentication will fail.

You can configure the appliance to display a page telling the client that their attempt to access the server failed because of a configuration issue.

NTLM will only work in Transparent Router or Transparent Bridge mode.

Non-compliant POST requests

You can configure how the appliance handles non-compliant HTTP POST requests.

A POST request is a request made by an HTTP client to send data to a server.

A non-compliant POST request can occur when the client (web browser) inserts non-compliant characters, such as line breaks, into the POST request.

Downgrade to HTTP 1.0

In order to communicate with older clients and servers, the appliance can be configured to use HTTP version 1.0 when dealing with HTTP requests or redirections.

TRACE and OPTIONS requests

You can configure how the appliance handles TRACE and OPTIONS requests. If you are unfamiliar with the TRACE and OPTIONS feature, refer to the HTTP RFC for more information.

Server for internal information pages

The appliance can be configured to include information about itself when issuing certain pages, such as Download Status Pages. This information includes the IP address of the appliance and the HTTP port number on which it receives HTTP traffic.

On large networks that have several appliances and use additional load balancing hardware, this information is used by the load balancing tool to make sure that HTTP responses are routed to the relevant appliance.

Caution

Do not change this option unless directed to do so by McAfee Technical Support or your network expert.

Scanning

The appliance can be set up to scan:

- Request headers.
- Request bodies.
- Response headers.
- Response bodies.

Streaming media

Streaming media is a technique for transferring data such that it can be processed as a steady and continuous stream. With streaming media, the user can view or listen to the data before the entire file has been transmitted.

The appliance cannot scan streaming media. To be able to scan a file for viruses, all of the data contained in that file must be available to the appliance. Streaming media is a continuous stream of data without a clear end-marker that the appliance can use to determine if the transmission is complete. The appliance does not know if it has received the whole file, and is therefore unable to complete the scan.

The appliance can be set up to allow streaming media to pass through it unscanned.

You can specify:

- If streaming media is allowed to pass through the appliance.
- Which types of data are considered to be streaming media by the appliance.
- From which type of servers the appliance will treat all data as streaming media.

Caution

All the data received from these servers types will be treated as streaming media and it will not be scanned by the appliance. This presents a serious security risk to your network. You should only configure this option at the request of McAfee Technical Support or your network expert.

Incoming streaming media must satisfy some security conditions before it passes to internal users. These security conditions are listed below:

- **Explicit Proxy mode**

- ◆ Streaming media on ports not scanned by the appliance will not be able to pass through the appliance. You should set up an alternative network route for this traffic.
- ◆ For streaming media arriving on port 80, the **Allow streaming media** checkbox must be selected. The media stream must also identify itself of being of a (MIME) type that is treated as streaming media by the appliance.

- **Transparent Router mode and Transparent Bridge mode**

- ◆ Streaming media on ports not scanned by the appliance, simply passes through the appliance to the users.
- ◆ For streaming media arriving on port 80, the **Allow streaming media** checkbox must be selected and the media stream must also identify itself as being of (MIME) type *audio/**.

Advanced configuration — If you are an expert user, you can also allow other types of streaming media, such as *video/**, to pass through the appliance. There is a Technical Note available to describe how to do this.

WARNING

We strongly discourage allowing streaming media of type *application/octet-stream* or *application/** to pass through the appliance as these MIME types are executable and present a significant security risk.

URL blocking and request permissions

The appliance can be set up to:

- Specify which HTTP verbs can be used.
- Specify which request schemes can be used.
- Specify which HTTP port numbers can be used.
- Specify which Secure Socket Layer (SSL) port numbers can be used.
- Deny access to web sites (URLs).
- Deny access to web sites that use special characters in their URL addresses.

HTTP verbs

You can specify which HTTP verbs can be used in the communication between the hosts and the appliance. Some examples of HTTP verbs include: GET, PUT, HEAD, OPTIONS, TRACE, POST, DELETE and CONNECT.

As soon as you have entered one or more HTTP verbs in the Permitted HTTP verbs list, by implication, all other HTTP verbs not in that list will be rejected. In a similar way, as soon as you enter one or more HTTP verbs in the Denied HTTP verbs list, by implication, you are permitting the use of all other HTTP verbs not in that list.

Request schemes

Uniform Resource Locators (URLs) include some text that defines which resource is being requested. For example:

Scheme	Requested resource
http	HyperText Transport Protocol
ftp	File Transfer Protocol

You can specify which schemes can be requested.

NOTE

As soon as you have entered one or more schemes in the **Permitted request schemes** list, by implication, all other schemes not in that list will be rejected.

In a similar way, as soon as you enter one or more schemes in the **Denied request schemes** list, by implication, you are permitting the use of all other schemes not in that list.

HTTP port numbers

For security reasons the appliance will only forward requests to certain port numbers. This prevents potential hackers "tunnelling" different protocols over an HTTP connection.

Use the **Permitted HTTP port numbers** option to set up port numbers for HTTP traffic that is not sent over the Secure Socket Layer (SSL). For HTTPS traffic that is sent over SSL, use the **Permitted SSL port numbers** option.

Caution

Changing port numbers can cause serious security and connectivity issues for you network. For this reason, we recommend that you only change these settings when requested to do so by McAfee Technical Support or by your network expert.

Secure Socket Layer (SSL) port numbers

For security reasons the appliance will only forward requests to certain port numbers. This prevents potential hackers "tunnelling" different protocols over an HTTP connection.

The port numbers that can be used depend on which HTTP verb is being used. Access using the CONNECT verb is more tightly restricted than other verbs. This is because once the CONNECT verb has been accepted, there is little restriction on the data that can be transferred.

Web browsers configured to operate in proxy mode use the CONNECT verb when trying to initiate an HTTPS connection running over the Secure Socket Layer (SSL).

For HTTPS traffic that is sent over SSL, use the Permitted SSL port numbers option. For all other HTTP traffic, use the Permitted HTTP port numbers option.

Blocking URLs

The appliance can block access to certain web sites. You can:

- Block specific URL addresses, by specifying those addresses in the **Denied URLs** list.

For example, to block access to any web site that contains the word "bad", you would add the word bad to the **Denied URLs** list.

- Use regular expressions to block all web sites that contain certain text patterns.

For example, to block access to any web sites with URLs that begin "http://" and contain the string "/images/", enter ^http://.* /images/* in the **Denied URLs (Via Regular Expressions)** list.

Denied (forbidden) URL characters

The use of certain characters with URL addresses can cause problems for some of the older browsers. You can set up the appliance to reject URL requests that contain these characters.

The appliance includes a File Transfer Protocol (FTP) proxy that is used for the secure transfer of files between computers. FTP is based on command-line instructions.

This section provides more information about FTP, and it includes the following topics:

- [General FTP configuration.](#)
- [Policy-based FTP configuration on page 184.](#)

General FTP configuration

The following feature, listed under the **Configure -> FTP** menu option, is not based on specific policies.

Connection settings (Advanced)

For TCP and UDP protocols, ports numbers are used to identify the ends of logical connections which carry specific services. Each service has an associated port number.

The default port number for FTP is port 21.

NOTE

We recommend that you do not change the port numbers unless you understand port assignments and the implications of changing the port number.

Intercept ports

The **Intercept ports** list only applies to appliances operating in Transparent Router mode or Transparent Bridge mode. You can specify the ports on which the appliance will intercept FTP traffic.

Listen ports

The **Listen ports** list only applies to appliances operating in Explicit Proxy mode. You can specify the ports on which the appliance will listen for FTP traffic.

Listeners, connections and memory

You can set up scanning resources on a protocol-by-protocol basis.

For each protocol, you need to set up the number of processes listening for the protocol-specific traffic. These processes are known as *Listeners*. Each listener can handle a number of connections. The number of scans that can be performed simultaneously for each protocol is equal to the number of listeners multiplied by the number of connections.

As there is a finite amount of resource available for scanning, there will always be a resource trade-off between:

- The number of listeners for that protocol.
- The number of connections handled by the listeners for the protocol.
- The amount of memory required for scanning that protocol.
- The scanning resources assigned to other protocols.

When you first use the appliance, some default values are in place. You can restore these settings at any time if the modified values are unsuitable.

Policy-based FTP configuration

This section briefly describes the purpose of each of the policies you can apply to FTP transfers.

Content policies

The **Policy -> FTP -> Content -> From Outside** menu options can be used to set up policies that tell the appliance how to handle the FTP requests and responses from hosts in your **Outside Networks** list.

The **Policy -> FTP -> Content -> From Inside** menu options can be used to set up policies that tell the appliance how to handle the FTP requests and responses from hosts in your **Inside Networks** list.

You can specify the actions that the appliance will perform in certain circumstances. See [FTP actions on page 185](#) for a list of possible actions.

You can configure the following FTP content policies:

- [Anti-virus settings on page 185](#).
- [Scanner control \(denial-of-service attacks\) on page 185](#).

Anti-virus settings

The appliance can be configured to detect viruses and other potentially harmful software. For more information about scanning for viruses and other potentially harmful software, see [Virus-scanning on page 163](#).

Scanner control (denial-of-service attacks)

Large or complex files such as compressed files or .ZIP files can take some time to scan. Such files can be used to attack your network, deliberately slowing its performance.

You can:

- Specify the amount of time that the appliance can spend scanning any file.
- Specify the action the appliance takes when that time limit is exceeded.

FTP actions

The appliance can be configured to perform specific actions when a scanner triggers. For example, when the anti-virus scanner detects a virus, you can tell the appliance to reject the original data. The FTP actions that the appliance can perform are described below:

- **Refuse the original data**

The file is rejected.

- **Allow through**

The appliance will let the file through. Any scanner detections will be logged but not acted upon.

Connection policies

The **Policy -> FTP -> Advanced Policies -> Connection -> From Outside** menu option can be used to set up a time-out policy for FTP connections initiated by hosts in your **Outside Networks** list.

The **Policy -> FTP -> Advanced Policies -> Connection -> From Inside** menu option can be used to set up a time-out policy for FTP connections initiated by hosts in your **Inside Networks** list.

You can configure the following FTP connection policy:

- [Time-outs on page 186](#).

Time-outs

You can configure the appliance so that it will close an FTP connection if an FTP command has not been received within a set time limit.

You can also specify how frequently the appliance should check the FTP connection.

Protocol policies

The **Policy -> FTP -> Advanced Policies -> Protocol -> From Outside** menu options can be used to set FTP-specific policies that control the communication between the appliance and hosts listed in your **Outside Networks** list.

The **Policy -> FTP -> Advanced Policies -> Protocol -> From Inside** menu options can be used to set FTP-specific policies that control the communication between the appliance and hosts listed in your **Inside Networks** list.

You can configure the following FTP protocol policies:

- *Data processing.*
- *Download status and data trickling.*
- *FTP handoff host.*
- *Upload status and data trickling.*

Data processing

This option allows you to configure:

- *Client messages.*
- *Keep-alive commands.*
- *Denied commands (advanced).*

Client messages

Client messages are generated by the appliance and shown to users who connect to the appliance using FTP. Their purpose is to prevent FTP downloads timing-out.

You can:

- Enable or disable the use of client messages.
- Change the help message that is displayed when a client sends an FTP request for help to the appliance.
- Change the welcome message that is displayed whenever a user connects to the appliance using FTP.

Keep-alive commands

When a user tries to download a file over FTP, the appliance contacts the relevant server and starts downloading the requested file to a storage area on the appliance. The appliance will typically download the whole file and scan it before sending the file to the user.

It can take a long time to transfer and scan large files, and the server software can time-out the connection if it does not see any activity from the appliance.

To prevent this problem, you can configure the appliance to repeatedly send an FTP command to the server software to prevent that software timing-out the connection. The command you choose is known as the keep-alive command.

You can specify which command to use as the keep-alive command and how frequently that command should be sent to keep the connection open.

Denied commands (advanced)

You can specify which FTP commands will not be accepted by the appliance.

Download status and data trickling

You can:

- Permit or deny the downloading of files over an FTP connection.
- Permit or deny the scanning of files downloaded over FTP.
- Specify if 8-bit data should be blocked in ASCII mode.
- Enable or disable the use of download status messages, and specify how frequently the download status messages should be sent, if enabled.
- Enable or disable the use of data trickling.

NOTE

You can use the download status message option or the data trickling option, but not both, as they are mutually exclusive.

Handling ASCII-mode FTP

The *File Transfer Protocol (FTP)* allows data to be passed between computers in two modes: binary and 8-bit ASCII (*American Standard Code for Information Interchange*). Binary is consistent across computer platforms, so its data can be scanned effectively. However, 8-bit ASCII can contain different character codes and formatting, depending on the computer systems in use, so viruses can be concealed easily within its data. For this reason, the appliance blocks this transfer mode by default, to maintain the security of your organization. You can configure the appliance to allow or block 8-bit data transfers.

Blocking 8-bit file transfers in ASCII mode prevents binary files being transferred in ASCII mode, but may also prevent legitimate text files in 8-bit character sets being transferred. If your users need to transfer text files in 8-bit character sets using FTP, the safest solution is to transfer the files in binary mode and convert the files to the appropriate local file format using utilities such as `unix2dos` or `dos2unix`.

NOTE

Some file transfer utilities use the 8-bit ASCII mode by default, so you must remember to change your utilities to binary mode, if the appliance is blocking the 8-bit ASCII mode.

Download status messages

When a user tries to download a file over FTP, the appliance contacts the relevant server and starts downloading the requested file to a storage area on the appliance. The appliance will download the whole file and scan it before sending the file to the user.

It can take a long time to download and scan large files, and the user might think that the connection has timed-out if they do not see any activity from the appliance.

You can configure the appliance to display status messages to let the user know that the download is still in progress. These messages are displayed on the FTP command line that the user sees when using command-line-driven FTP software.

Data trickling

You can configure the appliance to send the large file to the client as a series of much smaller data chunks. The large file can be *trickled* to the client in small amounts, before the whole file has been received from the server and scanned by the appliance. This is known as *Data Trickling*.

The advantages and disadvantages of using data trickling are described in the *Product Guide*.

Caution

Data trickling can leave your network vulnerable to viruses and other potentially harmful software.

FTP handoff host

An FTP handoff host is used to divert all client requests to a specific FTP proxy server. This server is then responsible for handling the client requests.

For example, if your firewall has an FTP proxy server, you can use the appliance's handoff host option to redirect FTP requests to the firewall.

Upload status and data trickling

You can:

- Permit or deny the uploading of files over an FTP connection.
- Permit or deny the scanning of files uploaded over FTP.
- Specify if 8-bit data should be blocked in ASCII mode.
- Enable or disable the use of upload status messages, and specify how frequently the upload status messages should be sent, if enabled.
- Enable or disable the use of data trickling.

NOTE

The handling of ASCII-mode FTP, is described in [Handling ASCII-mode FTP](#).

Upload status messages

When a user tries to upload a file over FTP, the appliance contacts the relevant server and start uploading the requested file to a storage area on the appliance. The appliance will upload the whole file and scan it before sending the file to the user.

It can take a long time to upload and scan large files, and the user might think that the connection has timed-out if they do not see any activity from the appliance.

You can configure the appliance to display status messages to let the user know that the download is still in progress. These messages are displayed on the FTP command line that the user sees when using command-line-driven FTP software.

Data trickling

When a user tries to upload a file over FTP, the appliance contacts the relevant server and starts uploading the file to a storage area on the appliance. The appliance will typically upload the whole file and scan it before sending the file to the server.

It can take a long time to upload and scan large files, and the user might think that the connection has timed-out if they do not see any activity from the appliance.

You can configure the appliance to upload the large file to the server as a series of much smaller data chunks. The large file can be *trickled* to the server in small amounts, before the whole file has been received from the client and scanned by the appliance. This is known as *Data Trickling*.

Caution

Data trickling can leave your network vulnerable to viruses and other potentially harmful software.

The advantages and disadvantages of using data trickling are described in the *Product Guide*.

This section provides more information about the POP3 (Post Office Protocol version 3) e-mail collection protocol. It includes the following topics:

- [General POP3 configuration on page 191.](#)
- [Policy-based POP3 configuration on page 193.](#)

General POP3 configuration

The following feature, listed under the **Configure** -> **POP3** menu option, is not based on specific policies.

Connection settings (Advanced)

For TCP and UDP protocols, ports numbers are used to identify the ends of logical connections which carry specific services. Each service has an associated port number.

The default port number for POP3 is 110.

NOTE

We recommend that you do not change the port numbers unless you understand port assignments and the implications of changing the port number.

Intercept ports

The **Intercept ports** list only applies to appliances operating in Transparent Router mode or Transparent Bridge mode. You can specify the ports on which the appliance will intercept POP3 traffic.

Listen ports

The **Listen ports** list only applies to appliances operating in Explicit Proxy mode. You can specify the ports on which the appliance will listen for POP3 traffic.

Dedicated Ports

The POP3 (Post Office Protocol version 3) e-mail collection protocol allows e-mail messages to be downloaded (pulled) from a mailbox on a remote server.

There are two modes of operation:

- **Generic connection** — allows connection to any POP3 server, but does not support APOP (Authenticated POP).

NOTE

As you configure the appliance with a port number for generic connections, your POP3 clients (software) do not need to specify that port number every time they make a generic POP3 connection through the appliance.

- **Dedicated connection** — allows connections to dedicated POP3 servers with APOP.

When a user makes a dedicated proxy connection through the appliance, the appliance uses a specified port to reach the POP3 server. You must specify a unique port number for each server. We recommend that you choose numbers between 1024–32767, because numbers up to 1024 are generally assigned to other protocols. The server will have a fully qualified domain name (FQDN), for example, pop3server.example.com.

NOTE

You can use the generic proxy port (110 by default) for a dedicated proxy connection. The dedicated connection overrides any generic connections.

Listeners, connections and memory

You can set up scanning resources on a protocol-by-protocol basis.

For each protocol, you need to set up the number of processes listening for the protocol-specific traffic. These processes are known as *Listeners*. Each listener can handle a number of connections. The number of scans that can be performed simultaneously for each protocol is equal to the number of listeners multiplied by the number of connections.

As there is a finite amount of resource available for scanning, there will always be a resource trade-off between:

- The number of listeners for that protocol.
- The number of connections handled by the listeners for the protocol.
- The amount of memory required for scanning that protocol.
- The scanning resources assigned to other protocols.

When you first use the appliance, some default values are in place. You can restore these settings at any time if the modified values are unsuitable.

Policy-based POP3 configuration

You can use policies to determine how the appliance handles e-mail in certain circumstances. For example, how the appliance handles an e-mail message could be determined by:

- Whether the e-mail message is from hosts in your Inside Networks Lists or Outside Networks list.
- Who sent the e-mail message.
- Who is going to receive the e-mail message.
- The content of the e-mail message.

Before you begin

This section describes some of the issues you should consider before setting up policies for POP3 e-mail messages.

Before setting up policies, you should:

- Familiarize yourself with the concepts described in this section.
- Spend some time thinking about how to organize users and computers into different policy groups. See [Policy Groups on page 108](#) for more information about policy groups.
- Consider which policies you want to assign to those policy groups.

There are some general guidelines that you should follow when setting up policies for e-mail messages:

- Set up the global policies to cover most scenarios. See [Global policies on page 104](#) for more information about global policies.
- Only set up a non-global policies if you need to create exceptions to the way that most e-mail messages are handled, as specified by the global policy.

See [Non-global policies on page 104](#) for more information about non-global policies.
- When the appliance is in Transparent Router or Transparent Bridge mode, the priority assigned to non-global content policies is important. See [Ordering non-global policies on page 105](#) for more information about prioritizing non-global policies.

Caution

Incorrect configuration of advanced policy settings can cause serious security and connectivity issues for your network. For this reason, we recommend that you do not change advanced settings unless instructed to do so by Technical Support or your network expert.

Content Policies

The **Policy -> POP3 -> Content -> From Outside** menu options can be used to set up policies that tell the appliance how to handle POP3 e-mail messages from hosts in your **Outside Networks** list.

The **Policy -> POP3 -> Content -> From Inside** menu options can be used to set up policies that tell the appliance how to handle POP3 e-mail messages from hosts in your **Inside Networks** list.

You can specify the actions that the appliance will perform in certain circumstances. See [POP3 Actions on page 199](#) for a list of possible actions.

You can configure the following SMTP content policies:

- [Alert settings on page 194.](#)
- [Anti-virus settings on page 195.](#)
- [Corrupt content on page 195.](#)
- [Encrypted content on page 195.](#)
- [Mail settings on page 195.](#)
- [Mail size filtering on page 196.](#)
- [Protected content on page 196.](#)
- [Scanner control \(denial-of-service attacks\) on page 197.](#)
- [Signed content \(digital signatures\) on page 198.](#)

Alert settings

The appliance will send an HTML message to clients when a specific event occurs. this known as an HTML alert.

You can:

- Change the text that appears at the start of the HTML alert, known as the *alert header*.
- Change the text that appears at the end of the HTML alert, known as the *alert footer*.

Anti-virus settings

The appliance can be configured to detect viruses and other potentially harmful software. For more information about scanning for viruses, see [Virus-scanning on page 163](#).

Corrupt content

If content is corrupt, the appliance might not be able to scan the file for viruses or banned content.

You can:

- Specify the action the appliance should take when corrupt content is detected.
- Specify which alert should be used, and if necessary, customize the alert text.

Encrypted content

If content is encrypted, it cannot be scanned. You can:

- Specify how encrypted content is handled by the appliance.
- Specify which alert message to use, and if necessary, customize the alert message text.

If you allow it through, it must be scanned after it is decrypted, and this typically occurs at the client computer.

Mail settings

You can specify how the appliance handles e-mail messages that use the MIME format. You can:

- Specify the action the appliance should take when a partial message is detected. A partial message is a message that has been divided into smaller parts for sending as several separate e-mail messages.
- Specify the action the appliance should take when a message contains a reference to an external resource and the scheme needed (usually FTP) to retrieve that resource. These messages are known as *external-body messages*.
- Specify which alert message should be used, and if necessary, customize the alert text.
- Add a prefix to the subject line of a message.
- Tell the appliance how to handle MIME messages that have corrupt header files.
- Tell the appliance where to position the alert and disclaimer attachments. The text can appear in the body text or be included as an attachment.

- Set the re-encoding options.
- Tell the appliance how to handle MIME header files that contain null characters.
- Tell the appliance how to handle inconsistent line endings.
- Specify the maximum number of MIME parts a message can have before the appliance considers it to be corrupt.
- Specify which MIME types should be treated as text attachments and which MIME types should be treated as binary attachments.

Mail size filtering

You can specify how the appliance handles e-mail messages that:

- Are larger than a pre-defined limit.
- Have attachments that are larger than a pre-defined limit.
- Have more attachments than are allowed.

An attachment, typically a graphic, a document, or a spreadsheet can greatly increase the size of a complete message — a typical memo of a few kilobytes can grow to many megabytes. Normally this flow of information is necessary for your organization to function, but problems arise when attachments are used excessively or when their use is abused.

For example, computer games are sometimes attached to e-mail messages. Each game typically consumes a few megabytes. Large audio or graphics files — whether for entertainment or business purposes — approach similar sizes. Popular items, when copied and forwarded many times over, can add a heavy load to your mail server. All users will suffer from the slower performance.

The appliance allows you to remove attachments from e-mail messages if they exceed a specified size or quantity. Discarded attachments can be replaced by a small text file, which informs the recipient that attachments were removed. You can also specify special actions against any e-mail message that exceeds a specified size overall.

Protected content

You can specify how the appliance handles e-mail messages that contain data that cannot be scanned because it is protected in some way. For example, it is password protected. You can:

- Specify the action the appliance should take when it detects protected content.
- Edit the replacement message.


Scanner control (denial-of-service attacks)

Large or complex files such as compressed files or .ZIP files can take some time to scan. Such files can be used to attack your network, deliberately slowing its performance. For these reasons, you can limit the size to which any file may be expanded and the depth of nesting.


You can also specify the amount of time that the appliance may spend scanning any file.

Depth of nesting in compressed files


To understand the effect of scanning to a depth of nesting, consider the next diagram. This shows a compressed file, which contains documents and a compressed file. That compressed file contains documents and another compressed file, and so on.



- A depth of 2 scans only the non-compressed files inside a compressed file (as shaded). The contents of any compressed files are not scanned.



- A depth of 3 scans the non-compressed files inside a compressed file, plus only the non-compressed files inside any compressed file that it contains (as shaded).



Signed content (digital signatures)

You can specify how the appliance handles signed e-mail messages. You can:

- Specify the action the appliance should take when an e-mail message is signed.
- Edit the replacement message.

Whenever information is sent electronically, it runs the risk of being accidentally or wilfully altered. To overcome this, some e-mail software uses a digital signature — the electronic form of a handwritten signature. A digital signature is extra information added to a sender's message, which identifies and authenticates the sender and the information in the message. It acts like a unique summary of the data. Typically, a long string of letters and numbers appears at the end of a received e-mail message. The e-mail software then re-examines the information in the sender's message, and creates a digital signature. If that signature is identical to the original, you can be sure that the data has not been altered.

While this method is useful most of the time, it can cause problems if the message violates your policy. For example, if the message contains a virus, bad content, or is too large, the appliance might clean or remove some part of the message. The original digital signature is now 'broken'. In other words, although the message is still valid (and usually readable), its signature is invalidated. Now the recipient cannot rely on the contents of the message at all because the contents might also have been altered in other ways.

You need to consider carefully how you handle signed messages. You can choose one of the following actions:

- **Refuse the original data and return a rejection code** — the appliance rejects the e-mail message and sends an SMTP 550 command to the mail server.
- **Accept and then drop the data** — the appliance accepts the e-mail message and discards it. It sends an SMTP 250 command to the mail server.
- **Allow changes to break the signed E-mail** — the appliance modifies the message, even if this breaks the signature. The modified e-mail message is sent to the recipient.
- **Do not allow changes to break the signed E-mail** — the appliance only performs actions that will not break the signed e-mail message signature. It then attempts to deliver the e-mail message to the original recipients.
- **Allow through** — some e-mail software might not accept any changes to the signed message, and therefore you cannot allow the appliance to alter the content. The danger here is that if you choose to allow all signed messages through, an undesirable item can escape detection if it is inside a signed message. If you allow all signed messages through, you need to be sure that the messages come from a trusted source, or that they will be scanned at a later stage. Any scanner detections will be logged but not acted upon.

In all cases, you can record the arrival of signed messages and notify an administrator.

POP3 Actions

The appliance can be configured to perform specific actions when a scanner triggers. For example, when the anti-virus scanner detects a virus, you can tell the appliance to attempt to clean the e-mail message.

The actions that are available will depend on the selected POP3 policy setting and on which scanner detected the issue.

These actions are described below:

- **Refuse the original data and return a rejection code**

The appliance rejects the e-mail message.

- **Accept and then drop the data**

The appliance accepts the e-mail message and discards it.

- **Do not allow changes to break the signed E-mail**

The appliance only performs actions that will not break the signed e-mail message signature. It then attempts to deliver the e-mail message to the original recipients.

- **Allow changes to break the signed e-mail**

If a scanner is triggered due to the content of an e-mail message, the appliance will modify the message, even if this breaks the signature. The modified e-mail message is sent to the recipient.

- **Replace the content with an HTML alert**

If a scanner is triggered due to the content of an e-mail message, it replaces the content with an HTML alert. The modified e-mail message is then sent to the recipient.

- **Remove the content**

You can configure the appliance to limit the number and size of attachments it will scan. If an e-mail message exceeds these limits, the appliance removes the excess content, scans the remaining e-mail message, and sends the modified e-mail message to the recipient.

- **Clean the content**

If the appliance detects a virus within the e-mail message, it will attempt to clean that virus. If the e-mail can be cleaned, the modified e-mail message is sent to the recipient.

You can also tell the appliance how to handle files that have zero bytes after they have been cleaned. The zero byte file options are accessed using the **Advanced** button in the anti-virus actions page.

You can:

- ◆ **Keep zero byte files** — the appliance will allow files that have been cleaned to have zero bytes.
 - ◆ **Remove zero byte files** — the appliance will remove any file that has zero bytes after it has been cleaned.
 - ◆ **Treat zero byte files as a failure to clean** — the appliance will treat the files as if it cannot clean them, and performs the action specified in If cleaning fails, take the following action.
- **Allow Through**

The appliance will let the e-mail message through. Any scanner detections will be logged but not acted upon. For example, you would select this action if you want to monitor the use of certain words in e-mail messages, without preventing their use.

Connection policies

The **Policy -> POP3-> Advanced Policies -> Connection -> From Outside** menu options can be used to set up policies for POP3 connections initiated by hosts in your **Outside Networks** list.

The **Policy -> POP3 -> Advanced Policies -> Connection -> From Inside** menu options can be used to set up policies for POP3 connections initiated by hosts in your **Inside Networks** list.

You can configure the following POP3 connection policy:

[Time-outs on page 201](#)

Time-outs

You can set up the following time-outs for POP3 connections:

- Inactivity time-out

The maximum number of seconds that the appliance will wait for a POP3 command before closing the connection.

- Client time-out

The maximum number of seconds that the appliance will wait for the client to complete the data transfer before closing a connection.

- Server time-out

The maximum number of seconds that the appliance will wait for the server to complete the data transfer before closing a connection.

- Server connection time-out

The maximum number of seconds that the appliance will wait for a response while trying to connect to an FTP server.

Protocol policies

The **Policy -> POP3 -> Advanced Policies -> Protocol -> From Outside** menu options can be used to set POP3-specific policies that control the communication between the appliance and hosts listed in your **Outside Networks** list.

The **Policy -> POP3 -> Advanced Policies -> Protocol -> From Inside** menu options can be used to set POP3-specific policies that control the communication between the appliance and hosts listed in your **Inside Networks** list.

You can configure the following POP3 protocol policies:

Protocol settings

You can configure:

- Server keep-alive commands
- Client keep-alive commands
- Address delimiter characters

Server keep-alive commands

The appliance can repeatedly send a POP3 command to prevent the connection between the appliance and the mail server timing-out.

Client keep-alive commands

The appliance can repeatedly send a POP3 command to prevent the connection between the appliance and the POP3 mail client timing-out.

Address delimiter characters

You can tell the appliance how to interpret the user's address when a generic proxy connection is made through the appliance. You can specifying which characters are used to identify each part of that address. By default the user name part of the address is separated from the host name by a hash (#), and the host information is separated from the port number by a colon (:). For example:

`<user name>#<host name>:<port number>`

You only need to change the delimiter characters if your POP3 provider uses different characters.

The appliance generates a large amount of information about its performance and status, and records this information in a log. This log can record activity over weeks or months. You can review the information contained in the log when necessary or convenient. You also have the option of configuring the appliance to respond to certain types of information by sending out an Alert (an alerting message) when specific events require an administrator or other person to be informed quickly.

You can select what types of data you want the log to record, the amount of detail, and whether to monitor specific events. The generated information can also be filtered. For example, you can restrict the information to show only the occurrences of one particular virus.

The appliance provides a selection of reports from which you can choose how you want information presented.

The information recorded by the log includes:

- Viruses detected and the action taken against them — whether the infected files have been cleaned, deleted or quarantined.
- Attempts to access web sites that are banned because they are considered inappropriate to business purposes.
- Spam (optional). Incidences of spam e-mail messages, including date, time, and sender.
- SMTP e-mail content rules. Content rules that have been triggered because of the banned content inside an SMTP e-mail message.
- Management events, such as failed login attempts and service failures.

NOTE

You must decide what level of detail you want the log to record, and configure this as soon as possible, using the **Configure -> Logging and Alerting** options.

Reports can only be based on data that has been saved in the log. If you do not set the log to record sufficient detail, you may find that your reports do not contain all the information that you need.

You can view and configure monitoring, reporting and alerting, from the **Monitor** section of the appliance window which provides you with a number of options. Alerting is configured from **Configure -> Logging and Alerting**.

Monitoring the appliance

The following options are available under the **Monitor** menu:

- **Status** — the **System Status** page
- **Performance**
- **Logs** — the **Logging and Alerting Reports** page
- **Charts**
- **Updates**
- **Resources**

Alerting is configured from **Configure** -> **Logging and Alerting**.

System Status

This page displays information about the status of the appliance, including the volume of traffic and detection rates.

The page shows the value of each parameter since the counters were last reset. These values are refreshed according to a schedule, but to see an immediate update click **Refresh**. The **Reset Counters** button resets the counters.

Clicking a link displays more information about the performance of that parameter.

The **Settings** button is used to define value ranges for status parameters, such as the protocol health parameters.

The following status information is shown:

Protocol Status

Shows protocol-specific scanning statistics collected since the appliance was last reset. It includes:

- Number of viruses detected for each protocol.
- Total number of viruses detected.
- Number of SMTP e-mail messages that have been deferred.
- Number of SMTP e-mail messages that have been quarantined because they contain a virus.
- Number of SMTP e-mail messages quarantined because they have triggered, for example, a content or spam rule.
- Amount of traffic received for each protocol.

- Number of SMTP e-mail messages containing spam that were detected by the SpamKiller for WebShield appliances software (optional).
- How many SMTP e-mail messages have been blocked by the SpamKiller for WebShield appliances software (optional).
- How many SMTP e-mail messages have been blocked using the Real Time Blocking (RBL) List.
- The number of e-mail messages deferred or quarantined.

Dashboard

Shows the:

- Health of each protocol. Health indicates that the protocol is able to accept a new connection.
- The rate at which the appliance swaps program pages between memory and the swap space on its hard disk. The swap rate is measured in pages per second. If the swap rate is persistently high, consider using additional appliances to share the workload.
- The load average, which is the number of processes awaiting execution on the CPU queue. The smaller the number of queued processes the better.
- Amount of storage space (as a percentage of the total storage space) that has been used for each type of partition on the hard disk, and the amount of space still available.
- CPU usage in megabytes and the number of free files.

General Status

Shows general status information. It includes:

- The time elapsed since the appliance was last restarted.
- The version number of the virus definition (DAT) files and anti-virus engine installed on the appliance.
- When the virus definition files and anti-virus engine files were last updated.
- The version number of the anti-spam rules and anti-spam engine (optional).
- The version number of the WebShield software.
- The user-friendly name of the appliance.
- The language used when generating reports and messages. It can be different from the language in which the interface is displayed.

Hardware Status

Shows information about the hardware configuration. It includes:

- The status of the Raid array, used for mirroring disks on some appliances with more than one disk.
- The physical address (MAC address), link status, speed, and connection type of the LAN1 and LAN2 ports.
- Transparent bridge status, if the appliance is configured as a bridge. Bridging can be either on or off. If bridging is turned on, the bridge will be shown as either forwarding or blocking traffic.

Load sharing status

- Shows information for up to ten load sharing servers. If load sharing is switched on, it includes the queue size on the controlling appliance, where the queue size is the number of connections that are still waiting to be virus-scanned. This also shows the name of the algorithm used to distribute the virus-scanning workload to the load sharing appliances. The default algorithm is Least Used.
- The name of each load sharing server.
- The connection status. If the status is **UP**, the appliance can contact the load sharing server.
- The number of concurrent connections that are supported by each load sharing appliance, and the number of connections still available on each of those appliances.

Performance

The **Performance Monitor** page allows you to monitor the performance of up to 4 appliances, each with its own tab.

The **New Chart** button can be used to select an additional appliance to monitor.

NOTE

You can only monitor additional appliances if they have the same logon name and password as the appliance you are currently working from. You can:

- Select the appliances you want to monitor.
- Set up a chart for each appliance.
- Save the settings for each chart, so that you do not have to re-enter those settings each time you want to view the chart.
- Select the counters you want to display for each of the appliances.

- Load the previously saved chart and counter settings.
- Save the current counter values to a file for later reference.

Clicking the **Configure Chart** button displays the **Chart Configuration** page. You can change:

- How often the chart will be updated.
- Which grids will be displayed.
- The name of the chart.
- The name used to label the vertical axis.
- The scale that should be used for the vertical axis.

Clicking the **Add** button displays the **Performance Counters** page. You can select:

- The performance object.
- The protocol you want to monitor.
- The performance counters you want to monitor.
- The color and scale of the line used to represent that counter.

For example, if you select a scale of 10 the appliance will multiply each value for that counter by ten. This is a useful feature if you want to view side-by-side counters that have very different value ranges.

If one counter always appears at the bottom of the chart and another counter always appears in the higher value ranges at the top of the chart, you might want to scale up the bottom counter. Both counters will then appear closer together, making it easier to monitor them side-by-side. The values of the scaled up counter will become relative values rather than absolute values.

The **Save** button saves the settings. The **Configure Chart -> Load** button can be used to display previously saved charts.

You can use the **Save Chart** button to save the chart details as a tab-separated-values (.TSV) file.

You can use the **Close Chart** button to close all the charts, except the original chart.

Logging and Alerting Reports

NOTE

To create and view reports, you must have first enabled logging using the **Logging and Alerting** option within the **Configure** menu.

To display the **Logging and Alerting Reports** page, select **Logs**. Use this page to select which charts and reports you want to generate. From here you can:

- Select the type of report you want to create.
- Specify which time period should be used when generating the report.
- Click the **Next** button to display the report.

Viewing the log

The log displays information according to the report type and time period you select. To view the log, select **Logs**, then choose the type of report you want to see. The following options are available: If there is more than one page of information you can scroll up and down.

Report type	Records
Resource and System	
User and User Interface	Events created through the use of the user interface.
Hardware and Resources	Events created by hardware configuration and use of resources.
Updates	When the appliance was updated.
Load Sharing and Bridging	Events created by load sharing and bridge configuration.
Protocol	
Conversation	Communication between the appliance and the clients/servers.
Processing	Events created by the processing of the protocols.
SMTP Transport	Events created while handling e-mail messages (SMTP).
Communications	
Network	Events created in network level communications.
Scanner	If any of the scanning engines cannot scan an object.
Detection	
Virus	Events created when anti-virus detection occurs.
Content	When banned content is detected.
Spam	When spam is detected.
URLs blocked	When web site addresses are blocked.
SMTP blocked	When e-mail messages are blocked.
Actions	Which actions have occurred as a result of a detection.

NOTE

Logs and reports relating to load sharing display an underscore in the host name when it is load sharing. For example, <name of appliance>_<name> indicates that the appliance was load sharing with the appliance called <name> at the time an event was recorded, whereas <name of appliance> without any underscores indicates that the appliance was not load sharing.


Displaying data graphically

The information which the appliance records in its log can be presented in pie charts or tables.

Showing the top 10

You can display the ten most frequently reported viruses or sources of spam-denied relay sources. These pie charts are known as *top 10* pie charts.

An example of a top 10 pie chart is shown in [Figure 15-1](#).



The top ten pie chart can show up to and including ten items. For example, if there are only two sources of spam, only two segments will appear in the pie chart.

Figure 15-1. Example of a pie chart

The top ten pie chart displays information for the selected report type and time period. To display a top ten pie chart, select the **Top Ten** chart format option.

Report using .TSV format

The information in the appliance's log can be exported as a text-based file in tab-separated values (.TSV) format.

A very simplified form of .TSV format is shown below:

Time	Event Date	Event Id	Event Text	Reason Text
12:55:07	12/12/2001	18100	"Found a virus"	Ripper
12:55:30	12/12/2001	18100	"Found a virus"	Love Letter

The .TSV format is ideal for importing the information into a spreadsheet or database tool, such as Microsoft Excel, Microsoft Access or Lotus 123. Using these tools, you can further manipulate the data and produce reports.

The .TSV file contains information for the selected report type and time period. To generate a .TSV file, select the **Report (TSV)** report format option.

Logging charts

Select **Monitor | Charts** to open this page, which you use to display logged information in charts.

To view these charts you must first have enabled logging, using the **Logging and Alerting** option in the **Configure** menu.

You can:

- Select the type of charts you want to display.
- Specify which time period should be used when generating the charts.
- Use **Maximum number of categories** to specify the maximum number of different scores that can be displayed in the pie chart.

The maximum number of categories should not be confused with the maximum number of top scorers, or the maximum number of segments a pie chart can have. Entries that have the same score are considered to be in the same category, although they can be displayed as separate segments in the pie chart. For example, if you set the maximum number of categories to 5 for event types, you could get the following results:

If there are 5 or less event types listed in the logs, each of these event types will have its own segment in the pie chart.

If there are 6 different event types - all with the same score - there is only one category, as they all have the same score. This means that all six event types will be displayed, each in its own segment.

If there are 10 event types - all with different scores - there are 10 categories. As only 5 categories can be displayed, only the top 5 scorers will be given their own segments. The rest will be grouped under a pie chart segment called **Others**. There will be 6 segments in the pie chart.

If there are 10 event types - some with the same score - the number of segments depends on the number of event types that can be grouped into the same category because they have the same score.

For example, if the scores are: a=10, b=10, c=10, d=9, e=8, f=8, g=7, h=6, i=5, j=4, there are 7 categories, as there are 7 different numbers (10, 9, 8, 7, 6, 5, 4).

If you only want to display a maximum of 5 categories, only event types with a score of 10, 9, 8, 7, or 6 will have their own segments in the pie chart. This means that event types a,b,c,d,e,f,g,h would have a segment each, and i and j would be grouped under a segment called **Others**. There will be 9 segments in all.

- Click the **Show Charts** button to display the charts.

Some reports are only available if you have McAfee SpamKiller for WebShield appliances installed and activated.

System Updates

Click **Monitor | Updates** to open this page which displays the anti-virus and anti-spam automatic update schedules. It shows:

- The name of each scheduled update.
- The current status of each scheduled task.
- When the last update took place.

System Resources

- Enable or disable the generation of disk usage events.
- Specify at which point the different disk usage events will trigger.
- Restart the selected protocols when there are no available connections.

Configuring Logging and Alerting

You can configure the appliance so that when an event occurs it will trigger a response. For example, you can set up the appliance so that when the "disk is nearing full" event occurs the appliance sends an e-mail message to your network administrator.

Click **Configure -> Logging and Alerting** to open the **Logging and Alerting Configuration** page. Use this to:

- Specify which events will be recorded by the appliance.
- Specify what type of response is required when specific events occur.
- Configure the way that the notification is presented. The notification can be sent as e-mail messages, SNMP traps, ePolicy Orchestrator alerts, and Syslog entries.

Refer to the *Product Guide* for further information about the logging and alerting parameters.

The appliance generates a large number of alerts arising from events such as:

- Detection of a virus.
- Detection of a banned word or phrase.
- Detection of a spam e-mail message.
- A failed attempt to log on.
- Resources becoming exhausted.

Distributing the alerts

The appliance creates a large amount of information. You can choose the method of distribution as follows:

- **E-mail messages**

This information is sent as e-mail messages to any number of recipients.

- **SNMP traps**

This information is sent as alerts to an SNMP (Simple Network Management Protocol) manager. The .MIB file supplied on the appliance tells the SNMP manager how to interpret the data in the traps.

- **McAfee ePolicy Orchestrator events**

An agent program installed on the appliance generates a file that contains system, virus detection, and anti-spam detection information. This information is passed to the ePolicy Orchestrator server for monitoring and reporting purposes.

Specifying the detail

You can refine your choice of events to record, as follows:

- Specify the severity of protocol and communication events you want to forward — such as a detection or a critical warning.
- Specify the type of detection events.
- Specify a code for the event, its event ID — this enables you to specify precisely what you want to monitor.

Monitoring critical levels

The appliance monitors several of its e-mail resources to ensure that they do not become exhausted, for example:

- Free space remaining in inbound and outbound virus quarantine areas.
- Free space remaining in the inbound and outbound content quarantine area.
- Size of the deferred mail queue.

You can configure the appliance to alert an administrator when the resources are close to exhaustion. To do this, you enable the e-mail alerting, and configure it to forward information about all high severity communication events.

Handling e-mail problems

The appliance can assign an administrator (or *postmaster*) to handle any queries from senders regarding e-mail messages that were returned because of a virus or their content. To do this click **Configure** -> **SMTP** and enter the postmaster's e-mail address in the relevant dialog box.

The postmaster must be someone who will read mail regularly. Besides using the names of single users, you can also use names of distribution lists, because the appliance does not distinguish between them. If you are using an e-mail tool such as Microsoft Outlook or Lotus Notes, users can set up rules to forward e-mail messages when they are out of the office.

This section describes load sharing, and includes the following topics:

- [Load sharing basic concepts.](#)
- [Configuration scenarios on page 216.](#)
- [Configuring load sharing on page 217.](#)
- [Load sharing examples on page 219.](#)
- [Viewing the load sharing status on page 221.](#)

Load sharing basic concepts

An appliance can be set up so that when it receives traffic from supported protocols, it can off-load some or all of its scanning workload to other appliances. This ability to share the scanning workload between appliances is called *load sharing*.

An appliance that off-loads its scanning workload is known as a *controlling appliance*.

Appliances that receive scanning work from a controlling appliance are known as *load sharing appliances*.

A controlling appliance can share its scanning workload with many load sharing appliances. A load sharing appliance can accept scanning work from more than one controlling appliance.

Load sharing can offer the advantage of handling larger numbers of concurrent connections, but at the expense of maximum throughput.

Configuration scenarios

There are several ways you can set up load sharing:

- Scenario 1 — single appliance with one scanner per connection.
- Scenario 2 — single appliance with multiple scanners per connection.
- Scenario 3 — multiple appliances with the controlling appliance off-loading some scanning workload to other appliances.
- Scenario 4 — multiple appliances with the controlling appliance off-loading all of its scanning workload to other appliances.

NOTE

Load sharing does not use the Content Vectoring Protocol (CVP).

Scenario 1 — single appliance in non-sharing mode

By default, the appliance still performs all the scanning itself, but each connection that needs to be scanned has a dedicated scanner. There is a one-to-one relationship between the connection and the scanners. The appliance can be thought of as being in a non-sharing mode, which reduces the number of concurrent connections that it can support.

Scenario 2 — single appliance in sharing mode

The appliance can be set up to effectively load share with itself. It does not off-load any of its scanning workload to other appliances. When the appliance receives traffic to be scanned it uses its own internal scanners to scan that traffic. The appliance can be thought of as sharing the scanning workload for each connection across a number of internal scanners. There is a one-to-many relationship between the connection and the scanners.

Scenario 3 — controlling appliance off-loads some workload

If you have more than one appliance, you can set up one of the appliances as a *controlling appliance* and the others as *load sharing appliances*. In this scenario the controlling appliance continues to perform some scanning itself, but also off-loads some of the scanning to other appliances.

The controlling appliance's scanning settings are used to control scanning on the load sharing appliances. That is, the scanning settings on the controlling appliance override any scanning settings on a load sharing appliance whenever that appliance receives traffic to scan from the controlling appliance.

Off-loading some of the scanning to other appliances allows the controlling appliance to dedicate more of its internal resources to managing incoming traffic.

Scenario 4 — controlling appliance off-loads all workload

This scenario is similar to scenario 3, except that the controlling appliance off-loads all of its scanning workload to the load sharing appliances. By off-loading all of its scanning workload, the controlling appliance can dedicate even more of its internal resources to managing incoming traffic.

NOTE

If there are five or more load sharing appliances, we recommend that the controlling appliance off-loads all of its scanning workload to those appliances to free up its own resources for traffic management.

Configuring load sharing

This section describes how to configure the appliances involved in load sharing.

NOTE

We recommend that you use appliances of the same type when load sharing.

Like all appliances, the appliances involved in load sharing, must be able to access the DAT and anti-virus engine update servers to ensure that they have access to the latest anti-virus software.

If load sharing anti-spam scanning, all of the appliances must have access to the anti-spam engine and anti-spam rules package server.

Configuring the controlling appliance

You must specify which appliances will share the scanning workload:

- 1 *On the controlling appliance*, from the **Network** menu select the **Load Sharing** option.
- 2 In **Requests**, select **Make**.
- 3 Within **Servers list** select **Add**.
- 4 Enter a unique name that can be used to identify the load sharing appliance.
- 5 Enter the load sharing appliance's IP address, or its fully qualified host name.

NOTE

If you enter a host name it must only resolve to a single IP address in DNS.

- 6 Click **OK**.
- 7 Repeat this procedure for each load sharing appliance that this controlling appliance will use.

- 8 If necessary, use the **Move Up** and **Move Down** buttons to change the order in which the controlling appliance will contact the load sharing appliances.

Configuring the load sharing appliance

To configure the load sharing appliance:

- 1 Enable load sharing on the load sharing appliance. For security reasons, the ability to load share with other appliances is disabled by default. To enable load sharing:
 - a On the load sharing appliance, from the **Network** menu, select **Load Sharing**.
 - b Select **Accept**.
- 2 You must list the controlling appliance in the load sharing appliance's *Inside Networks* list. To list the controlling appliance:
 - a On the load sharing appliance, from the **Network** menu, select **Settings**.
 - b Click the **Inside Networks** tab.
 - c If the controlling appliance is not already listed, add it by selecting **Add**.
 - d Enter the IP address, subnet mask, or host name of the controlling appliance.
- 3 If necessary, you can change the total number of scans that can be performed simultaneously by this load sharing appliance. Although in most cases the default values should be sufficient and should not be changed.

You can change the total number of scans by changing the **Listeners** and **Connections** settings on the **Load Sharing Configuration** page.

The total number of scans is equal to the number of listeners multiplied by the number of connections, and is the total number of scans for all protocols. For example, if SMTP and HTTP scanning is enabled on the controlling appliance, and the load sharing appliance is set up with 4 listeners and 35 connections, the maximum number of SMTP/HTTP scans that can be performed simultaneously by that load sharing appliance is 140. The total number of scans also depends on the amount of memory that is available.

- 4 Repeat steps 1 and 3 for each load sharing appliance, or use the **System -> Manage Appliances -> Manage a group of appliances** option to set up the other load sharing appliances. For more information on managing groups of appliances, see the *Product Guide* for more information about managing a group of appliances.

Load sharing examples

This section tells you how to configure the appliances for each scenario described in [Configuration scenarios on page 216](#).

Example 1 — single appliance in non-sharing mode

By default, the appliance is in non-sharing mode. In non-sharing mode each connection will have a dedicated scanner.

If the appliance has been previously configured to operate in sharing mode, it can be returned to its default configuration by deselecting the **Make** option in the **Load Sharing Configuration** page.

Example 2 — single appliance in sharing mode

The appliance is load sharing with itself, as indicated by the entry **localhost 127.0.0.1** in its **Servers list**.

The appliance can be configured to share with itself by entering **localhost 127.0.0.1** in its **Servers list**, as described below:

- 1 From the **Network** menu, select **Load Sharing**.
- 2 In **Requests**, select **Make**.
- 3 In the **Servers list** section select **Add**.
- 4 Enter **localhost** as the name of the server.
- 5 Enter **127.0.0.1** as its IP address. You **must not** enter any other IP address for the **localhost** entry.
- 6 Click **OK**. The appliance will restart when the changes are applied.

Example 3 — controlling appliance off-loads some workload

To configure the appliances so that the controlling appliance off-loads some of its scanning workload to other appliances:

- 1 Configure the appliances as described in [Configuring load sharing on page 217](#).
- 2 Make sure that the controlling appliance is listed in its own **Servers list**. The appliance should be represented by the entry **localhost 127.0.0.1**. If the controlling appliance is not listed, add it as described below:
 - a *On the controlling appliance*, from the **Network** menu, select **Load Sharing**.
 - b In **Requests**, select **Make**.
 - c In the **Servers list** section select **Add**.
 - d Enter **localhost** as the name of the server.
 - e Enter **127.0.0.1** as its IP address. You *must not* enter any other IP address for the **localhost** entry.
 - f Click **OK**. The appliance will then restart when the changes are applied.

Example 4 — controlling appliance off-loads all workload

To configure the appliances so that the controlling appliance off-loads all of its scanning workload to the load sharing appliances:

- 1 Configure the appliances as described in [Configuring load sharing on page 217](#).
- 2 Stop the appliance load sharing with itself:
 - a *On the controlling appliance*, select **Load Sharing** from the **System** menu.
 - b In the **Servers list** section select **localhost 127.0.0.1**.
 - c Click **Delete**. The appliance will restart when the changes are applied.

Viewing the load sharing status

To view the load sharing status, from the **Monitor** menu, select **Status**.

The **System Status** pages displays the general status of the appliance, and the load sharing status of up to nine other appliances involved in load sharing. The following load sharing information is displayed:

- The queue size on the controlling appliance. The queue size is the number of connections on the controlling appliances that are still waiting to be scanned.
- The algorithm used by the controlling appliance to distribute work to the load sharing appliances. By default the controlling appliance uses the **Least used** algorithm.
- The number of concurrent connections that are supported by each load sharing appliance.
- The number of connections still available on each of the load sharing appliances.

This chapter describes the changes that you might need to make to the appliance to ensure correct configuration and operation.

Changing the password

By default, the password used to access the appliance is:

webshieldchangeme

For security reasons you should change this password.

Use the **System -> Manage Appliances -> Manage this appliance -> Set the password** option to change the password.

Turning off or rebooting the appliance

The appliance can be rebooted remotely, or turned off completely. To prevent tampering and the accidental rebooting or stopping of the appliance, these features only work if a password is given.

Use the **System -> Manage Appliances -> Manage this appliance -> Reboot the appliance** option to reboot the appliance.

Use the **System -> Manage Appliances -> Manage this appliance -> Stop the appliance** option to turn off the appliance.

Setting the system date and time

You can set the system date and time used by the appliance for reporting and other purposes.

Use the **System -> Manage Appliances -> Manage this appliance -> Set the date and time** options to change the system date and time zone used by the appliance.

Use the **System -> Manage Appliances -> Manage this appliance -> Time zone** option to specify which time zone the appliance should use.

Setting the operational language

You can change the language used for internal reports and error messages.

Use the **System -> Manage Appliances -> Manage this appliance -> Set the operational language -> Operational Language** option to change the language used for internal reporting and error messages.

Installing service packs and HotFixes

McAfee occasionally release Service Packs and HotFixes to enhance the appliance software. The new software can be accessed from a web site, or from files on another computer.

Use the **System -> Manage Components -> Component Install** option to install service packs and HotFixes.

Evaluating SpamKiller

You can evaluate the *McAfee SpamKiller for WebShield appliances* optional software for up to 30 days.

Use the **System -> Manage Components -> Component Activation -> Start Evaluation** option to start the evaluation.

If you stop the evaluation, or the evaluation period is over, the button name will change to say **Complete**, and you will be unable to restart the evaluation.

Activating SpamKiller

To activate the *McAfee SpamKiller for WebShield appliances* optional software, you must first acquire the activation CD.

Information about obtaining the activation CD can be found on our web site. For contact information, see [Contacting McAfee Security & Network Associates on page 18](#).

When you have acquired the activation CD, you can use the **System -> Manage Components -> Component Activation -> Browse** option to locate it and begin the activation process.

Enabling the ePolicy Orchestrator Agent

ePolicy Orchestrator enables you to monitor virus activity and distribute anti-virus software from a single point.

To communicate with ePolicy Orchestrator you must enable the ePolicy Orchestrator agent on the appliance. The appliance can then send status information to the ePolicy Orchestrator server, so that the appliance can be monitored remotely with ePolicy Orchestrator.

Use the **System -> Manage Components -> Component Activation -> Enable ePolicy Orchestrator Agent** option to enable the agent.

Refer to the *ePolicy Orchestrator Configuration Guide* that accompanies your appliance for more information about configuring ePolicy Orchestrator.

Uninstalling the ePolicy Orchestrator Agent

ePolicy Orchestrator Agent can be uninstalled.

Use the **System -> Manage Components -> Component Uninstall** option to uninstall the agent.

Refer to the *ePolicy Orchestrator Configuration Guide* that accompanies your appliance for more information about configuring ePolicy Orchestrator.

Saving the logs

The appliance maintains a number of logs that record changes. For example:

- **System logs** — the appliance stores information about system events, such as failed log on attempts, in its system logs. You cannot disable the system logs.
- **Appliance logs** — the appliance maintains other logs that record appliance events, such as virus detections. These other logs will be referred to as *appliance logs*. You cannot disable the appliance logs.

You can choose which types of events you want to capture in the appliance logs. You can use the **Configure -> Logging and Alerting** option to set up the appliance logs.

NOTE

Charts and reports can only show events that have been logged. See [Monitoring the appliance on page 203](#) for more information about charts and reports.

Use the **System -> Backup and Restore -> Save logs** option to save the logs to a computer.

Off-box logging with Syslog

You can send logging information to an *off-box Syslog* file.

Use the **Configure -> Logging and Alerting -> Channel Filters -> Syslog** options to specify which event types should be sent to the Syslog.

Use the **Configure -> Logging and Alerting -> Channel Settings -> Syslog** options to:

- Enable or disable logging to the Syslog.
- Specify which computer is being used for off-box logging.

Saving system configuration

You can safely store details about the appliance's configuration offline, and restore that information later.

Use the **System -> Backup and Restore -> Save Configuration** option to save the logs to a computer.

System configuration settings are saved to a .ZIP file, which contains mainly .XML files.

Restoring system settings

You can restore previously saved settings to the appliance. You might do this for the following reasons:

- You have upgraded the appliance's software and want to use the same settings as before.
- You have restored the appliance's software (because of a problem) and want to use the same settings as before.
- You have another appliance and you want to copy its settings so that they are similar. A better way of copying configuration from one appliance to another is described in [Copying configuration on page 227](#).

You can use the **System -> Backup and Restore -> Restore -> Restore Configuration** option to restore the configuration to the appliance.

NOTE

If you use the same system settings file for more than one appliance, they will all have the same appliance names (host names) and IP addresses. Once you have selected the **Restore Configuration** button, you must change the appliance name and IP addresses so that they are unique, before clicking **Apply All Changes**.

If you have installed new software on the appliance, or turned on the appliance and logged on for the first time, you will see a special version of the appliance's home page. The **System -> Backup and Restore** option is also available on that page.

NOTE

The user name and password are not saved from the previous configuration. You need to log on to the appliance using its user name and default password, then change the password.

The user name for all appliances is *webshield*, and the default password is *webshieldchangeme*.

Restoring default settings

You can use the **System -> Backup and Restore -> Restore -> Restore Defaults** option to restore the appliance to its default state.

To view a list of default settings, from the **Links Bar** select the **Resource** tab. Click on the **Default Settings** link.

Accessing the MIB definition file

To view the MIB definition file, from the **Links Bar** select the **Resource** tab. Click on the **MIB File** link.

Copying configuration

The configuration of one appliance can be applied to a group of appliances.

Not all configuration parameters are pushed to the other appliances. Appliance specific parameters, such as its IP address, will not be distributed around the group.

Use the **System -> Manage Appliances -> Manage a group of appliances** option to copy the configuration of an appliance to a group of appliances.

NOTE

The appliances in the group must have the same logon password as the appliance whose configuration will be distributed to the group.

Removing old files

To prevent the appliance running out of resources, you should regularly remove unwanted:

- E-mail messages that are stored in the quarantine area.
- E-mail messages that have been deferred.
- Information in the System logs.

Quarantine maintenance

You can configure the appliance to automatically remove unwanted e-mail messages at regular intervals. By default quarantined e-mail messages are removed after 14 days.

Use the **E-Mail -> Quarantine Maintenance** option to set quarantine maintenance.

Deferred e-mail maintenance

Deferred e-mail messages are sent to the *deferred* area on the appliance. You can use the **Delete** or **Delete All** buttons available under the **E-Mail -> Deferred** option to remove unwanted messages.

Restricting the number of log files

As new logs are generated, the number of logs stored on the appliance increases. To manage the amount of resource taken up by the log files, you can restrict the number of days the logs will be kept before being deleted.

Use the **Configure -> Logging and Alerting -> Channel Settings -> XML -> Keep logs** option to specify how long the logs will be kept.

Word Separators

A

When you create content scanning rules you need to know how the appliance treats the different word separators that it finds when it scans e-mail headers, body content, and attachments.

This appendix lists the Unicode and ASCII characters that the appliance recognizes as word separators when it scans e-mail messages.

NOTE

When the text being scanned is in ASCII format, only the Latin characters with decimal values up to and including 127 are used.

This appendix does not show the actual characters. The actual characters can be viewed at the Unicode Consortium web site.

Characters are grouped into charts according to their hexadecimal range: with each range typically being a regional character set, such as Latin, or a functional grouping, such as Symbols. To view the charts, go to:

<http://www.unicode.org/charts>

You might find it easier to use this appendix to find the name of the character that you want to view, then use the character index at:

<http://www.unicode.org/charts/charindex.html>

The character index lists the character names in alphabetical order, and provides links to the relevant chart.

For each character, this appendix shows the:

- Character name.
- Hexadecimal code — used to generate the character on a computer.
- Decimal code — used to generate the character on a computer.
- Type of character — see *Type key on page 230*.

The appliance recognizes punctuation, separators, and math symbols as word separator character types within content rules.

Type key

Character list on page 231 uses the following acronyms for character *Type*:

Pc	Punctuation, Connect
Pd	Punctuation, Dash
Ps	Punctuation, Open
Pe	Punctuation, Close
Pi	Punctuation, Initial quote
Pf	Punctuation, Final quote
Po	Punctuation, Other
Zs	Separator, Space
Zl	Separator, Line
Zp	Separator, Paragraph
Sm	Math Symbol

Character list

Character Name	Hexadecimal Code	Decimal Code	Type
HORIZONTAL TABULATION	0x0009	09	
LINE FEED	0x000a	10	
CARRIAGE RETURN	0x000d	13	
SPACE	0x0020	32	Zs
EXCLAMATION MARK	0x0021	33	Po
QUOTATION MARK	0x0022	34	Po
NUMBER SIGN	0x0023	35	Po
PERCENT SIGN	0x0025	37	Po
AMPERSAND	0x0026	38	Po
APOSTROPHE	0x0027	39	Po
LEFT PARENTHESIS	0x0028	40	Ps
RIGHT PARENTHESIS	0x0029	41	Pe
ASTERISK	0x002a	42	Po
PLUS SIGN	0x002b	43	Sm
COMMA	0x002c	44	Po
HYPHEN-MINUS	0x002d	45	Pd
FULL STOP	0x002e	46	Po
SOLIDUS	0x002f	47	Po
COLON	0x003a	58	Po
SEMICOLON	0x003b	59	Po
LESS-THAN SIGN	0x003c	60	Sm
EQUALS SIGN	0x003d	61	Sm
GREATER-THAN SIGN	0x003e	62	Sm
QUESTION MARK	0x003f	63	Po
COMMERCIAL AT	0x0040	64	Po
LEFT SQUARE BRACKET	0x005b	91	Ps
REVERSE SOLIDUS	0x005c	92	Po
RIGHT SQUARE BRACKET	0x005d	93	Pe

Character Name	Hexadecimal Code	Decimal Code	Type
LOW LINE	0x005f	95	Pc
LEFT CURLY BRACKET	0x007b	123	Ps
VERTICAL LINE	0x007c	124	Sm
RIGHT CURLY BRACKET	0x007d	125	Pe
TILDE	0x007e	126	Sm
NO-BREAK SPACE	0x00a0	160	Zs
INVERTED EXCLAMATION MARK	0x00a1	161	Po
LEFT-POINTING DOUBLE ANGLE QUOTATION MARK	0x00ab	171	Pi
NOT SIGN	0x00ac	172	Sm
SOFT HYPHEN	0x00ad	173	Pd
PLUS-MINUS SIGN	0x00b1	177	Sm
MIDDLE DOT	0x00b7	183	Po
RIGHT-POINTING DOUBLE ANGLE QUOTATION MARK	0x00bb	187	Pf
INVERTED QUESTION MARK	0x00bf	191	Po
MULTIPLICATION SIGN	0x00d7	215	Sm
DIVISION SIGN	0x00f7	247	Sm
GREEK QUESTION MARK	0x037e	894	Po
GREEK ANO TELEIA	0x0387	903	Po
ARMENIAN APOSTROPHE	0x055a	1370	Po
ARMENIAN EMPHASIS MARK	0x055b	1371	Po
ARMENIAN EXCLAMATION MARK	0x055c	1372	Po
ARMENIAN COMMA	0x055d	1373	Po
ARMENIAN QUESTION MARK	0x055e	1374	Po
ARMENIAN ABBREVIATION MARK	0x055f	1375	Po
ARMENIAN FULL STOP	0x0589	1417	Po
ARMENIAN HYPHEN	0x058a	1418	Pd
HEBREW PUNCTUATION MAQAF	0x05be	1470	Po
HEBREW PUNCTUATION PASEQ	0x05c0	1472	Po
HEBREW PUNCTUATION SOF PASUQ	0x05c3	1475	Po
HEBREW PUNCTUATION GERESH	0x05f3	1523	Po

Character Name	Hexadecimal Code	Decimal Code	Type
HEBREW PUNCTUATION GERSHAYIM	0x05f4	1524	Po
ARABIC COMMA	0x060c	1548	Po
ARABIC SEMICOLON	0x061b	1563	Po
ARABIC QUESTION MARK	0x061f	1567	Po
ARABIC PERCENT SIGN	0x066a	1642	Po
ARABIC DECIMAL SEPARATOR	0x066b	1643	Po
ARABIC THOUSANDS SEPARATOR	0x066c	1644	Po
ARABIC FIVE POINTED STAR	0x066d	1645	Po
ARABIC FULL STOP	0x06d4	1748	Po
SYRIAC END OF PARAGRAPH	0x0700	1792	Po
SYRIAC SUPRALINEAR FULL STOP	0x0701	1793	Po
SYRIAC SUBLINEAR FULL STOP	0x0702	1794	Po
SYRIAC SUPRALINEAR COLON	0x0703	1795	Po
SYRIAC SUBLINEAR COLON	0x0704	1796	Po
SYRIAC HORIZONTAL COLON	0x0705	1797	Po
SYRIAC COLON SKEWED LEFT	0x0706	1798	Po
SYRIAC COLON SKEWED RIGHT	0x0707	1799	Po
SYRIAC SUPRALINEAR COLON SKEWED LEFT	0x0708	1800	Po
SYRIAC SUBLINEAR COLON SKEWED RIGHT	0x0709	1801	Po
SYRIAC CONTRACTION	0x070a	1802	Po
SYRIAC HARKLEAN OBELUS	0x070b	1803	Po
SYRIAC HARKLEAN METOBELUS	0x070c	1804	Po
SYRIAC HARKLEAN ASTERISCUS	0x070d	1805	Po
DEVANAGARI DANDA	0x0964	2404	Po
DEVANAGARI DOUBLE DANDA	0x0965	2405	Po
DEVANAGARI ABBREVIATION SIGN	0x0970	2416	Po
SINHALA PUNCTUATION KUNDDALIYA	0x0df4	3572	Po
THAI CHARACTER FONGMAN	0x0e4f	3663	Po
THAI CHARACTER ANGKHANKHU	0x0e5a	3674	Po
THAI CHARACTER KHOMUT	0x0e5b	3675	Po

Character Name	Hexadecimal Code	Decimal Code	Type
TIBETAN MARK INITIAL YIG MGO MDUN MA	0x0f04	3844	Po
TIBETAN MARK CLOSING YIG MGO SGAB MA	0x0f05	3845	Po
TIBETAN MARK CARET YIG MGO PHUR SHAD MA	0x0f06	3846	Po
TIBETAN MARK YIG MGO TSHEG SHAD MA	0x0f07	3847	Po
TIBETAN MARK SBRUL SHAD	0x0f08	3848	Po
TIBETAN MARK BSKUR YIG MGO	0x0f09	3849	Po
TIBETAN MARK BKA- SHOG YIG MGO	0x0f0a	3850	Po
TIBETAN MARK INTERSYLLABIC TSHEG	0x0f0b	3851	Po
TIBETAN MARK DELIMITER TSHEG BSTAR	0x0f0c	3852	Po
TIBETAN MARK SHAD	0x0f0d	3853	Po
TIBETAN MARK NYIS SHAD	0x0f0e	3854	Po
TIBETAN MARK TSHEG SHAD	0x0f0f	3855	Po
TIBETAN MARK NYIS TSHEG SHAD	0x0f10	3856	Po
TIBETAN MARK RIN CHEN SPUNGS SHAD	0x0f11	3857	Po
TIBETAN MARK RGYA GRAM SHAD	0x0f12	3858	Po
TIBETAN MARK GUG RTAGS GYON	0x0f3a	3898	Ps
TIBETAN MARK GUG RTAGS GYAS	0x0f3b	3899	Pe
TIBETAN MARK ANG KHANG GYON	0x0f3c	3900	Ps
TIBETAN MARK ANG KHANG GYAS	0x0f3d	3901	Pe
TIBETAN MARK PALUTA	0x0f85	3973	Po
MYANMAR SIGN LITTLE SECTION	0x104a	4170	Po
MYANMAR SIGN SECTION	0x104b	4171	Po
MYANMAR SYMBOL LOCATIVE	0x104c	4172	Po
MYANMAR SYMBOL COMPLETED	0x104d	4173	Po
MYANMAR SYMBOL AFOREMENTIONED	0x104e	4174	Po
MYANMAR SYMBOL GENITIVE	0x104f	4175	Po
GEORGIAN PARAGRAPH SEPARATOR	0x10fb	4347	Po
ETHIOPIC WORDSPACE	0x1361	4961	Po
ETHIOPIC FULL STOP	0x1362	4962	Po
ETHIOPIC COMMA	0x1363	4963	Po

Character Name	Hexadecimal Code	Decimal Code	Type
ETHIOPIC SEMICOLON	0x1364	4964	Po
ETHIOPIC COLON	0x1365	4965	Po
ETHIOPIC PREFACE COLON	0x1366	4966	Po
ETHIOPIC QUESTION MARK	0x1367	4967	Po
ETHIOPIC PARAGRAPH SEPARATOR	0x1368	4968	Po
CANADIAN SYLLABICS CHI SIGN	0x166d	5741	Po
CANADIAN SYLLABICS FULL STOP	0x166e	5742	Po
OGHAM SPACE MARK	0x1680	5760	Zs
OGHAM FEATHER MARK	0x169b	5787	Ps
OGHAM REVERSED FEATHER MARK	0x169c	5788	Pe
RUNIC SINGLE PUNCTUATION	0x16eb	5867	Po
RUNIC MULTIPLE PUNCTUATION	0x16ec	5868	Po
RUNIC CROSS PUNCTUATION	0x16ed	5869	Po
KHMER SIGN KHAN	0x17d4	6100	Po
KHMER SIGN BARIYOOSAN	0x17d5	6101	Po
KHMER SIGN CAMNUC PII KUUH	0x17d6	6102	Po
KHMER SIGN LEK TOO	0x17d7	6103	Po
KHMER SIGN BEYYAL	0x17d8	6104	Po
KHMER SIGN PHNAEK MUAN	0x17d9	6105	Po
KHMER SIGN KOOMUUT	0x17da	6106	Po
KHMER SIGN AVAKRAHASANYA	0x17dc	6108	Po
MONGOLIAN BIRGA	0x1800	6144	Po
MONGOLIAN ELLIPSIS	0x1801	6145	Po
MONGOLIAN COMMA	0x1802	6146	Po
MONGOLIAN FULL STOP	0x1803	6147	Po
MONGOLIAN COLON	0x1804	6148	Po
MONGOLIAN FOUR DOTS	0x1805	6149	Po
MONGOLIAN TODO SOFT HYPHEN	0x1806	6150	Pd
MONGOLIAN SIBE SYLLABLE BOUNDARY MARKER	0x1807	6151	Po
MONGOLIAN MANCHU COMMA	0x1808	6152	Po

Character Name	Hexadecimal Code	Decimal Code	Type
MONGOLIAN MANCHU FULL STOP	0x1809	6153	Po
MONGOLIAN NIRUGU	0x180a	6154	Po
EN QUAD	0x2000	8192	Zs
EM QUAD	0x2001	8193	Zs
EN SPACE	0x2002	8194	Zs
EM SPACE	0x2003	8195	Zs
THREE-PER-EM SPACE	0x2004	8196	Zs
FOUR-PER-EM SPACE	0x2005	8197	Zs
SIX-PER-EM SPACE	0x2006	8198	Zs
FIGURE SPACE	0x2007	8199	Zs
PUNCTUATION SPACE	0x2008	8200	Zs
THIN SPACE	0x2009	8201	Zs
HAIR SPACE	0x200a	8202	Zs
ZERO WIDTH SPACE	0x200b	8203	Zs
HYPHEN	0x2010	8208	Pd
NON-BREAKING HYPHEN	0x2011	8209	Pd
FIGURE DASH	0x2012	8210	Pd
EN DASH	0x2013	8211	Pd
EM DASH	0x2014	8212	Pd
HORIZONTAL BAR	0x2015	8213	Pd
DOUBLE VERTICAL LINE	0x2016	8214	Po
DOUBLE LOW LINE	0x2017	8215	Po
LEFT SINGLE QUOTATION MARK	0x2018	8216	Pi
RIGHT SINGLE QUOTATION MARK	0x2019	8217	Pf
SINGLE LOW-9 QUOTATION MARK	0x201a	8218	Ps
SINGLE HIGH-REVERSED-9 QUOTATION MARK	0x201b	8219	Pi
LEFT DOUBLE QUOTATION MARK	0x201c	8220	Pi
RIGHT DOUBLE QUOTATION MARK	0x201d	8221	Pf
DOUBLE LOW-9 QUOTATION MARK	0x201e	8222	Ps
DOUBLE HIGH-REVERSED-9 QUOTATION MARK	0x201f	8223	Pi

Character Name	Hexadecimal Code	Decimal Code	Type
DAGGER	0x2020	8224	Po
DOUBLE DAGGER	0x2021	8225	Po
BULLET	0x2022	8226	Po
TRIANGULAR BULLET	0x2023	8227	Po
ONE DOT LEADER	0x2024	8228	Po
TWO DOT LEADER	0x2025	8229	Po
HORIZONTAL ELLIPSIS	0x2026	8230	Po
HYPHENATION POINT	0x2027	8231	Po
LINE SEPARATOR	0x2028	8232	Zl
PARAGRAPH SEPARATOR	0x2029	8233	Zp
NARROW NO-BREAK SPACE	0x202f	8239	Zs
PER MILLE SIGN	0x2030	8240	Po
PER TEN THOUSAND SIGN	0x2031	8241	Po
PRIME	0x2032	8242	Po
DOUBLE PRIME	0x2033	8243	Po
TRIPLE PRIME	0x2034	8244	Po
REVERSED PRIME	0x2035	8245	Po
REVERSED DOUBLE PRIME	0x2036	8246	Po
REVERSED TRIPLE PRIME	0x2037	8247	Po
CARET	0x2038	8248	Po
SINGLE LEFT-POINTING ANGLE QUOTATION MARK	0x2039	8249	Pi
SINGLE RIGHT-POINTING ANGLE QUOTATION MARK	0x203a	8250	Pf
REFERENCE MARK	0x203b	8251	Po
DOUBLE EXCLAMATION MARK	0x203c	8252	Po
INTERROBANG	0x203d	8253	Po
OVERLINE	0x203e	8254	Po
UNDERTIE	0x203f	8255	Pc
CHARACTER TIE	0x2040	8256	Pc
CARET INSERTION POINT	0x2041	8257	Po
ASTERISM	0x2042	8258	Po

Character Name	Hexadecimal Code	Decimal Code	Type
HYPHEN BULLET	0x2043	8259	Po
LEFT SQUARE BRACKET WITH QUILL	0x2045	8261	Ps
RIGHT SQUARE BRACKET WITH QUILL	0x2046	8262	Pe
QUESTION EXCLAMATION MARK	0x2048	8264	Po
EXCLAMATION QUESTION MARK	0x2049	8265	Po
TIRONIAN SIGN ET	0x204a	8266	Po
REVERSED PILCROW SIGN	0x204b	8267	Po
BLACK LEFTWARDS BULLET	0x204c	8268	Po
BLACK RIGHTWARDS BULLET	0x204d	8269	Po
SUPERSCRIFT LEFT PARENTHESIS	0x207d	8317	Ps
SUPERSCRIFT RIGHT PARENTHESIS	0x207e	8318	Pe
SUBSCRIPT LEFT PARENTHESIS	0x208d	8333	Ps
SUBSCRIPT RIGHT PARENTHESIS	0x208e	8334	Pe
LEFT-POINTING ANGLE BRACKET	0x2329	9001	Ps
RIGHT-POINTING ANGLE BRACKET	0x232a	9002	Pe
IDEOGRAPHIC SPACE	0x3000	12288	Zs
IDEOGRAPHIC COMMA	0x3001	12289	Po
IDEOGRAPHIC FULL STOP	0x3002	12290	Po
DITTO MARK	0x3003	12291	Po
LEFT ANGLE BRACKET	0x3008	12296	Ps
RIGHT ANGLE BRACKET	0x3009	12297	Pe
LEFT DOUBLE ANGLE BRACKET	0x300a	12298	Ps
RIGHT DOUBLE ANGLE BRACKET	0x300b	12299	Pe
LEFT CORNER BRACKET	0x300c	12300	Ps
RIGHT CORNER BRACKET	0x300d	12301	Pe
LEFT WHITE CORNER BRACKET	0x300e	12302	Ps
RIGHT WHITE CORNER BRACKET	0x300f	12303	Pe
LEFT BLACK LENTICULAR BRACKET	0x3010	12304	Ps
RIGHT BLACK LENTICULAR BRACKET	0x3011	12305	Pe
LEFT TORTOISE SHELL BRACKET	0x3014	12308	Ps

Character Name	Hexadecimal Code	Decimal Code	Type
RIGHT TORTOISE SHELL BRACKET	0x3015	12309	Pe
LEFT WHITE LENTICULAR BRACKET	0x3016	12310	Ps
RIGHT WHITE LENTICULAR BRACKET	0x3017	12311	Pe
LEFT WHITE TORTOISE SHELL BRACKET	0x3018	12312	Ps
RIGHT WHITE TORTOISE SHELL BRACKET	0x3019	12313	Pe
LEFT WHITE SQUARE BRACKET	0x301a	12314	Ps
RIGHT WHITE SQUARE BRACKET	0x301b	12315	Pe
WAVE DASH	0x301c	12316	Pd
REVERSED DOUBLE PRIME QUOTATION MARK	0x301d	12317	Ps
DOUBLE PRIME QUOTATION MARK	0x301e	12318	Pe
LOW DOUBLE PRIME QUOTATION MARK	0x301f	12319	Pe
WAVY DASH	0x3030	12336	Pd
KATAKANA MIDDLE DOT	0x30fb	12539	Pc
ORNATE LEFT PARENTHESIS	0xfd3e	64830	Ps
ORNATE RIGHT PARENTHESIS	0xfd3f	64831	Pe
PRESENTATION FORM FOR VERTICAL TWO DOT LEADER	0xfe30	65072	Po
PRESENTATION FORM FOR VERTICAL EM DASH	0xfe31	65073	Pd
PRESENTATION FORM FOR VERTICAL EN DASH	0xfe32	65074	Pd
PRESENTATION FORM FOR VERTICAL LOW LINE	0xfe33	65075	Pc
PRESENTATION FORM FOR VERTICAL WAVY LOW LINE	0xfe34	65076	Pc
PRESENTATION FORM FOR VERTICAL LEFT PARENTHESIS	0xfe35	65077	Ps
PRESENTATION FORM FOR VERTICAL RIGHT PARENTHESIS	0xfe36	65078	Pe
PRESENTATION FORM FOR VERTICAL LEFT CURLY BRACKET	0xfe37	65079	Ps
PRESENTATION FORM FOR VERTICAL RIGHT CURLY BRACKET	0xfe38	65080	Pe
PRESENTATION FORM FOR VERTICAL LEFT TORTOISE SHELL BRACKET	0xfe39	65081	Ps
PRESENTATION FORM FOR VERTICAL RIGHT TORTOISE SHELL BRACKET	0xfe3a	65082	Pe

Character Name	Hexadecimal Code	Decimal Code	Type
PRESENTATION FORM FOR VERTICAL LEFT BLACK LENTICULAR BRACKET	0xfe3b	65083	Ps
PRESENTATION FORM FOR VERTICAL RIGHT BLACK LENTICULAR BRACKET	0xfe3c	65084	Pe
PRESENTATION FORM FOR VERTICAL LEFT DOUBLE ANGLE BRACKET	0xfe3d	65085	Ps
PRESENTATION FORM FOR VERTICAL RIGHT DOUBLE ANGLE BRACKET	0xfe3e	65086	Pe
PRESENTATION FORM FOR VERTICAL LEFT ANGLE BRACKET	0xfe3f	65087	Ps
PRESENTATION FORM FOR VERTICAL RIGHT ANGLE BRACKET	0xfe40	65088	Pe
PRESENTATION FORM FOR VERTICAL LEFT CORNER BRACKET	0xfe41	65089	Ps
PRESENTATION FORM FOR VERTICAL RIGHT CORNER BRACKET	0xfe42	65090	Pe
PRESENTATION FORM FOR VERTICAL LEFT WHITE CORNER BRACKET	0xfe43	65091	Ps
PRESENTATION FORM FOR VERTICAL RIGHT WHITE CORNER BRACKET	0xfe44	65092	Pe
DASHED OVERLINE	0xfe49	65097	Po
CENTRELINE OVERLINE	0xfe4a	65098	Po
WAVY OVERLINE	0xfe4b	65099	Po
DOUBLE WAVY OVERLINE	0xfe4c	65100	Po
DASHED LOW LINE	0xfe4d	65101	Pc
CENTRELINE LOW LINE	0xfe4e	65102	Pc
WAVY LOW LINE	0xfe4f	65103	Pc
SMALL COMMA	0xfe50	65104	Po
SMALL IDEOGRAPHIC COMMA	0xfe51	65105	Po
SMALL FULL STOP	0xfe52	65106	Po
SMALL SEMICOLON	0xfe54	65108	Po
SMALL COLON	0xfe55	65109	Po
SMALL QUESTION MARK	0xfe56	65110	Po
SMALL EXCLAMATION MARK	0xfe57	65111	Po

Character Name	Hexadecimal Code	Decimal Code	Type
SMALL EM DASH	0xfe58	65112	Pd
SMALL LEFT PARENTHESIS	0xfe59	65113	Ps
SMALL RIGHT PARENTHESIS	0xfe5a	65114	Pe
SMALL LEFT CURLY BRACKET	0xfe5b	65115	Ps
SMALL RIGHT CURLY BRACKET	0xfe5c	65116	Pe
SMALL LEFT TORTOISE SHELL BRACKET	0xfe5d	65117	Ps
SMALL RIGHT TORTOISE SHELL BRACKET	0xfe5e	65118	Pe
SMALL NUMBER SIGN	0xfe5f	65119	Po
SMALL AMPERSAND	0xfe60	65120	Po
SMALL ASTERISK	0xfe61	65121	Po
SMALL HYPHEN-MINUS	0xfe63	65123	Pd
SMALL REVERSE SOLIDUS	0xfe68	65128	Po
SMALL PERCENT SIGN	0xfe6a	65130	Po
SMALL COMMERCIAL AT	0xfe6b	65131	Po
FULLWIDTH EXCLAMATION MARK	0xff01	65281	Po
FULLWIDTH QUOTATION MARK	0xff02	65282	Po
FULLWIDTH NUMBER SIGN	0xff03	65283	Po
FULLWIDTH PERCENT SIGN	0xff05	65285	Po
FULLWIDTH AMPERSAND	0xff06	65286	Po
FULLWIDTH APOSTROPHE	0xff07	65287	Po
FULLWIDTH LEFT PARENTHESIS	0xff08	65288	Ps
FULLWIDTH RIGHT PARENTHESIS	0xff09	65289	Pe
FULLWIDTH ASTERISK	0xff0a	65290	Po
FULLWIDTH COMMA	0xff0c	65292	Po
FULLWIDTH HYPHEN-MINUS	0xff0d	65293	Pd
FULLWIDTH FULL STOP	0xff0e	65294	Po
FULLWIDTH SOLIDUS	0xff0f	65295	Po
FULLWIDTH COLON	0xff1a	65306	Po
FULLWIDTH SEMICOLON	0xff1b	65307	Po
FULLWIDTH QUESTION MARK	0xff1f	65311	Po

Character Name	Hexadecimal Code	Decimal Code	Type
FULLWIDTH COMMERCIAL AT	0xff20	65312	Po
FULLWIDTH LEFT SQUARE BRACKET	0xff3b	65339	Ps
FULLWIDTH REVERSE SOLIDUS	0xff3c	65340	Po
FULLWIDTH RIGHT SQUARE BRACKET	0xff3d	65341	Pe
FULLWIDTH LOW LINE	0xff3f	65343	Pc
FULLWIDTH LEFT CURLY BRACKET	0xff5b	65371	Ps
FULLWIDTH RIGHT CURLY BRACKET	0xff5d	65373	Pe
HALFWIDTH IDEOGRAPHIC FULL STOP	0xff61	65377	Po
HALFWIDTH LEFT CORNER BRACKET	0xff62	65378	Ps
HALFWIDTH RIGHT CORNER BRACKET	0xff63	65379	Pe
HALFWIDTH IDEOGRAPHIC COMMA	0xff64	65380	Po
HALFWIDTH KATAKANA MIDDLE DOT	0xff65	65381	Pc

Index

Symbols

- ! routing character, [131](#)
- % routing character, [131](#)

Numerics

- 8-bit file transfer, [188](#)
- 8BITMIME, [136](#)

A

actions

- FTP, [185](#)
- HTTP, [172](#)
- POP3, [199](#)
- SMTP, [141](#)
- primary actions, [142](#)
- secondary actions, [143](#)

ActiveX, [172](#)

- MacroMedia Flash, [172](#)

alerts, [212](#)

- e-mail, [212](#)
- ePolicy Orchestrator, [212](#)
- SNMP, [212](#)

aliases, [98](#)

allow through

- FTP, [185](#)
- HTTP, [172](#)
- POP3, [200](#)
- SMTP, [130](#), [143](#), [198](#)

amazing offers! , spam phrases, [150](#)

annotated e-mail, [144](#)

anti-relay, [113](#)

- blocking actions, [115](#)
- deny domains, [114](#)
- local domains, [100](#), [114](#)
- permit domains, [114](#)
- response, [115](#)
- routing characters, [113](#)

anti-spam

- engine, [152](#)
- extra rules, [148](#), [152](#)
- rules, [152](#)
- scheduling updates, [152](#)
- updates web site, [18](#)

anti-virus software, [164](#)

- scheduling updates, [165](#)
- updates web site, [18](#)
- updating, [164](#)

APOP, [192](#)

appliance

- automatic updates, [164](#)
- anti-spam, [152](#)
- anti-virus, [164](#)
- default IP addresses, [20](#)
- default password, [20](#)
- default system name, [20](#)
- defaults
 - IP addresses, [20](#)
 - password, [102](#)
- domain name, [96](#)
- handling traffic, [23](#)
- host name, [20](#), [96](#)
- HTTPS, [20](#), [24](#)
- load sharing, [79](#), [215](#)
- logs, [225](#)
- management interface, [24](#)
- operational modes, [22](#), [31](#)
- overview, [19](#)
- password, [102](#)
- ports, [34](#), [38](#), [41](#), [43](#), [120](#), [169](#), [183](#), [191](#)
- positioning, [23](#)

- protocol support, [22, 34, 38](#)
 - APOP, [192](#)
 - Content Vectoring Protocol (CVP), [22](#)
 - disabling, [23, 26, 97](#)
 - enabling, [23, 26, 97](#)
 - FTP, [22, 183](#)
 - HTTP, [22](#)
 - HTTPS, [22](#)
 - IPX, [34, 38](#)
 - multicast IP, [34](#)
 - NetBEUI, [34, 38](#)
 - OSPF, [101](#)
 - POP3, [22](#)
 - Real Player, [22](#)
 - RIP, [101](#)
 - streaming media, [22](#)
- rebooting, [223](#)
- recommended topologies, [23](#)
- restarting, [223](#)
- restoring configuration, [226](#)
- restrictions, [24](#)
- saving system settings, [226](#)
- scanning traffic, [26](#)
- security, [20, 23](#)
- system language, [102](#)
- time zone settings, [102](#)
- virus scanning, see virus scanning, [163](#)
- what is it?, [19](#)

ASCII, [229](#)

- blocking 8-bit ASCII, [187](#)

audience for this manual, [15](#)

AUTH, [136](#)

AVERT (Anti-Virus Emergency Response Team),
contacting, [18](#)

B

- beta program, contacting, [18](#)
- bitmap, see BMP, [161](#)
- blackhole lists, [118](#)
- BMP, [161](#)
- bridgehead mail server, [52, 71, 86](#)
- Bubbleboy, [168](#)

C

- characters
 - not detected, [156](#)
 - used as delimiters, [158](#)
- check for client
 - HTTP, [173](#)
- cleaning e-mail viruses, [142](#)
- client alert messages, [174](#)
- client download status message, [174](#)
- command-line FTP clients, [56](#)
- commands
 - NOOP, [136](#)
- comma-separated value files
 - see CSV files, [114](#)
- confidential information, [160](#)
- connection policies
 - FTP, [185](#)
 - HTTP, [173](#)
 - POP3, [200](#)
 - SMTP, [131](#)
- connections
 - load sharing, [218](#)
- contacting McAfee Security, [18](#)
- content policies
 - FTP, [184](#)
 - HTTP, [171](#)
 - POP3, [194](#)
 - SMTP, [122, 141](#)
- content rules, [103](#)
 - and rule groups, [108, 153](#)
 - definition, [103, 108, 153](#)
 - using, [108, 153](#)
 - warning about names of, [155](#)
- content scanning, [28](#)
- Content Vectoring Protocol (CVP), [22](#)
- controlling appliance, [29, 215](#)
- conventions used in this manual, [16](#)
- copying system settings, [226](#)
- CSV files, [114](#)
- currency symbols, [156](#)
- customer service, contacting, [18](#)
- CVP, see Content Vectoring Protocol, [22](#)

D

DAT file
 updates web site, 18

DAT files
 automatic updates, 164
 definition, 164

data trickling
 FTP, 189
 HTTP, 175

dedicated POP3 proxy, 192

default settings
 IP addresses, 20
 system name, 20
 user name and password, 20

delimiter characters, 158

delivery methods (SMTP), 111

denial-of-service prevention, 128, 134, 172, 174, 185, 197

deny domains, 114

depth of nesting, 128, 197
 in compressed files, 129, 197

detection events, 212

directional scanning, 26

disclaimers, 124

distractions, 162

DNS servers, 53, 57, 112
 disabling, 112
 enabling, 112
 listing, 101
 use forward only mode, 101

documentation for the product, 17

domain name, 96

domains
 appliance domain, 96
 Inside Networks, 25
 Outside Networks, 25

domains relays, 112

dos2unix, 188

downgrade to HTTP 1.0, 178

download status page, 175

download web site, 18

downloading the latest anti-spam files, 152

DSN, 136

dynamic routes, 101
 disabling, 102
 enabling, 102
 routing table, 102

E

e-mail
 allowing through, 130, 143, 198
 annotated e-mail, 144
 cleaning, 142
 forwarding, 144
 logging, 212
 messages, 111
 modify message, 143
 quarantine, 103
 replace content, 142
 signed e-mails, 130, 142, 198
 store and forward, 121
 unsolicited messages, 147

e-mail delivery methods, 111

engine
 anti-spam
 automatic updates, 152
 definition, 152
 anti-virus
 automatic updates, 164
 definition, 164

ePolicy Orchestrator, 96, 225
 logging events, 212

ESMTP, 136

evaluating SpamKiller, 224

events, 212
 codes, 212
 severity, 212

explicit
 definition, 32

Explicit Proxy mode, 22, 34, 66
 basic concepts, 32
 changing modes, 43
 configuration examples, 36, 45
 definition of proxy, 32
 FTP configuration examples, 56
 FTP inbound examples, 57 to 58

- FTP outbound example, 56
- HTTP configuration examples, 59
- HTTP inbound not recommended, 62
- HTTP outbound examples, 59 to 60, 62
- load sharing example, 66
- POP3 configuration examples, 63
- POP3 dedicated proxy example, 64
- POP3 generic proxy example, 63
- POP3 multiple servers example, 65
- ports, 34
 - load sharing, 66
- positioning the appliance, 34
- recommended topologies, 45
- SMTP configuration examples, 46
- SMTP demilitarized zone example, 49 to 50
- SMTP directional scanning example, 47
- SMTP handling network failure, 53
 - fail-closed, 54
 - fail-open, 55
 - fail-over, 53
- SMTP international organization example, 52
- SMTP multiple sites example, 49
- SMTP one site example, 46
- streaming media, 180
- external networks, 25
- EXTRA DAT files
 - definition, 164
- extra rules, 152
- F**
 - fail-closed configuration, 54
 - fail-open configuration, 55
 - fail-over configuration, 53
 - fail-safe configuration, 54
 - fallback relays, 111 to 112
 - field delimiters, 158
 - file format, 126
 - firewall, 46, 81 to 82
 - Flash (MacroMedia), 172
 - FTP, 183
 - ASCII-mode, 187
 - client messages, 186
 - connection policies, 185
 - content policies, 184
 - data processing, 186
 - data trickling, 189
 - definition, 183
 - handoff Host, 188
 - intercept ports, 183
 - keep-alive commands, 187
 - listen ports, 183
 - port number 21, 183
 - protocol policies, 186
 - upload status message, 189
 - FTP commands
 - denied, 187
 - FTP over HTTP, 176
- G**
 - generic POP3 proxy, 202
 - getting information, 17
 - global policies
 - definition, 103
 - global policy
 - defined, 104
- H**
 - handoff host, 57
 - FTP, 188
 - HTTP, 176
 - header blocking, 177
 - heuristic analysis, 163, 166
 - host name, 96
 - HotFix, 224
 - HotFixes, 224
 - HTML alert, 142, 172
 - HTML settings
 - http, 172
 - HTTP
 - actions, 172
 - blocking
 - HTML elements, 172
 - check for client, 173
 - client alert messages, 174
 - client download status message, 174
 - connection policies, 173

- content policies, 171
- data trickling, 175
- downgrade to 1.0, 178
- download status page, 175
- handoff host, 176
- header blocking, 177
- intercept ports, 169
- listen ports, 169
- non-compliant POSTs, 178
- NTLM failure pages, 178
- OPTIONS, 178
- port number 80, 169
- port numbers, 181
- protocol policies, 174
- request permissions, 180
- request schemes, 181
- server internal information pages, 178
- SSL ports, 181
- streaming media, 179
- time-outs, 173
- TRACE, 178
- URL blocking, 180
- verbs, 181

HTTPS, 20, 22, 24

I

- inheritance
 - policies, 105
- Inside Networks, 25
- Insult 23, 155
- Intercept ports
 - FTP, 183
 - HTTP, 169
 - POP3, 191
 - SMTP, 120
- intercept ports
 - FTP, 183
 - HTTP, 169
 - POP3, 191
 - SMTP, 120
- internal networks, 25, 98
- IP addresses, 98
 - aliases, 98

- changing modes, 43
- defaults, 20
- primary IP addresses, 98

IPX, 34, 38

J

- Java applets, 172
- JavaScript, 172

K

- keep-alive commands
 - FTP, 187
- keep-alive headers, 175
- KnowledgeBase search, 18

L

- legal implications, use of e-mail and Internet, 110
- liability, limiting, 124
- listen ports
 - FTP, 183
 - HTTP, 169
 - POP3, 191
 - SMTP, 120
- listeners, 120, 170, 184, 192
 - load sharing, 218
- load sharing, 29, 66, 79, 94, 215
 - connections, 218
 - controlling appliance, 29, 215
 - definition, 29, 215
 - examples, 219
 - Explicit Proxy mode, 66
 - listeners, 218
 - load sharing appliance, 29, 215
 - scenarios, 216
 - Transparent Bridge mode, 94
 - Transparent Router mode, 79
- load sharing appliance, 29, 215
- local domains, 114
- localhost, 219
- logging
 - detail, 212
 - distribution, 212
- logs

overview, [225](#)

Lotus Notes, [213](#)

M

MacroMedia Flash, [172](#)

manuals, [17](#)

MAPS, [113](#)

McAfee Security University, contacting, [18](#)

Melissa, W97M/Melissa@MM virus, [168](#)

messages

example of spam, [150](#)

MIB file

SNMP, [212](#)

Microsoft Outlook, [213](#)

MIME, [159](#)

8BITMIME, [136](#)

modify e-mail message, [143](#)

modifying the subject line, [150](#)

MPEGs, [161](#)

multicast IP, [34](#), [38](#)

N

nesting, see depth of nesting, [128](#), [197](#)

NetBEUI, [34](#), [38](#)

network

addresses, [98](#)

load, [161](#)

non-compliant post requests, [178](#)

non-global policies

definition, [103](#)

See also policies

NOOP, [136](#)

notification e-mail, [143](#)

NTLM failure pages, [178](#)

nuisance e-mail, [161](#)

O

offensive words, [161](#)

operational modes, [22](#), [97](#)

changing modes, [43](#)

differences between, [31](#) to [32](#)

explicit or transparent?, [32](#)

which mode, [31](#)

OPTIONS, [178](#)

ordering policies, [105](#)

OSPF, [101](#)

outbound scenario, HTTP, [90](#)

Outside Networks, [25](#)

P

password

default password, [20](#)

permit domains, [114](#)

permit sender, [148](#)

policies

actions

FTP, [185](#)

POP3, [199](#)

content rules. See content rules

defined, [103](#)

FTP connection policies, [185](#)

FTP content policies, [184](#)

FTP protocol policies, [186](#)

global

definition, [103](#)

HTTP actions, [172](#)

HTTP connection policies, [173](#)

HTTP content policies, [171](#)

HTTP protocol policies, [174](#)

inheritance, [105](#)

non-global

definition, [103](#)

ordering, [105](#)

time restrictions, [106](#)

policy groups. See policy groups

POP3 connection policies, [200](#)

POP3 content policies, [194](#)

POP3 protocol policies, [201](#)

rule groups. See rule groups

SMTP connection policies, [131](#)

SMTP content policies, [122](#)

SMTP protocol policies, [133](#)

what is a policy?, [29](#), [103](#)

policy groups, [103](#)

creating, [108](#)

definition, [103](#)

- POP3, 191 to 192
 - connection policies, 200
 - content policies, 194
 - dedicated proxy, 192
 - definition, 191 to 192
 - generic proxy, 202
 - intercept ports, 191
 - listen ports, 191
 - port number 110, 191
 - protocol policies, 201
 - port numbers, 181
 - streaming media, 180
 - ports
 - Explicit Proxy mode, 34
 - intercept ports, 120, 169, 183, 191
 - listen ports, 120, 169, 183, 191
 - Transparent Bridge mode, 41
 - Transparent Router mode, 38
 - postmaster, 113, 213
 - prefix
 - spam messages identified, 150
 - primary actions
 - defined, 141
 - primary IP addresses, 98
 - PrimeSupport, 18
 - product documentation, 17
 - product overview, 19
 - product training, contacting, 18
 - protocol policies
 - FTP, 186
 - HTTP, 174
 - POP3, 201
 - SMTP, 133
 - protocol support, 22
 - APOP, 192
 - Content Vectoring Protocol (CVP), 22
 - disabling, 23, 26, 97
 - enabling, 23, 26, 97
 - FTP, 22, 183
 - HTTP, 22
 - HTTPS, 22
 - IPX, 34, 38
 - multicast IP, 34, 38
 - NetBEUI, 34, 38
 - OSPF, 101
 - POP3, 22
 - Real Player, 22
 - RIP, 101
 - SMTP, 22
 - streaming media, 22
 - proxy
 - definition, 32
- ## Q
- quarantine, 103, 160
 - specifying action, 155
 - quarantine area
 - areas at critical level, 213
- ## R
- RBL servers list, 118
 - Real Player traffic, 22
 - rebooting the appliance, 223
 - relaying e-mail, 112 to 113
 - request permissions, 180
 - request schemes, 181
 - restarting the appliance, 223
 - retryer, 121
 - retryer settings, 121
 - RIP, 101
 - routing characters, 113, 115, 131
 - rule groups, 108, 153
 - rules, 160
 - description, 155
 - e-mail forwarding, 113
 - extra rules, 148
 - name, 155
 - problems with complex rules, 159
- ## S
- saving system settings, 226
 - scanning
 - affect on performance, 166
 - scanning options
 - enable macro heuristics, 166
 - enable program file heuristics, 166

- expand archive files, [167](#)
 - expand compressed program files, [167](#)
 - jokes, [167](#)
 - scan all files, [167](#)
 - scan for suspicious programs, [167](#)
 - score, see spam score, [150](#)
 - secondary actions
 - defined, [141](#)
 - Secure Socket Layer (SSL), [181](#)
 - security headquarters, contacting AVERT, [18](#)
 - separators, for words, [158](#)
 - server for internal information pages, [178](#)
 - service packs, [224](#)
 - service portal, PrimeSupport, [18](#)
 - signed e-mail messages, [130](#), [142](#), [198](#)
 - SMTP, [111](#)
 - actions, [141](#)
 - primary, [142](#)
 - secondary, [143](#)
 - allow through, [130](#), [143](#), [198](#)
 - anti-relay, [113](#)
 - anti-spam
 - RBL servers list, [118](#)
 - connection policies, [131](#)
 - content policies, [122](#)
 - content scanning, [28](#), [229](#)
 - delivery methods, [111](#)
 - Explicit communication example, [32](#)
 - Inside Networks, [114](#)
 - intercept ports, [120](#)
 - listen ports, [120](#)
 - local domains, [114](#)
 - notification e-mail, [143](#)
 - port number 25, [120](#)
 - protocol policies, [133](#)
 - transparent communication example, [33](#)
 - using .CSV file, [114](#)
 - zero bytes, [143](#)
 - SNMP
 - definition, [212](#)
 - logging, [212](#)
 - MIB file, [212](#)
 - traps, [212](#)
 - ** SPAM ** prefix, [150](#)
 - spam, [150](#)
 - blacklist, [149](#)
 - disabling some rules against, [151](#)
 - example message, [150](#)
 - getting the balance right, [147](#)
 - overview, [147](#)
 - phrases, [150](#)
 - rules, [149](#)
 - score, [149](#) to [150](#)
 - tips for avoiding, [151](#)
 - whitelist, [149](#)
 - spam score
 - indicator, [151](#)
 - typical value of 5, [150](#)
 - SpamKiller for WebShield appliances
 - activating, [224](#)
 - evaluating, [224](#)
 - overview, [147](#)
 - static routes, [101](#)
 - streaming media, [179](#)
 - allow pass through, [180](#)
 - port number, [180](#)
 - submitting a sample virus, [18](#)
 - substitution variables, [104](#)
 - system language, [102](#)
 - system settings
 - restoring, [226](#)
 - saving, [226](#)
- T**
- TCP/IP, [120](#), [169](#), [183](#), [191](#)
 - Technical Support, [18](#)
 - third-party
 - relaying, [113](#)
 - time restrictions
 - non-global policies, [106](#)
 - time-outs
 - HTTP, [173](#)
 - tokens, [104](#)
 - TRACE, [178](#)
 - training web site, [18](#)
 - transparent

definition, 33

Transparent Bridge mode, 22, 41, 94

- basic concepts, 33
- changing modes, 43
- configuration examples, 81
- definition of transparent, 33
- FTP configuration examples, 87
- FTP inbound example, 88
- FTP outbound example, 87
- HTTP configuration examples, 88
- HTTP inbound examples, 92
- HTTP outbound examples, 89 to 91
- load sharing, 94
- POP3 configuration example, 93
- positioning the appliance, 41
- recommended topologies, 81
- SMTP configuration examples, 81
- SMTP dedicated appliances example, 83
- SMTP demilitarized zone not recommended, 85
- SMTP international organization example, 86
- SMTP multiple sites example, 84
- SMTP one site example, 82
- streaming media, 180

Transparent Router mode, 22, 38, 79, 94

- basic concepts, 33
- changing modes, 43
- configuration example, 40
- configuration examples, 67
- definition of transparent, 33
- FTP configuration examples, 72
- FTP inbound FTP example, 73
- FTP outbound examples, 72
- HTTP configuration examples, 73
- HTTP inbound example, 77
- HTTP outbound examples, 74 to 76
- load sharing example, 79
- POP3 configuration example, 78
- positioning the appliance, 39
- recommended topologies, 67
- SMTP configuration examples, 67
- SMTP demilitarized zone not recommended, 70
- SMTP international organization example, 71

- SMTP multiple sites example, 69
- SMTP one site example, 68
- streaming media, 180
- true FTP clients, 56
- TSV files, 209

U

- UDP, 120, 169, 183, 191
- unicode, 229
- unix2dos, 188
- update web site
 - anti-spam, 18
 - anti-virus, 18
- upgrades, 18
- upload status message, 189
- URL blocking, 180, 182
- user name
 - appliance, 227
 - FTP, 57
 - POP3 account, 63, 202

V

- VBS/Bubbleboy@MM virus, 168
- VBScript, 172
- Virus Information Library, 18
- virus scanning
 - across appliances, 29, 215
 - anti-virus software, 164
 - DAT files, 26
 - automatic updates, 164
 - directional scanning, 26
 - engine, 26, 164
 - heuristic analysis, 163
 - load sharing, 29
 - overview, 26
 - virus definition files (see DAT Files), 26
 - virus signatures, 163
- virus, submitting a sample, 18
- viruses
 - hoaxes, 147
- virus-scanning engine
 - see engine, 164
- Visual Basic script, 172

W

WAN link, [52](#), [71](#), [86](#)

warnings

 complex rules, [159](#)

 disclaimer, [124](#)

WebShield appliance

 policies explained, [29](#), [103](#)

whitelists, [149](#)

wildcards, [156](#)

word delimiters, [158](#)

word separators, [229](#)

word, defined, [158](#)

X

XML, [226](#)

Z

zero bytes, [143](#)