

Gauntlet[®] for IRIX[®] Administrator's Guide

Version 4.1

Document Number 007-2826-007

CONTRIBUTORS

Written by John Raithel with updates by Pam Sogard, Susan Thomas, Renate Kempf, and Terry Schultz based on materials from Trusted Information Systems, Inc.
Production by Mary Macanek, Heather Hermstad, and Amy Swenson.
Engineering contributions by Jessica Humphreys, Ed Mascarenhas, Dj Padzensky, and Mayank Vasa.
St. Peter's Basilica image courtesy of ENEL SpA and InfoByte SpA. Disk Thrower image courtesy of Xavier Berenguer, Animatica.

© 1998, 1999, Silicon Graphics, Inc.— All Rights Reserved

The contents of this document may not be copied or duplicated in any form, in whole or in part, without the prior written permission of Silicon Graphics, Inc.

LIMITED AND RESTRICTED RIGHTS LEGEND

Use, duplication, or disclosure by the Government is subject to restrictions as set forth in the Rights in Data clause at FAR 52.227-14 and/or in similar or successor clauses in the FAR, or in the DOD, DOE or NASA FAR Supplements. Unpublished rights reserved under the Copyright Laws of the United States.

Contractor/manufacture is SGI, 1600 Amphitheatre Parkway, Mountain View, CA 94043-1389.

Silicon Graphics is a registered trademark and SGI and the SGI logo are trademarks of Silicon Graphics, Inc. UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company, Ltd. Gauntlet is a trademark of Trusted Information Systems, Inc. Lotus Notes is a registered trademark of Lotus Development Corporation. Windows is a registered trademark and Windows NT and NetShow are trademarks of Microsoft Corporation. RealAudio is a registered trademark and RealVideo is a trademark of Real Networks, Inc. VDOLive is a trademark of VDOnet Corporation. Netscape Navigator is a registered trademark of Netscape Communications, Inc. Java and JavaScript are trademarks of Sun Microsystems, Inc.

Contents

List of Figures	xxv
About This Guide	xxix
Audience for This Guide	xxix
Contents of This Guide	xxix
Conventions Used in This Guide	xxx
Additional Documentation	xxx
Additional Resources	xxx
Books	xxxi
Newsgroups	xxxi
Mailing Lists	xxxi
Frequently Asked Questions Lists	xxxi
White Papers	xxxii
How to Get Latest Security Patches	xxxii

PART I Introducing the Gauntlet Firewall

1. Understanding the Gauntlet Firewall	1
Understanding Gauntlet Firewall Concepts	1
Gauntlet Firewall Design Philosophy	1
Establishing a Security Perimeter	2
Trusted and Untrusted Networks	2
Trusted Networks	2
Untrusted Networks	3
Unknown Networks	3
Service Groups	3
Transparency	4

- Understanding Gauntlet Firewall Components 5
 - Gauntlet Firewall Hardware 5
 - Gauntlet Firewall Software 5
 - Gauntlet Firewall Operating System 5
 - Application Level Security Services (Proxies) 5
 - IP Screening Utility 7
 - Management Utilities 8
 - How the Gauntlet Firewall Works 8
 - Dual-Homed Bastion Host 10
 - Processing Packets and Requests 12
 - Receiving Packets 12
 - Checking Source and Destination 13
 - Checking the Request Type 13
 - Calling the Appropriate Program 13
 - Processing the Request 14
- 2. Managing the Gauntlet Firewall 15**
 - General Management and Maintenance 15
 - Creating User Accounts 15
 - Backing Up and Restoring 16
 - Managing Gauntlet Firewall Options 16
 - Understanding the Gauntlet Firewall Manager 17
 - Firewall Manager Graphical Interface 17
 - Firewall Server 17

- Using the Gauntlet Firewall Manager 18
 - Planning the Firewall 18
 - Configuring the Firewall 19
 - Configuring Your System 20
 - Configuring Your Web Browser 20
 - Accessing the Gauntlet Firewall Manager 20
 - Exiting the Gauntlet Firewall Manager 22
 - Saving Your Changes 22
 - Quit Menu Command 22
 - Save Menu Command 23
 - Save and Apply Menu Command 23
 - Save, Apply, and Reboot Menu Command 23
 - Reboot Menu Command 23
 - Cancel 24
 - Getting Help 24

PART II Managing the Gauntlet Firewall

- 3. **Service Groups and Service Group Rules 27**
 - Understanding Service Groups 27
 - Default Service Groups 28
 - Trusted Service Group 29
 - Untrusted Service Group 29
 - Service Group Membership 30
 - When to Add New Service Groups 30
 - Accessing Service Group Configuration 31
 - Configuring Service Groups 32
 - Planning Service Groups 32
 - Creating Service Groups 33
 - Modifying Service Groups 34
 - Deleting Service Groups 35
 - Understanding Service Group Rules 35
 - Order of Precedence 35
 - Accessing Service Group Rules Configuration 36

- Adding Service Group Rules 37
- Modifying Service Group Rules 38
- Deleting Service Group Rules 38
- Changing Order of Precedence 38
- 4. Destination Access 39**
 - Understanding Destination Access 39
 - How Destination Access Works 40
 - Default Destination Access Rules 40
 - Order of Precedence 40
 - Considerations for Specifying Destination Addresses 42
 - Unknown Keyword 42
 - Accessing Destination Access Configuration 43
 - Configuring Destination Access Rules 44
 - Planning Destination Access Rules 45
 - Creating Destination Access Rules 45
 - Modifying Destination Access Rules 46
 - Deleting Destination Access Rules 46
 - Changing Order of Precedence 46
- 5. Networks and Network Groups 47**
 - Understanding Networks and Network Groups 47
 - Default Network Groups 48
 - Trusted Network Group 48
 - Untrusted Network Group 49
 - When to Add New Network Groups 49
 - Accessing Network Configuration 50
 - Configuring Networks 50
 - Planning Networks 50
 - Creating Networks 51
 - Modifying Networks 52
 - Deleting Networks 53
 - Accessing Network Group Configuration 53

Configuring Network Groups	54
Planning Network Groups	54
Creating Network Groups	54
Modifying Network Groups	56
Deleting Network Groups	56
6. Users and User Groups	57
Understanding the User Authentication Management System	57
How the Firewall Uses Authentication Information	58
How Other Services Use Authentication Information	58
Understanding Users and Groups	59
Users in the Authentication Management System	59
Groups in the Authentication Management System	59
Understanding Strong Authentication	60
Access Key II Authentication	60
APOP	60
CRYPTOCARD RB-1	61
Digipass Authentication	61
SafeWord Authentication Server	61
SecurID Authentication	62
S/Key Authentication	62
RADIUS Authentication	62
Reusable Passwords	63

- Configuring Users 63
 - Accessing User Configuration 63
 - Creating Users 64
 - Modifying Users 64
 - Changing User Names 65
 - Changing User IDs 65
 - Changing Group Membership 65
 - Changing Authentication Method 66
 - Changing Passwords 66
 - Allowing Users to Change Their Password 67
 - Enabling Users 68
 - Disabling Users 68
 - Deleting Users 69
- Managing Groups 69
 - Creating Groups 69
 - Disabling Groups 70
 - Deleting Groups 70
- 7. User Restrictions 71**
 - Understanding User Restrictions 71
 - How User Restrictions Work 72
 - Order of Precedence 72
 - Accessing User Restriction Configuration 73
 - Configuring User Restriction Rules 75
 - Planning User Restriction Rules 75
 - Creating User Restriction Rules 75
 - Modifying User Restriction Rules 76
 - Deleting User Restriction Rules 76
 - Changing Order of Precedence 77

PART III Configuring and Using Proxy Services

- 8. Managing Proxy Services 81**
 - Understanding Proxy Services 81

-
- Configuring Proxy Services 82
 - Creating Multiple Configurations for a Proxy 82
 - Working With Configuration Sets 83
 - Planning Configuration Sets 83
 - Creating Configuration Sets 84
 - Modifying Configuration Sets 85
 - Deleting Configuration Sets 85
 - 9. Managing FTP Services 87**
 - Understanding the FTP Proxy 87
 - How the FTP Proxy Works 88
 - Accessing FTP Proxy Configuration 89
 - Configuring the Firewall for FTP Services 89
 - Planning FTP Proxy Settings 90
 - Configuring FTP Proxy Settings 90
 - Enabling FTP Proxy Services 90
 - Creating Authentication User Entries 90
 - Verifying Your Setup 91
 - Using FTP Services 91
 - Using Authentication 91
 - Using Authentication With Some GUI FTP Tools 92
 - Running an Anonymous FTP Server 94
 - 10. Managing LDAP Services 95**
 - Understanding the LDAP Proxy 95
 - How the LDAP Proxy Works 96
 - Planning the LDAP Proxy 97
 - Configuring LDAP Clients 97
 - Configuring LDAP Clients With Transparency 97
 - Configuring LDAP Clients Without Transparency 97
 - Configuring LDAP Proxy Settings 97
 - Enabling LDAP Services 99
 - 11. Managing Microsoft SQL Services 101**
 - Understanding the SQL Server Proxy 101
 - How the SQL Server Proxy Works 102

- Accessing SQL Server Proxy Configuration 103
- Configuring the Firewall for Microsoft SQL Services 104
 - Planning SQL Server Proxy Settings 104
 - Configuring SQL Server Proxy Settings 105
- Enabling SQL Server Proxy Services 107
- Configuring Microsoft SQL Clients 108
- Verifying Your Setup 108
- 12. Managing Multimedia Services 109**
 - Understanding the Multimedia Proxy 110
 - How the Multimedia Proxy Works 110
 - Accessing Multimedia Proxy Configuration 111
 - Configuring the Firewall for Multimedia Services 112
 - Planning the Firewall for Multimedia Services 112
 - Configuring Multimedia Proxy Services 113
 - Enabling Multimedia Proxy Services 113
 - Verifying Your Setup 113
 - Using the NetShow Proxy 114
 - Using the NetShow Proxy With Transparency 114
 - Using the NetShow Proxy Without Transparency 114
 - Using the RealPlayer Proxy 115
 - Using the RealPlayer Proxy With Transparency 115
 - Using the Proxy for Real Audio and Real Video Proxy Without Transparency 116
 - Using the StreamWorks Proxy 117
 - Using the VDOLive Proxy 118
- 13. Managing Network Management Services 121**
 - Understanding the SNMP Proxy 121
 - How the SNMP Proxy Works 122
 - SNMP Requests 122
 - SNMP Trap Requests 123
 - Accessing SNMP Proxy Configuration 123

- Configuring the Firewall for SNMP Services 124
 - Planning the Firewall for SNMP Services 124
 - Configuring SNMP Proxy Settings 125
 - Enabling SNMP Proxy Services 126
- Configuring SNMP Agents 127
- 14. Managing News Services 129**
 - Understanding the News Proxy 130
 - How the News Proxy Works 130
 - Configuring the Firewall for News Services 131
 - Planning News Settings 131
 - Configuring News Settings 131
 - Enabling News Proxy Services 133
 - Informing Your News Feed 134
 - Configuring Your News Server 134
 - Using News 135
- 15. Managing Print Services 137**
 - Understanding the lp Proxy 137
 - How the lp Proxy Works 138
 - Configuring the Print Client 140
 - Configuring the Print Server 140
 - Accessing lp Proxy Configuration 141
 - Configuring the Firewall for lp Services 142
 - Planning lp Proxy Settings 142
 - Configuring lp Proxy Settings 142
 - Enabling the *lp* Proxy 144
 - Using lp Services 144
- 16. Managing rsh Services 145**
 - Understanding the rsh Proxy 145
 - How the rsh Proxy Works 146
 - Accessing rsh Proxy Configuration 147

- Configuring the Firewall for rsh Services 147
 - Planning rsh Proxy Settings 147
 - Configuring rsh Proxy Settings 148
 - Enabling rsh Proxy Services 148
 - Verifying Your Setup 148
- Using rsh Services 149
 - Configuring the Remote System 149
- 17. Managing Sybase Services 151**
 - Understanding the Sybase Proxy 151
 - How the Sybase Proxy Works 152
 - Accessing Sybase Proxy Configuration 153
 - Configuring the Firewall for Sybase Services 154
 - Planning Sybase Proxy Settings 154
 - Configuring Sybase Proxy Settings 154
 - Enabling Sybase Proxy Services 155
 - Configuring Sybase Clients 156
 - Verifying Your Setup 156
- 18. Managing Terminal Services 157**
 - Understanding the TELNET and rlogin Proxies 158
 - How the TELNET and rlogin Proxies Work 159
 - Accessing TELNET and rlogin Proxy Configuration 160
 - Configuring the Firewall for Terminal Services 161
 - Planning TELNET and rlogin Proxy Settings 161
 - Configuring TELNET and rlogin Proxy Settings 162
 - Enabling TELNET and rlogin Proxy Services 162
 - Creating Authentication User Entries 162
 - Verifying Your Setup 162
 - Using Terminal Services 163
 - TELNET, rlogin, and TN3270 Without Authentication 163
 - TELNET and rlogin with Authentication 163
 - TN3270 With Authentication 165

-
- 19. **Managing WWW and Gopher Services** 167
 - Understanding the HTTP and Gopher Proxies 168
 - How the HTTP, Gopher, and SSL Proxies Work 168
 - Changing the HTTP Proxy Port 169
 - Authenticated HTTP 169
 - Gopher and FTP Services 170
 - SHTTP and SSL Services 170
 - Accessing HTTP, SSL, and Gopher Proxy Configuration 171
 - Configuring the Firewall for Web and Gopher Services 173
 - Planning Web and Gopher Service Proxy Settings 173
 - Configuring Web and Gopher Service Proxy Settings 174
 - Configuring HTTP Proxy Settings 174
 - Configuring Authenticated HTTP Proxy Settings 174
 - Configuring SSL Proxy Settings 174
 - Configuring Gopher Proxy Settings 174
 - Enabling Proxy Services 175
 - Creating User Authentication Entries 175
 - Verifying Your Setup 175
 - Using Web Services 175
 - Non-Transparent Access 175
 - Authenticated HTTP 176
 - Transparent Access 176
 - Using Proxy-Aware Browsers 176
 - Configuring Web Browsers 176
 - Accessing Web Services without Authentication 178
 - Accessing Web Services with Authentication 178
 - Using Non-Proxy-Aware Browsers 179
 - Configuring Web Browsers 179
 - Accessing Web Services 179
 - Using Gopher Services 180
 - Running a Web Server 180
 - 20. **Managing X Window Services** 181
 - Understanding the X11 Proxy 181

- How the X11 Proxy Works 182
- Accessing X11 Proxy Configuration 183
- Configuring the Firewall for X11 Services 183
 - Planning X11 Proxy Settings 184
 - Configuring X11 Proxy Settings 184
 - Enabling X11 Proxy Services 185
 - Verifying Your X11 Proxy Setup 185
- Using X11 Services 185
- 21. Managing Custom Services 189**
 - Understanding the Plug Proxy 190
 - How the Plug Proxy Works 191
 - Accessing Plug Proxy Service Configuration 192
 - Configuring the Firewall for Plug Proxy Services 192
 - Planning 193
 - Configuring Plug Proxy Settings 193
 - Enabling Plug Proxy Services 195
 - Configuring Your Service 195
- 22. Managing Custom Services With Authentication 197**
 - Understanding the Circuit Proxy 198
 - How the Circuit Proxy Works 199
 - Accessing Circuit Proxy Configuration 200
 - Configuring the Firewall for Circuit Proxy Services 201
 - Planning 201
 - Configuring Circuit Proxy Settings 202
 - Enabling Circuit Proxy Services 204
 - Verifying Your Setup 204
 - Using the Circuit Proxy 204
- 23. Managing MediaBase Services 207**
 - Understanding the MediaBase Proxy 207
 - How It Works 208

Configuring the Firewall to Use the MediaBase Proxy	208
Planning	208
Configuring MediaBase Proxy Settings	209
Enabling MediaBase Services	209
Using the MediaBase Proxy	209
Verifying Your Setup	210

PART IV Managing the Firewall Environment

24. Mail Services	213
Understanding the SMTP Proxy	214
How the SMTP Proxy Works	214
Accessing SMTP Proxy Configuration	215
Configuring the Firewall for SMTP Services	216
Planning SMTP Proxy Services	216
Configuring SMTP Proxy Services	216
Enabling SMTP Proxy Services	217
Configuring Other Settings	218
Advertising the Firewall as a Mail Exchanger	218
Configuring Your Internal Mail Hub	218
Verifying Your Setup	218
Using Mail	219
Understanding the POP3 Proxy	219
How the POP3 Proxy Works	220
Accessing POP3 Proxy Configuration	220
Configuring the Firewall for POP3 Services	221
Planning	221
Configuring POP3 Proxy Settings	222
Creating User Authentication Entries	222
Enabling POP3 Proxy Services	222
Configuring Your Internal POP3 Mail Server	223
Using POP3 to Exchange Mail	223
Using the POP3 Proxy with Multiple POP3 Servers	224

- 25. **Managing Packet Screening** 225
 - Understanding Packet Screening 226
 - How Packet Screening Works 227
 - How Packet Screening Rules Work 227
 - Packet Screening Action Rules 227
 - Deny Rules 227
 - Permit Rules 228
 - Absorb Rules 228
 - Packet Screening Field Rules 229
 - Source IP Address and Destination IP Address Rules 229
 - Interface Rules 230
 - Protocol Rules 230
 - Source Port and Destination Port Rules 231
 - Order of Precedence of Packet Screening Rules 231
 - Accessing Packet Screening Configuration 232
 - Adding Packet Screening Rules 233
 - Planning Packet Screening Rules 233
 - Creating Packet Screening Rules 233
 - Loading Packet Screening Rules 235
 - Modifying Packet Screening Rules 235
 - Deleting Packet Screening Rules 235
 - Changing Order of Precedence 236
 - Verifying Your Configuration 236

- PART V**
- Managing Additional Firewall Services**
- 26. **Managing Logging and Reporting** 239
 - Understanding Logging and Reporting 239
 - Creating Logs 240
 - Configuring Logs 241
 - Configuring Proxy Logging 241
 - Configuring Log Retention Time 241

- Creating Reports 242
 - Creating Service Summary Reports 243
 - Creating Exception Reports 243
 - Creating Other Reports 244
- Configuring Reports 244
 - Accessing Report Configuration 244
 - Configuring Report Recipients 244
 - Configuring Report Frequency 245
 - Configuring Events to Ignore in Exception Reports 245
 - Security Alerts 247
- Reading Logs and Reports 248
 - Reading Logs 248
 - Reading Service Summary Reports 249
 - Reading Exception Reports 249
 - Exception Reports: The Security Alerts Section 250
 - Exception Reports: The System Warnings Section 250
 - Exception Reports: Possible Items of Interest Section 250
 - Exception Reports: Example 251
- 27. **Verifying Integrity** 253
 - Understanding System Integrity 253
 - How System Integrity Works 254
 - Configuring Integrity Checks 254
 - Accessing Integrity Checking 254
 - Configuring Files to Ignore 255
 - Creating an Integrity Database 256
 - Planning an Integrity Database 256
 - Creating the Database 256
 - Protecting the Integrity Database 257
 - Updating the Integrity Database 257
 - Updating the Database 257
 - Verifying System Integrity 257
 - Check Integrity 258
 - Viewing and Understanding the Results 258

- 28. Managing Virtual Private Networks 259**
 - Understanding VPNs 260
 - Privacy With Trust (Trusted Link) 262
 - Privacy Without Trust (Private Link) 262
 - Encryption Through Multiple Firewalls (Passthrough Link) 263
 - How Virtual Private Networks Work 264
 - Encrypting the Data 264
 - Decrypting the Data 264
 - Routing the Packet 265
 - Accessing Encryption Key Configuration 265
 - Working With Encryption Keys 265
 - Planning Encryption Keys 265
 - Creating Encryption Keys 266
 - Modifying Encryption Keys 267
 - Deleting Encryption Keys 267
 - Accessing the VPN Configuration 268
 - Working With the VPN Configuration 268
 - Planning VPNs 268
 - Creating a Trusted or Private VPN 269
 - Creating a Passthrough Link 270
 - Modifying VPNs 271
 - Deleting VPNs 272
 - Starting Your VPN 272
 - Testing Your VPN 272
 - Stopping Your VPN 273
- 29. Managing Web and Gopher Servers 275**
 - Understanding the Info Server 275
 - How the Info Server Works 276
 - How the Database Works 276
 - Info Server Directories 277
 - Info Server Data Files 277
 - Info Server Queries and Executable Programs 278
 - Info Server Gopher Menu Files 278

-
- Accessing Info Server Configuration 279
 - Configuring the Firewall to Run an Info Server 280
 - Planning an Info Server 280
 - Creating Files for an Info Server 280
 - Placing Info Server Files on the Firewall 280
 - Adding Files to the Info Server Database 281
 - Adding Text Files to the Info Server Database 281
 - Adding Binary Files to the Info Server Database 282
 - Adding Query Files to the Info Server Database 283
 - Creating Info Server Gopher Menu Files 283
 - Enabling the Info Server 283
 - Verifying Your Info Server Setup 284
 - Using the Info Server 284
 - Updating Info Server Files 284
 - 30. Managing Content Scanning 285**
 - Understanding Content Scanning 285
 - FTP 286
 - SMTP Mail 286
 - HTTP 286
 - Content Scanning Engine 287
 - How Content Scanning Works 287
 - Accessing Content Scanning Configuration 288
 - Configuring the Firewall to Use Content Scanning 289
 - Planning the Firewall Configuration for Content Scanning 290
 - Configuring and Enabling Content Scanning 290
 - Configuring the Content Scanner 291
 - 31. Managing URL Filtering 293**
 - Understanding URL Filtering 293
 - Denying and Allowing Access to Specified URLs 294
 - Filtering URL Headers 294

- Configuring URL Filtering 294
 - Planning URL Filtering 294
 - Configuring URL Filtering Settings 295
 - Enabling URL Filtering Services 296
- Understanding Cyber Patrol 297
- Configuring Cyber Patrol 299
 - Planning 299
 - Configuring Cyber Patrol Settings 300
 - Enabling Cyber Patrol Services 304
- 32. Managing the Network Management Agent 305**
 - Understanding the SNMP Agent 306
 - How the SNMP Agent Works 306
 - Accessing SNMP Agent Configuration 307
 - Configuring the Firewall as an SNMP Agent 307
 - Planning SNMP Agent Settings 308
 - Configuring SNMP Agent Settings 308
 - Enabling the SNMP Agent 308
 - Configuring SNMP Network Managers 309
 - Understanding SNMP Agent Replies 309
 - SNMP Agent and Management Information Base 309
 - SNMP Agent Community 309
 - SNMP Agent Object ID 310
 - The value provided for the mib-2.system.sysObjectID object is .1.3.6.1.4.1.602.1.1.1.5.1. 310
 - Trap 310
- 33. Login Shell 311**
 - Understanding the Login Shell Program 311
 - How the Login Shell Program Works 312

Configuring the Firewall to Use the Login Shell Program	312
Planning Remote Login	312
Enabling Remote Login	312
Adding Support for the Login Shell	313
Creating User Accounts	313
Configuring the Proxy Rules	313
Configuring the Shell	314
Creating User Authentication Records	314
Securing Other Applications	315
Verifying Your Setup	315
Using the Login Shell Program	315
Accessing the Firewall From Trusted Networks	315
Accessing the Firewall From Untrusted Networks	316
Changing Password for User Account	316

PART VI Appendices, Glossary, and Index

A. Installing and Upgrading to Gauntlet 4.1	319
Gauntlet Execution Environment Subsystems	320
Gauntlet Disk Space Requirements	321
Prerequisites for Installing Gauntlet	321
Software License	321
IRIX Software Prerequisites	322
Installing the Software	322
Files In This Release	324
Gauntlet Configuration	324
Activating Your Gauntlet License	325
A Few Definitions	325
Obtaining and Installing a Software License	326
Upgrading to Gauntlet 4.1	328
Upgrade Instructions	328
How the Upgrade Program Works	328
Upgrade Considerations	329

- B. Initializing Strong Authentication Tokens 331**
 - Access Key II Authentication 331
 - How Access Key II Authentication Works 331
 - Configuring the Access Key II 332
 - Adding an Access Key II User 332
 - Using Access Key II With the Gauntlet Firewall 333
 - Authenticating With Access Key II 333
 - Access Key II Example 334
 - CRYPTOCard RB-I 334
 - How It Works 334
 - Configuring the CRYPTOCard 335
 - Generating the Shared Secret 335
 - Initializing the CRYPTOCard 335
 - Adding a CRYPTOCard User 336
 - Using the CRYPTOCard with the Gauntlet Firewall 337
 - Using the CRYPTOCard for the First Time 337
 - Authenticating With CRYPTOCard 338
 - Changing Your CRYPTOCard PIN 338
 - CRYPTOCard Example 339
 - Digipass Authentication 339
 - How Digipass Authentication Works 339
 - Configuring the Digipass 340
 - Adding a Digipass User 340
 - Using Digipass With the Gauntlet Firewall 341
 - Authenticating With Digipass 341
 - Digipass Example 342

SafeWord Authentication Server	342
How SafeWord Authentication Works	343
Configuring the SafeWord Authentication Server	343
Configuring the Firewall for SafeWord Authentication	343
Adding a SafeWord User	344
Using SafeWord Authentication Server With the Gauntlet Firewall	345
Authenticating With SafeWord Authentication Server	345
SafeWord Authentication Server Example	346
SecurID System Authentication	346
How SecurID System Authentication Works	346
Configuring the ACE/Server	347
Configuring the Firewall	347
Adding SecurID Users	348
Adding Individual SecurID Users	348
SecurID Default Users	349
Using SecurID with the Gauntlet Firewall	349
Authenticating With SecurID	349
SecurID Example	350
S/Key System	351
How the S/Key System Works	351
Adding an S/Key User	351
Generating One-Time Access Passwords	352
Determining the Key Value	352
Generating the One-Time Passwords	353
Allowing Users to Generate One-Time Passwords	354
Using S/Key With the Gauntlet Firewall	354
Authenticating With the S/Key System	354
S/Key Example	355

- RADIUS Authentication 356
 - How RADIUS Authentication Works 356
 - Configuring the RADIUS Authentication Server 357
 - Enabling RADIUS Support 357
 - Adding a RADIUS User 358
 - Using RADIUS With the Gauntlet Firewall 358
 - Authenticating With RADIUS 359
 - RADIUS Authentication Example 359
- Reusable Passwords 360
 - How Reusable Passwords Work 360
 - Adding a User with Reusable Passwords 360
 - Using Passwords with Gauntlet 361
 - Authenticating With Reusable Passwords 361
 - Reusable Password Authentication Example 362
- Glossary 363**
- Index 373**

List of Figures

Figure 1-1	Gauntlet Internet Firewall Standard Configuration	9
Figure 1-2	Dual-Homed Bastion Host	11
Figure 2-1	Authentication Window	21
Figure 2-2	Gauntlet Firewall Manager	21
Figure 3-1	Service Groups Window	31
Figure 3-2	Add Service Groups Window	33
Figure 3-3	Rules Window	37
Figure 4-1	Destination Access Window	44
Figure 5-1	Networks Window	50
Figure 5-2	Add Network Definition Window	51
Figure 5-3	Network Groups Window	53
Figure 5-4	Add Network Group Window	55
Figure 6-1	Users Window	64
Figure 7-1	User Restrictions Window	74
Figure 9-1	FTP Window	89
Figure 10-1	LDAP Window	98
Figure 11-1	SQL Server Window	104
Figure 11-2	Add SQL Server Services Window	106
Figure 12-1	NetShow Window	111
Figure 12-2	Microsoft NetShow Player Advanced Properties Window	115
Figure 12-3	RealPlayer Proxy Preferences Window	116
Figure 12-4	StreamWorks Player Network Settings Window	117
Figure 12-5	VDOLive Setup Settings Window	119
Figure 13-1	SNMP Window	124
Figure 13-2	Add SNMP Agents Window	125
Figure 14-1	News Window	132
Figure 14-2	News Feed Server Identification Screen	132

Figure 14-3	External News Server Window	133
Figure 14-4	News Reader Window	135
Figure 15-1	Example Firewall Ip Configuration	139
Figure 15-2	Lp Window	141
Figure 15-3	Modify Printer Services Window	143
Figure 16-1	Rsh Window	147
Figure 17-1	Sybase Window	153
Figure 17-2	Modify Sybase Services Window	155
Figure 18-1	TELNET Window	160
Figure 18-2	rlogin Window	161
Figure 18-3	TELNET Connect Window	163
Figure 19-1	HTTP Window	171
Figure 19-2	SSL Window	172
Figure 19-3	Gopher Window	173
Figure 19-4	Browser Settings Window	177
Figure 19-5	Network Password Window	178
Figure 19-6	User Password Window	179
Figure 20-1	X Proxy Configuration Window	183
Figure 20-2	Modify TELNET Proxy Configuration Window	184
Figure 20-3	X Status Window	187
Figure 20-4	X Connection Confirmation Window	187
Figure 21-1	Plug Configuration Window	192
Figure 21-2	Add Plug Services Window	194
Figure 22-1	Circuit Configuration Window	200
Figure 22-2	Circuit Service Definition Window	202
Figure 22-3	Add Circuit Gateway Configuration Window	203
Figure 24-1	Mail Window	215
Figure 24-2	POP3 Mail and Proxy Configuration Window	221
Figure 24-3	Eudora Pro Configuration for APOP	224
Figure 25-1	Packet Screening Window	232
Figure 25-2	Add Packet Screening Rule Window	234
Figure 26-1	Configure Window	242
Figure 26-2	Items of Interest Window	246

Figure 26-3	Security Alerts Window	247
Figure 27-1	Integrity Check Window	255
Figure 28-1	Example VPN	261
Figure 28-2	Add Swipe Key Window	266
Figure 28-3	VPNs Window	268
Figure 28-4	Add VPN Link Configuration Window	269
Figure 29-1	Info Window	279
Figure 30-1	Mail Window	289
Figure 31-1	URL Filtering Window	295
Figure 31-2	Add HTTP Services Window	300
Figure 31-3	Cyber Patrol Configuration Window	301
Figure 31-4	Cyber Patrol Renewal/Registration Window	302
Figure 31-5	Work Time Range Configuration Window	303
Figure 32-1	SNMP Agent Window	307
Figure B-1	One-Time Password Window	355

About This Guide

Connecting your private, internal network to an outside, untrusted network can be both an asset and a liability. It is an asset because you can exchange computerized information with a variety of organizations. It can be a liability because you may be exposing your network resources to unwanted probing and spying. The Gauntlet firewall is an important component in a well-designed network security structure to combat these threats.

This introduction gives some overview information and also discusses “How to Get Latest Security Patches” on page xxxii.

Audience for This Guide

This guide is intended for firewall administrators. It assumes that you are familiar with IRIX system and networking administration and with basic firewall concepts. System administrators should be familiar with TCP/IP, domain name service, *sendmail*, and router configuration. Consult your local library, bookstore, network resources, and IRIX administrator for additional references.

Contents of This Guide

This guide consists of six parts:

- Part I, “Introducing the Gauntlet Firewall,” presents the initial information about the firewall and firewall administration.
- Part II, “Managing the Gauntlet Firewall,” describes how to manage the components of the network.
- Part III, “Configuring and Using Proxy Services,” describes how to configure the various proxies.

- Part IV, “Managing the Firewall Environment,” describes how to manage those services within the firewall environment.
- Part V, “Managing Additional Firewall Services,” describes the various services that the Gauntlet firewall provides.
- Part VI, “Appendices, Glossary, and Index,” provides supplementary information.

Conventions Used in This Guide

These type conventions and symbols are used in this guide:

Bold—keywords and command line options.

Italics—executable names, filenames, IRIX commands, manual/book titles, new terms, utilities, variable command-line arguments, and variables to be supplied by the user in examples, code, and syntax statements.

`Fixed-width type`—Code examples, prompts, and onscreen text.

Bold fixed-width type—User input, including keyboard keys, printing and nonprinting (see also <>).

Additional Documentation

Refer to the following documentation for additional information about the Gauntlet Firewall product:

- Check the release notes for the most recent information and software and hardware requirements.
- *Gauntlet Netperm Table Reference Guide* (part number 007-3822-003) describes how to edit the netperm table using the command-line interface.

Additional Resources

The collection of resources in this section is presented for your information only. It is not an endorsement of any of the products or organizations.

Books

Building Internet Firewalls. Chapman, D. Brent & Zwicky, Elizabeth. O'Reilly & Associates, Inc. ISBN 1-56592-124-0.

Firewalls and Internet Security: Repelling the Wily Hacker. Cheswick, Steven M. & Bellovin, William R. Addison Wesley. ISBN 0-201-63357-4.

Newsgroups

comp.security.firewalls—Discussions of anything regarding network security firewalls.

Mailing Lists

The Firewalls mailing list is for discussions of Internet firewall security systems and related issues. Relevant topics include the design, construction, operation, maintenance, and philosophy of Internet firewall security systems.

To subscribe to the regular mailing list, send the following command in the body of an e-mail message (*not* on the "Subject:" line!) to majordomo@greatcircle.com:

```
subscribe firewalls
```

To subscribe to the digest version of the mailing list, send the following command in the body of an email message (*not* on the "Subject:" line!) to majordomo@greatcircle.com:

```
subscribe firewalls-digest
```

Frequently Asked Questions Lists

The Internet Firewalls Frequently Asked Questions list is maintained by Marcus J. Ranum and located at:

<http://www.clark.net/pub/mjr/pubs/fwfaq/index.html>

White Papers

Application Gateways and Stateful Packet Filters

<http://www.nai.com/products/security/prodserv/gauntlet/firewallcomp.asp>

Firewalls Are Not Enough

<http://www.nai.com/products/security/prodserv/gauntlet/FirewallsNotEnough.asp>

Thinking About Firewalls

<http://www.nai.com/products/security/prodserv/gauntlet/fwovervw/index.asp>

How to Get Latest Security Patches

The CD-ROM containing the Gauntlet firewall software contains necessary security patches (if any) at the time of product release, so be sure to install those patches. Stay in touch with the WWW site for SGI Security Headquarters at <http://www.sgi.com/Support/Secur/security.html> for new security patches and security advisories. Be sure to install any security patches that replace patches found on your CD-ROM.

PART ONE

Introducing the Gauntlet Firewall

Chapter 1

Understanding the Gauntlet Firewall

Chapter 2

Managing the Gauntlet Firewall

Understanding the Gauntlet Firewall

A Gauntlet Firewall is a hardware- and software-based firewall system. It provides secure access and internetwork communications between private networks and public networks (such as the Internet), and between subnets of private networks. It offers application-level security services for both incoming and outgoing communications based on existing security practices or an organization's security policies.

This chapter provides an overview of the Gauntlet Firewall and how it works. It is not a thorough discussion of firewalls or security practices. Refer to "Additional Resources" on page xxx for a list of other resources that provide excellent introductory and advanced discussions of firewalls. The chapter contains these sections:

- "Understanding Gauntlet Firewall Concepts" on page 1
- "Understanding Gauntlet Firewall Components" on page 5
- "How the Gauntlet Firewall Works" on page 8

Understanding Gauntlet Firewall Concepts

Simply put, a firewall is a single point of defense that protects one side of a network from the other side. This usually means protecting your company's private network from other networks to which it is connected, such as the Internet. Firewalls can be as simple as a router that filters packets or as complex as a multisystem, multirouter solution that combines packet filtering with application gateways.

Gauntlet Firewall Design Philosophy

The Gauntlet Firewall follows this paradigm:

That which is not expressly permitted is prohibited.

The firewall will allow only those activities that are explicitly set, either through system defaults or through your own configurations. New services don't get through the firewall unless you allow them through. You must be able to identify and remove any "backdoors" you have left in place that no longer match your security policy.

Recognizing that most security breaches occur through a compromised user account, the Gauntlet Firewall generally has no user accounts. While you can set up an administrator account, users do not need to log into the firewall to access information on the other side of your firewall.

The Gauntlet Firewall is auditable, controllable, and configurable. It logs many activities and processes.

Establishing a Security Perimeter

Establishing a network security perimeter involves designating a network of systems you wish to protect and defining the mechanisms used to protect them. The firewall is the communications gateway for all hosts inside the perimeter. To have a successful network security perimeter, *all* communications from outside the security perimeter to hosts inside the perimeter must pass through the firewall.

Trusted and Untrusted Networks

Your firewall must be configured to differentiate between the "good guys" and the "bad guys." The firewall makes this determination using information you provide about different networks. It understands three types of networks: trusted, untrusted, and unknown.

Trusted Networks

Trusted networks are the networks inside your security perimeter. Trusted networks are usually the ones you are trying to protect. Often, you or someone in your organization administers the systems on these networks. Your organization controls the security measures for these networks. Usually, they are within the physical security perimeter. They can also be connected in a virtual private network, as explained in Chapter 28, "Managing Virtual Private Networks" on page 259.

When you set up the firewall, you explicitly configure the networks your firewall can trust. After initial configuration, the trusted networks usually include the firewall itself and all networks inside your security perimeter.

Untrusted Networks

Untrusted networks are the networks outside your security perimeter. They are untrusted because they are usually outside your control or knowledge. You have no control over the administration or security policies for these sites. They are the ones from which you are trying to protect your network. However, you still need and want to communicate with these networks, even though they are untrusted.

When you set up the firewall, you explicitly configure the networks from which your firewall can accept requests, but which it does not trust. By default, after initial configuration, the untrusted networks are all networks outside the perimeter.

The firewall uses different service groups for requests from untrusted networks than it does for requests from trusted networks. For example, the default configuration allows HTTP requests from trusted networks but not from untrusted networks. For some types of requests (including TELNET, FTP, rlogin, rsh, HTTP, and POP3), the firewall may use additional authentication before processing the request.

Unknown Networks

Unknown networks are those networks that are neither trusted nor untrusted. They are unknown quantities to the firewall because you have not explicitly told the firewall that this network is a trusted or an untrusted network. By default, the list of untrusted networks contains all networks that are not trusted (*). In that case, there are no unknown networks. If you change the list of untrusted networks to contain some specific networks, all networks not in the trusted or untrusted list become unknown.

Consider a company that lists its own networks as the trusted network. The company lists the networks for three clients as the untrusted networks. Every other network on the Internet is now an unknown network and cannot pass requests through the firewall.

Service Groups

Your security policy provides a list of allowed and prohibited activities. Similarly, the Gauntlet Firewall uses service groups to specify allowed activities. Service groups are

collections of rules about what the firewall can and cannot do in particular situations. Service groups indicate which proxies can run, and whether they require authentication, special logging, or other general settings. The service groups you create and rules you assign should be based on your site's security policies.

By default, the Gauntlet Firewall includes one service group for requests from trusted networks and one for requests from untrusted networks. The firewall uses the source IP address of the request to determine which service group to use.

The default service group for trusted networks allows for a variety of services (including TELNET, FTP, NNTP, and HTTP). It does not require users to authenticate. The default service group for untrusted networks allows for only a few services (including FTP, TELNET, NNTP, and POP3). It requires users to authenticate before accessing hosts inside your security perimeter. You can, of course, modify these default service groups to match your security policy.

Transparency

Transparency means that your firewall is not visible to your users as they work. They can continue to use TELNET, for example, to access client sites without having to explicitly connect to the firewall.

The default Gauntlet Firewall configuration implements transparency inside your firewall for your trusted networks. This is accomplished by creating default routes that send all requests to untrusted networks through the firewall and by certain configuration options on the firewall.

In contrast, the firewall does not implement transparency for requests from untrusted networks. In this case, you want users outside your security perimeter to be aware that they are entering your site through your firewall.

The advantage of transparent access is that you do not need to reconfigure your client systems or learn new procedures in order to use supported services. Nontransparent access is supported, but users must learn procedures to perform their tasks.

Understanding Gauntlet Firewall Components

The Gauntlet Firewall uses hardware and software to protect your network.

Gauntlet Firewall Hardware

The specific hardware components of the Gauntlet Firewall are the network interface cards. These cards physically separate networks from one another.

Gauntlet Firewall Software

The software components of the Gauntlet Firewall include a hardened operating system, application-level security services, security programs, and other management utilities.

Gauntlet Firewall Operating System

The operating system is a version of the standard Silicon Graphic's IRIX operating system, "hardened" by the Gauntlet software. As part of the firewall, the operating system has been tailored to provide support for only the services necessary to run the firewall. For example, the operating system on the firewall does not support IP packet forwarding, source routed packets, or ICMP redirects. These services change the directions that packets flow, and could direct networks to circumvent the firewall. Services such as NFS, NIS, and RPC cannot easily be made secure and so should be disabled.

Unsupported network services do not just report an error to the requesting site. The operating system logs these access attempts, providing information about such probes.

Application Level Security Services (Proxies)

The software on the Gauntlet Firewall includes security services on a per-application basis. As noted above, all packets, and therefore all application requests go to the firewall. On the firewall, proxy software relays information from one side of the firewall to the other. The proxy prevents the applications on outside networks from talking directly with the applications on your inside network, and vice versa. No IP packets pass from one side of the firewall to the other. All data is passed at the application level. (The "trusted ports" feature in this implementation is an exception to this generalization.)

Each application generally talks through a different proxy that understands the protocol for that application. Currently, the Gauntlet Firewall includes proxies for the following types of services:

- Terminal services (rlogin and TELNET)
- Electronic mail (POP3 and SMTP)
- File transfer services (FTP)
- Remote execution (rsh)
- Web services (HTTP, SHTTP)
- Gopher services (Gopher, Gopher+)
- X Window services (X11)
- Printing services (lp)
- SQL services (Microsoft SQL Server and Sybase SQL Server)
- Multimedia services (MediaBase, NetShow, RealAudio/RealVideo, StreamWorks, VDOLive)
- Network management services (SNMP)
- Certificate management (LDAP)

In addition, the Gauntlet Firewall includes a generic plug-board proxy. This proxy patches TCP traffic from a particular port on one side of the firewall to a particular TCP port on another system on the other side of the firewall. As with the service-specific proxies, no IP packets pass from one side of the firewall to the other. If you have not installed a proxy for a service, that type of traffic does not pass through the firewall.

The Gauntlet Firewall also includes an authenticating circuit proxy. This proxy also patches TCP traffic from a particular port on one side of the firewall to a particular TCP port on another system on the other side of the firewall. The circuit proxy requires users to authenticate first; the plug proxy does not.

The Gauntlet Firewall includes configured versions of the plug proxy for:

- AOL
- CompuServe
- finger
- LDAP

- Lotus Notes
- Usenet news (NNTP)
- Web services (SSL)
- whois

Because the non-plugin proxies use the same protocols to communicate as the applications, you do not need to modify the original client or server applications. For example, when the TELNET application connects to the firewall, the application follows the standard TELNET protocol in RFCs 764 and 854. You can continue to use the same TELNET application to connect to remote sites.

All of the proxies are configurable. You can accept or reject requests to or from certain sites and networks, or set up other rules the proxies use when passing requests through the firewall. You can also enable or disable individual proxies and run only the ones you need. You can easily translate your security policies into configuration rules.

The proxies log all activities to and through the firewall. You can use the logs to gather usage statistics or to look for potential attacks.

In addition, several of the proxies support strong user authentication systems. These one-time passwords or security token systems provide additional security because users use a different password each time they access the network.

IP Screening Utility

The Gauntlet Firewall includes additional security software in the form of an IP screening utility. This feature checks IP packets based on several criteria (for example, address and protocol) and processes or rejects the packets. It detects spoofed packets claiming to be from one network that are actually from another network.

Using the IP screening utility, you can configure the firewall so that the firewall is transparent to your users for most activities. The IP screening facility provides a method for permitting high bandwidth or unsupported protocols in situations where the security requirements are not as stringent as with an Internet firewall.

Management Utilities

The Gauntlet Firewall also contains several programs that ease the job of administering the firewall. These include management tools for configuring the firewall, and scripts for reporting activity through the firewall and performing general administration.

The graphical Gauntlet Firewall Manager utility provides an easy-to-use interface for performing most standard configuration activities. You do not need to modify system files or configuration files unless you want to customize uncommon features of your configuration.

The Gauntlet Firewall also includes shell scripts that assist in upgrading, creating backups, checking integrity, and other general administrative tasks.

How the Gauntlet Firewall Works

Consider a company, Yoyodyne, that has a connection to the Internet via an Internet service provider (ISP). They have installed a Gauntlet Firewall to protect their corporate network (yoyodyne.com) from every other host on the Internet. They are using the configuration shown in Figure 1-1.




Figure 1-1 Gauntlet Internet Firewall Standard Configuration

The Yoyodyne network is first separated from the rest of the Internet by a router. The router passes traffic from the Internet to the Gauntlet firewall only when that traffic is bound for some part of the Yoyodyne internal network. More sophisticated routers can

additionally strengthen a company's security perimeter by implementing certain security functions such as "IP spoofing filters."

The firewall is helping to establish a security perimeter to protect the internal (trusted) network. It screens all requests that need to pass from one side of the firewall to the other. Using rules Yoyodyne created based on their security policies, the firewall determines whether to accept or pass requests through (at the application or TCP level) to the other side.

Dual-Homed Bastion Host

In order to protect the inside network, the firewall must be able to see all of the packets intended for hosts on the inside network. While there are a number of ways to physically and logically accomplish this, the recommended configuration is the firewall system installed as a dual-homed bastion host.

As a dual-homed bastion host, the firewall system has two network interface cards, thus two connections—one to your network and one to the outside, as shown in Figure 1-2.




Figure 1-2 Dual-Homed Bastion Host

All outside network traffic enters and exits the firewall through one network interface, such as ec0. All inside network traffic enters and exits through a second network interface, such as ec1. To accomplish this, each interface has a separate IP address.

Yoyodyne was assigned the 204.254.155 network, and chose 204.254.155.253 as the outside IP address; it chose 10.0.1.253 for the inside IP address.

Note: You can also use two firewalls to create a virtual private network (or a virtual network perimeter), exchanging encrypted information across an untrusted network, such as the Internet. Because of United States government export regulations, this feature is generally not available outside the United States and Canada. Refer to Chapter 28, “Managing Virtual Private Networks” on page 259, for information about the encryption features of the Gauntlet Firewall.

Processing Packets and Requests

The firewall follows a standard set of steps for the packets it receives on either interface. The steps are discussed in these sections:

1. “Receiving Packets” on page 12
2. “Checking Source and Destination” on page 13
3. “Checking the Request Type” on page 13
4. “Calling the Appropriate Program” on page 13
5. “Processing the Request” on page 14

As we examine each step of the process, consider a Yoyodyne employee working at a client site (outside the security perimeter) who needs access to their system at work via TELNET.

Receiving Packets

Routing information on outside hosts and at the ISP directs all requests to hosts on the inside network to the firewall. In addition, the domain name system (DNS) on the firewall and other outside DNS servers advertise the outside IP address of the firewall as the *only* way to connect to anything on the inside network. Hosts on the inside network use routing information to direct all requests for outside networks to the inside address of the firewall.

For example, the client company systems consult their routing information and pass the TELNET request along until it reaches the Yoyodyne firewall.

Checking Source and Destination

Once the firewall receives a packet it must determine what to do with it. First, the operating system kernel examines the destination of the packet and determines whether the packet is destined for a host inside the firewall. The firewall accepts the packet and redirects it to the appropriate proxy. If there is no proxy configured to accept a packet, the firewall drops the packet and logs the failed access.

Next, the firewall examines the source address of the packet and the interface on which it received the packet. This process verifies the information against configuration tables, and prevents the firewall from accepting IP spoofed packets. If this check indicates that the current request could not possibly have come in through the associated interface, the firewall rejects the packet and logs it. For example, if the Yoyodyne firewall receives a packet on ec0 (the outside interface) claiming to be from 10.0.1.10 (an inside address), the firewall rejects the packet and logs the spoof attempt.

In our TELNET example, the destination of the packet is the firewall. The firewall receives a request on ec0, the outside interface. The address does not indicate that it came from an inside network. The firewall accepts the packet for local delivery and processing.

Checking the Request Type

Now that the firewall knows to deliver the packet locally, it looks at the contents of the packet. The operating system checks various tables on the firewall to determine if it offers the requested service on the requested port. If it does not, it logs the attempt as a potential security alert and rejects the request.

In our TELNET example, the packet indicates that it is a TELNET request on TCP port 23. The configuration tables indicate that the firewall supports this type of service.

Calling the Appropriate Program

Now that the firewall knows to offer the requested service, the operating system uses other configuration information to forward the request to the appropriate daemon. In our TELNET example, the firewall forwards the request to the TELNET proxy, which processes the TELNET request.

Processing the Request

The proxy now processes the request. It first checks its configuration information. The proxy determines how to handle the request based on the source (IP address) of the request. By default, it uses one service group (set of rules) for trusted networks and another service group for untrusted networks.

Once configured, the proxy processes the request as the standard application would. The proxies follow the same protocols and handshakes as indicated in the RFCs or other documents. Requesting applications think they are talking to an actual server, not a proxy.

The proxies also check whether the request is permitted for the destination. For some services, the proxies can perform the additional step of authenticating the user. This verification provides additional assurance that users really are who they say they are. The proxies then pass each request to the appropriate program on the other side of the firewall using the standard protocol for that service.

In our TELNET example, the TELNET proxy uses the generic untrusted service group because the request came from an outside network. The untrusted service group and rule permits TELNET to access systems on the inside network, but requires authentication. The firewall prompts the user to authenticate. Once the user authenticates, the proxy provides a command-line interface allowing the user to indicate the internal system to which they wish to connect. The proxy then uses the standard TELNET protocol to pass packets back and forth between the host on the outside network and the host on the inside network.

Managing the Gauntlet Firewall

Your Gauntlet Firewall is an integral part of your corporate network. Follow the same maintenance and management practices for the firewall as you do for any other mission-critical application or machinery. You also need to manage the Gauntlet specific elements of the firewall. The Gauntlet Firewall includes a graphical user interface, called the Gauntlet Firewall Manager, that allows you to configure the firewall easily.

This chapter explores considerations for managing your Gauntlet Firewall, and explains how to use the Gauntlet Firewall Manager in the following sections:

- “General Management and Maintenance” on page 15
- “Understanding the Gauntlet Firewall Manager” on page 17
- “Using the Gauntlet Firewall Manager” on page 18
- “Getting Help” on page 24

General Management and Maintenance

Consider the Gauntlet Firewall to be another IRIX system you must maintain. For most activities, you can continue to use the tools you are familiar with. The exceptions are noted in the sections that follow:

- “Creating User Accounts” on page 15
- “Backing Up and Restoring” on page 16
- “Managing Gauntlet Firewall Options” on page 16

Creating User Accounts

If you need to create user accounts on the firewall for yourself or another administrator, work with the tools you would normally use, such as the *addUserAccount* command or the System Manager from the Toolchest. Remember that you only need to create one

account on the firewall: the account for the administrator. You do not need to create any user accounts on your firewall.

Consider using the login shell included with the Gauntlet Firewall as the shell for the administrator accounts you create on the firewall. This login shell requires that you configure a user in the Gauntlet user authentication system, allowing users to use strong authentication to log in.

Backing Up and Restoring

Back up your firewall. In most cases, you can use the same backup scheme for your firewall as you use for other IRIX systems on your network.

Remember that the firewall includes a hardened version of the IRIX operating system. You cannot use NFS to make the drives on the firewall visible to a backup system on another system.

Managing Gauntlet Firewall Options

To manage Gauntlet Firewall settings you must use the tools included with the firewall. These include a graphical and a text-based interface.

The graphical interface, the Gauntlet Firewall Manager, is the recommended interface and allows you to configure your firewall easily. The Gauntlet Firewall Manager provides flexibility by allowing you to manage your firewall from a variety of locations. The ability to configure your firewall from remote locations is especially useful when your office is located across campus, across town, or across the country from your firewall.

The text-based interface allows you to configure most features of the Gauntlet Firewall. However, the Gauntlet Firewall Manager is still the recommended interface. Consider using the text-based management interface only if your security policy requires you to manage the firewall from the console. You may also need to modify configuration files manually if you have a particularly complex configuration, such as when using four network interface cards.

Note: Do not attempt to use both management interfaces at the same time. You risk overwriting changes you made with one interface with changes you made with the other interface. Take particular care if you have multiple firewall administrators. Be sure that none of your fellow administrators is using the text-mode interface when you use the Gauntlet Firewall Manager.

Understanding the Gauntlet Firewall Manager

The Gauntlet Firewall Manager allows you to configure your firewall quickly and easily. Because of its graphical nature, you can easily view your current configuration. The Gauntlet Firewall Manager consists of two parts: the interface and the firewall server.

Firewall Manager Graphical Interface

The graphical interface to Gauntlet that you see in your Web browser is an applet written in Java. You do not need to install any additional software on your client to use the Gauntlet Firewall Manager; your Web browser includes the necessary code to run the applet.

Firewall Server

The server on the firewall that serves the graphical interface is a modified version of the Gauntlet Information Server. This server is written with security in mind. The server portion of the Gauntlet Firewall Manager runs as a daemon listening for requests on TCP port 21000. When the firewall receives requests for services on this port, the server checks its configuration information and determines whether the initiating host has permission to use the Gauntlet Firewall Manager as a server. If the host does not have permission, the server logs the connection and displays the error message `Unauthorized to use gateway`.

If the host has permission, the server displays a user authentication page. The server authenticates the user. If the user does not provide the proper authentication, the server logs the connection and displays an error message.

If the user provides the proper authentication, the server stops listening on port 21000 and begins listening on a random port. The server now displays the interface.

The server portion of the Gauntlet Firewall Manager implements your configuration by taking the values you enter into the interface and placing them into the appropriate configuration files.

Using the Gauntlet Firewall Manager

Using the Gauntlet Firewall Manager involves planning, configuring the firewall, configuring your system, configuring your Web browser, and accessing the Gauntlet Firewall Manager.

Planning the Firewall

To plan the firewall:

1. Choose the host from which you will manage the firewall. You can access the Gauntlet Firewall Manager from any qualified host running a supported Web browser. These hosts include:
 - a host on your inside network
 - a host on your inside network using PC Extender for Windows 95
 - a host on your outside network using PC Extender for Windows 95

Carefully consider your choice of the host from which you manage the firewall. Remember that, by default, the traffic between the Web browser and the firewall is *not* encrypted.

2. Consider whether the host from which you manage the firewall will be a dedicated management workstation. With a dedicated management host, you can physically secure the system to help avoid unauthorized access to the firewall.
3. Consider whether you wish to use PC Extender for Windows 95 on the host from which you manage the firewall. PC Extender helps protect the confidentiality on the link. Refer to the PC Extender documentation for more information on the benefits of PC Extender.

Configuring the Firewall

You must tell the firewall which hosts can use the Gauntlet Firewall Manager. To configure the firewall:

1. Log in to the firewall and become root.
2. Start the text-based Gauntlet administrative tool:

```
# /usr/local/etc/gauntlet-admin
```
3. Select "Fast Setup for GUI admin tool".
4. Enter the information about your inside network interface. You can use the Tab key and the Up and Down arrows to navigate in the text-based interface.
5. Select "Next Screen" to move to the next screen.
6. Enter the IP address of the host or hosts from which you will manage your Gauntlet firewall. The wildcard * is valid, so for example you could enter 10.0.1.*.
7. Select "User Authentication Management".
8. Select "Add New User".
9. Configure a user with a user ID of fwadmin and select "Save these Changes".
10. Select "Quit", then "Return" to return to the Main Menu of the Gauntlet administrative interface.
11. Select "Update Configuration Menus".
12. Select "Quit and Update Configuration Database". When the administrative interface asks you if you wish to rebuild system. configuration files, type "y".

The Gauntlet Firewall Manager daemon starts when you boot the firewall. You do not need to explicitly start the Gauntlet Firewall Manager.

Configuring Your System

To use the Gauntlet Firewall Manager, you must be using certain configuration settings on your system.

To configure your system:

1. Make sure the color setting for your windowing display is set to at least 16-bit color (referred to as 65,536 color in the X Window System).
2. Make sure the screen resolution is set to at least 800 by 600 pixels.

Configuring Your Web Browser

To use the Gauntlet Firewall Manager, you must configure your Web browser as follows:

1. Make sure you have a browser which supports Java. Recommended browsers are Netscape Communicator version 4.0 or later, and Internet Explorer version 4.0 or later.
2. Add the IP address of the internal interface of your firewall to the list of hosts for which the browser should not use a proxy. This ensures that your browser doesn't try to send your requests to the HTTP proxy on the firewall.
3. Enable Java access. The Gauntlet Firewall Manager is written in Java. Therefore, your Web browser must support Java.

Accessing the Gauntlet Firewall Manager

To access the Gauntlet Firewall Manager:

1. Open the following URL:
`http://firewall:21000/auth/gui.html`
firewall is the hostname or IP address of the inside interface of the firewall.
2. Authenticate using the account you created in the text interface.



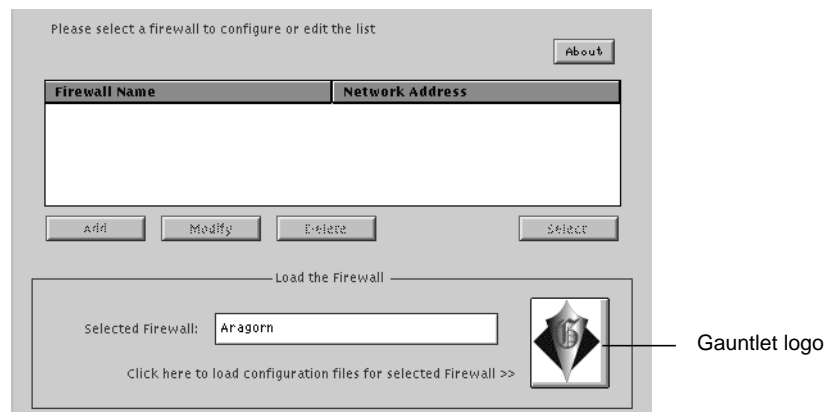
Please enter your username below

Username

Password:

Figure 2-1 Authentication Window


3. Wait as the Web browser starts the Gauntlet Firewall Manager. This process can take several minutes. Note that the Netscape Navigator logo is not animated while the Web browser is loading the configuration.



Please select a firewall to configure or edit the list About

Firewall Name	Network Address

— Load the Firewall —

Selected Firewall: 

Click here to load configuration files for selected Firewall >>

Gauntlet logo

Figure 2-2 Gauntlet Firewall Manager

4. Verify that the hostname or IP address of your firewall is shown as the Selected Firewall.
5. Click the Gauntlet logo to load the configuration settings for your firewall.

This loads the Gauntlet Firewall Manager with your configuration.

6. Minimize your Web browser.

Note: Do not close your browser, because this also closes the Gauntlet Firewall Manager.

Exiting the Gauntlet Firewall Manager

Note: Make sure you exit the Gauntlet Firewall Manager properly (using the *Exit* button in the upper left corner of the Manager window, not the Close or Exit functions of your web browser). If you do not exit properly, you may lose any changes you have made and your Firewall Server may need to be restarted.

To exit the Gauntlet Firewall Manager:

1. Click Exit.
2. If you have made any changes to your configuration you have not saved, answer the Gauntlet Firewall Manager prompt before exiting.
3. If you have not made any changes to your configuration, click Quit.

Saving Your Changes

The Gauntlet Firewall Manager allows you to choose when your changes take effect. When you save your changes, you have several options:

- Quit
- Save
- Save and Apply
- Save, Apply, and Reboot
- Reboot
- Cancel

Quit Menu Command

Quit exits from the Gauntlet Firewall Manager without saving any changes that you have made.

Save Menu Command

The Save command writes the changes you have made to the Gauntlet Firewall. However, when you save your changes, the Gauntlet Firewall Manager does not make your configuration changes take effect. In other words, the next time you use the Gauntlet Firewall Manager, you will see your changes reflected in the screens and setting. However, the firewall will not use these changes until you apply them.

For example, suppose you make a number of changes to the firewall's configuration. However, you do not want to reboot your firewall in the middle of the day while there are many users using the services of the firewall. You save your changes and exit the Gauntlet Firewall Manager. After your office has closed, you start the Gauntlet Firewall Manager to displays your changes. You apply your changes and reboot the firewall to make your configuration take effect. This delayed application of changes is also useful if you are making changes and need to leave your workstation to attend a meeting or go to lunch.

Save and Apply Menu Command

Choosing Save and Apply writes the changes you have made to the Gauntlet Firewall and makes them take effect. Most types of changes take effect immediately when you save and apply. For example, a destination access rule takes effect as soon as you choose Save and Apply from the Gauntlet Firewall Manager Exit window.

Save, Apply, and Reboot Menu Command

Choosing Save, Apply, and Reboot saves your changes, and makes them take effect immediately by rebooting the firewall and restarting all services. You must reboot your firewall to make some changes take effect, including:

- Enabling or disabling proxy services
- Starting or stopping virtual private networks (VPNs)
- Making changes to an interface
- Adding or deleting packet screening rules

Reboot Menu Command

Choosing Reboot reboots the firewall *without* saving or applying any changes that you have made.

Cancel

Cancel returns you to the Gauntlet Firewall Manager.

Getting Help

The Gauntlet Firewall Manager includes an online help system. This system provides information about all of the screens and options in the Gauntlet Firewall Manager.

The Gauntlet Firewall Manager displays the online help in a separate browser screen.

- To access the main help page, click the *Help* button from within the Gauntlet Firewall Manager.
- To access help for a specific screen, click the question mark in the lower right corner on the screen where you need help.

PART TWO

Managing the Gauntlet Firewall

Chapter 3

Service Groups and Service Group Rules

Chapter 4

Destination Access

Chapter 5

Networks and Network Groups

Chapter 6

Users and User Groups

Chapter 7

User Restrictions

Service Groups and Service Group Rules

Your company uses many different applications to accomplish tasks, and the Gauntlet Firewall provides proxy services for these applications. You want to be able to create rules that allow certain groups to use certain services. You could create rules for each service you are offering. However, creating rules for each service would be quite a time consuming and detailed effort. The Gauntlet Firewall includes the idea of service groups, which allow you to create rules for a set of services.

The following sections explain the concepts of service groups and service group rules and describe how to configure them:

- “Understanding Service Groups” on page 27
- “Accessing Service Group Configuration” on page 31
- “Configuring Service Groups” on page 32
- “Understanding Service Group Rules” on page 35
- “Adding Service Group Rules” on page 37
- “Accessing Service Group Rules Configuration” on page 36
- “Modifying Service Group Rules” on page 38
- “Deleting Service Group Rules” on page 38
- “Changing Order of Precedence” on page 38

Understanding Service Groups

Data traffic passing through the Gauntlet Firewall is generally handled through the use of services (also known as proxies). Services are applications running on the firewall that take a data stream from one interface and pass it to another interface. Most communications through the firewall occur through services. Each service is designed to handle one specific TCP or UDP protocol or port number. All services log the service

requests and the results of those requests. Some services provide blocking of certain data elements, such as Java or ActiveX.

The firewall determines whether to permit or deny traffic on the basis of the source address and the type of protocol used. When a host system attempts to communicate across the firewall using a specific protocol such as TELNET or FTP, the firewall uses the source address of the client host to determine what services have been specified for this address or the network groups to which it may belong. The firewall applies sets of rules to decide:

- whether the service is permitted or denied for this host (for example, the TELNET proxy service for the TELNET protocol)
- further requirements, such as authentication or permitted destinations

Instead of defining a group of individual services for one set of hosts and then defining the same group of services for another set of hosts, you can create a service group. A service group is a collection of services that are defined as a unit and that can be applied against a network group or specific hosts.

The service group can specify:

- Which services are available (for example, allow TELNET, FTP, and rlogin proxy services)
- What destination hosts are permitted or denied (for example, deny access to all hosts in the bigu.edu domain)
- Authentication requirements (for example, require authentication)

Default Service Groups

The initial Gauntlet configuration defines two service groups: a trusted service group and an untrusted service group.

Trusted Service Group

By default, the trusted service group is applied to that network group made up of the trusted networks, which are usually inside your security perimeter. This service group:

- Allows requests to be sent to any destination
- Permits access to some of the more commonly used proxies: finger, FTP, Gopher, HTTP, aHTTP, the Info Server, LDAP, lp, NetShow, NNTP, POP3, RealAudio/RealVideo, Rlogin, SSL, TELNET, VDOLive and whois.
- Does not generally require users to authenticate, but simply passes their requests through the firewall
- Allows users to change their passwords for strong authentication systems

Untrusted Service Group

By default, the untrusted service group is applied to that network group made up of the outside (untrusted) networks. This service group:

- Allows requests to be sent to any destination
- Permits access to a restricted group of protocols: TELNET, rlogin, FTP, NNTP, the Info Server, and POP3
- Requires users to authenticate with the authentication server that is on the firewall

Notice that the untrusted service group does not allow access to the HTTP protocol. You do not want outside users, especially people all over the Internet, using the HTTP protocol to gain access to Web servers on your internal network.

In addition, the untrusted service group does not allow users to change their passwords for strong authentication systems or allow access to the Gauntlet Firewall Manager. Because of the widespread use of packet sniffers and a multitude of other hacker tools, configuring the firewall and typing reusable passwords are extremely risky activities when done from an untrusted network. You run the risk of exposing this information to everyone and severely compromising your network security.

Service Group Membership

Services that are included within a service group are members of that group. Services that are not included within a service group are non-members. For example, by default, the HTTP service is a member of the trusted service group; this allows trusted hosts to use the HTTP proxy to access Web servers on an outside network. The HTTP service is a non-member of the untrusted service group; this prohibits untrusted hosts on an outside network from using the service to access a Web server on an inside network.

When to Add New Service Groups

Define a new service group whenever it is convenient to specify a set of general services to govern a single host or group of hosts. Often, changing business conditions or new requirements necessitate changes to the flow of information to or from an organization.

For example, until recently, the hosts in the research department at Big University have been governed by the default trusted service group and rules, which allows access to any destination on the outside networks. The research department has been providing advance research information to a group within Yoyodyne Corporation. Because of a contract dispute, the management at Big University has decided that the research department cannot give any more information to Yoyodyne. To ensure compliance, the firewall must not allow hosts in the research department to make network connections to any systems at Yoyodyne.

This example situation is an appropriate time to create a new service group. To implement the new service group, the firewall administrator at Big University:

- Creates a new service group called “no-yoyo”
- Adds all of the same services to no-yoyo as are in the trusted service group
- Adds a condition to the no-yoyo group that denies access to the destination yoyodyne.com
- Removes the service group rule that makes hosts in the research department use the trusted service group
- Adds a service group rule that makes hosts in the research department use the no-yoyo service group

The hosts in the research department continue to operate under the same services as the rest of the university, except they are now prohibited from accessing the yoyodyne.com systems.

In the future, when management at Big University determines that the hosts in the math department need to be prohibited from accessing yoyodyne.com as well, the firewall administrators don't create a new service group. Instead, they apply the "no-yoyo" service group to the math department hosts.

Accessing Service Group Configuration

To access service group configuration:

1. From within the Gauntlet Firewall Manager, select Firewall Rules.
2. Select the Service Groups tab.

The Service Groups window displays.

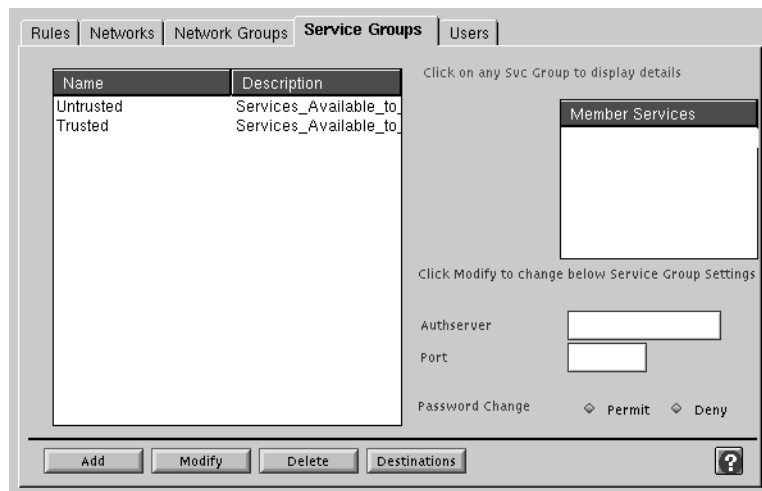


Figure 3-1 Service Groups Window

Configuring Service Groups

This section explains how to configure service groups. It discusses:

- “Planning Service Groups” on page 32
- “Creating Service Groups” on page 33
- “Modifying Service Groups” on page 34
- “Deleting Service Groups” on page 35

Planning Service Groups

When defining and setting up service groups, firewall administrators should plan for a number of groups that span a range from most restrictive to least restrictive services. The exact number of groups required will vary, depending on the size and diversity of the organization and the kinds of communications necessary to accommodate business requirements.

For example, in a small organization that has a limited number of users with uniform trust and levels of responsibilities assumed among all, the minimum trusted and untrusted service groups may be adequate. However, a large institution with multiple departments and diverse levels of trust (for example, Finance versus Engineering) may require a tailored set of service groups that will be applied on a departmental need-to-know basis.

Service groups should be designed so that they provide a required set of member services and associated rules that cover particular business or technical requirements, and no more. In particular, a service group that will be applied to control an untrusted network group or hosts should be limited to the absolute minimum set of services required.

In most cases, firewall administrators find it much easier to initially create a service group with the minimum number of member services required and add new services or liberalize existing service group rules later in response to a particular requirement. It is more difficult to grant a liberal set of services at the beginning and then attempt to remove access to a particular service or destination when users have grown accustomed to having that access.

Creating Service Groups

To create a service group:

1. In the Service Groups window, click Add.
The Add Service Group window displays.

Figure 3-2 Add Service Groups Window

2. Provide information about the new service group.
 - Group Name Name of the new service group.
 - Description Description for the service group.
3. Select the services you want to include in (make members of) your service group:
 - Click the service you would like to add from the list of services not included.
 - Click >> to add the service to the list of included services.

4. Set other options for your service group.

Enforce Authentication	Specifies whether this service group requires authentication for those service members that support authentication.
Authserver	IP address of the host running the authentication server. By default, this is 127.0.0.1, indicating that the authentication server is running on the firewall.
Port	TCP port on which the authentication server is running. By default, this is port 7777.
Allow Password Change	Specifies whether users can change their passwords when connecting from hosts that use this service group.
5. Add destination restrictions for your service group. Refer to Chapter 4, “Destination Access,” on page 39 for more information about destination restrictions.
6. Click OK.
7. Proceed to “Adding Service Group Rules” on page 37 to create rules specifying which networks will use your new service group. Remember to enable the services, as well.

Modifying Service Groups

When one or more changes need to be made to an existing service group, you can modify the settings for that service group.

To modify a service group:

1. Select the service group you wish to modify.
2. Click Modify.

The Modify Service Group window displays.
3. Change the settings for this service group.
4. Click OK.

Deleting Service Groups

When an existing service group is no longer needed, you may delete that group.

To delete a service group:

1. In the Service Groups window, select the service group you wish to delete.
2. Click Delete.

Understanding Service Group Rules

Once a service group has been defined along with its additional parameters and destinations, it may be applied to a network group, network, or specific host. The network or host may be *permitted* to use the service group, which means that it may use the member services, subject to the additional parameter requirements and destinations. Or the network may be *denied* the right to use the service group, which means that it is specifically denied access to the member services in the group.

Order of Precedence

Order of precedence is important when dealing with firewall rules. Applications and services read tables from the top to the bottom. They use the first rule that applies for a particular attribute.

Note: If there are multiple rules in the table that could apply for an attribute, the first one found is the one used. Any subsequent conflicting rules are ignored.

Rules that are higher in the list have a higher order of precedence than rules that are lower in the list. In other words, a higher rule will be interpreted by the firewall before rules further down, and if two rules are encountered that match a given situation but contradict each other, the first one encountered will apply. In general, the more specific rules need to be listed first.

For example, consider the rules menu in Table 3-1:

Table 3-1 Example Rules Menu 1

Rule	Source	Services	Access
1	Research	no-yoyo	Permit
2	Research	*	Deny

In Rule 1, Research is permitted to use the services specified in the no-yoyo service group. Rule 2 denies access to any other services not covered in Rule 1. Assume you reverse the rules, as shown in Table 3-2:

Table 3-2 Example Rules Menu 2

Rule	Source	Services	Access
1	Research	*	Deny
2	Research	no-yoyo	Permit

Research is now denied access to all services in Rule 1. Any services that may have been granted by Rule 2 are ignored.

Accessing Service Group Rules Configuration

To access service group rules configuration:

1. From the Gauntlet Firewall Manager, select Firewall Rules.
2. Click the Rules tab.

The Rules window displays.

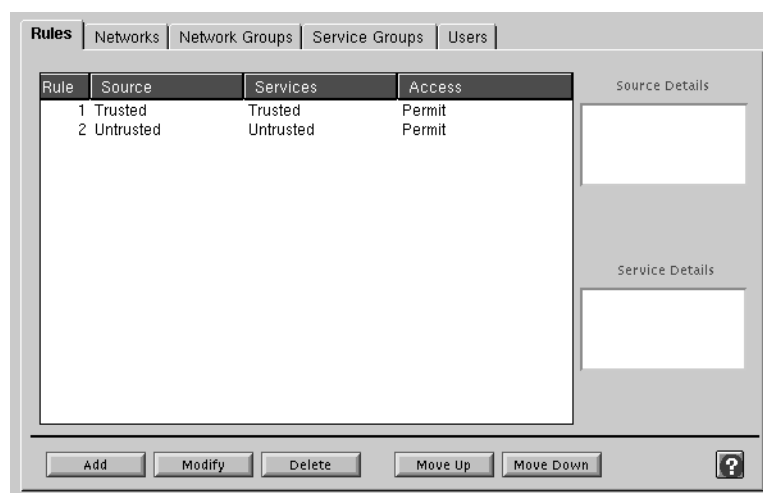


Figure 3-3 Rules Window

Adding Service Group Rules

To add a rule:

1. In the Rules window, click Add.

The Add Security Rule Definition window displays.

2. Provide information about your rule.

Network Source Network group, network, or host for which this rule applies.

Assign Access Specifies whether this rule permits or denies access. If permitted, the services specified in the next section will be available for use by the source system or network. If denied, the services in the next section will be denied to the source.

Service Configurations Services or service groups for which this rule applies.

3. Click OK.
4. Order your new service group rule, as described below, so that your firewall uses your rule in the right order.

Modifying Service Group Rules

To modify an existing rule:

1. In the Rules window, select the rule you wish to modify.
2. Click Modify.

The Modify Security Rule Definition window displays.

3. Change settings as needed.
4. Click OK.

Deleting Service Group Rules

To delete a rule:

1. In the Rules window, select the rule you wish to delete.
2. Click Delete.

Changing Order of Precedence

To move a rule up or down in the list and change its precedence:

1. In the Rules window, select the rule you wish to move.
2. Click Move Up or Move Down as many times as necessary to move the rule to its desired position in the list.

Destination Access

Assume you have configured your firewall to allow users to access the Internet without restriction. Because of a change in corporate security policy, you must now place some restrictions on what sites employees can visit. The Gauntlet Firewall includes the ability to restrict access to certain destinations.

This chapter discusses the concepts of creating destination access rules and explains how to configure the firewall to use the destination access rules. The chapter consists of these sections:

- “Understanding Destination Access” on page 39
- “How Destination Access Works” on page 40
- “Accessing Destination Access Configuration” on page 43
- “Configuring Destination Access Rules” on page 44

Understanding Destination Access

The Gauntlet Firewall allows you to specify the sites you do or do not want users to connect to. You can permit or deny access to a destination by:

- IP address
- hostname

You can also permit or deny access by service or service group. For example, Yoyodyne’s security policy restricts access to several systems on the development network inside their firewall. They are concerned about employees accessing these systems when they come in through the firewall. To implement the security policy, Yoyodyne could create a destination restriction rule that denies access to these systems for the untrusted service group. Or, they could create destination access rules that deny access to these systems for specific services.

Note: The SMTP, X Window System, and SNMP proxy services do not use destination access rules. The SMTP service simply delivers mail between the internal and external mail servers. You cannot permit or deny access to a particular host.

How Destination Access Works

When the firewall receives a request for a particular service, it uses the source address of the request to determine which rule applies. Once the firewall has determined which rule applies (and which service group), it then checks the destination of the request against the destination access rules.

The firewall reads the destination access rules from top to bottom. The firewall uses the first rule that matches. If the destination host in the request matches one of the destination access rules, and that rule indicates that the destination is denied, the firewall denies the request and logs the attempt. If the destination host matches a rule that explicitly permits that destination, the firewall passes the request to the destination host.

Destination access rules follow the rule “That which is not expressly permitted is denied.” This is an important concept to remember, especially when creating access rules that deny access to a particular destination. For example, you want to deny access from your Accounting network to all hosts at Big University. You create a destination access rule that denies access from the accounting service group to *.bigu.edu. This denies access to all hosts at Big University. This rule has the effect of denying access from the Accounting network to all other destinations as well. You have not created a rule that expressly permits access to other destinations, so the firewall denies these requests.

Default Destination Access Rules

The Gauntlet Firewall includes two default destination access rules. These destination access rules permit the trusted network service group and the untrusted network service group to access any destination.

Order of Precedence

Order of precedence is extremely important when dealing with destination rules. Applications and services read tables from top to bottom. They use the first rule that applies for a particular attribute. If there are multiple rules that could apply for an

attribute in the table, the first one found is the one used. Any subsequent conflicting rules are ignored.

Rules that are higher in the list have a higher order of precedence than rules that are lower in the list. In other words, a higher rule will be interpreted by the firewall before rules further down, and if two rules are encountered that match a given situation but contradict each other, the first one encountered will apply. In general, the more specific rules need to be listed first.

For example, consider the list of destination access rules in Table 4-1.

Table 4-1 Destination Access Rules 1

Rule	Name	Destination	Access
1	Trusted (Group)	*.bigu.edu	Deny
2	Trusted(Group)	*	Permit

Rule 1 denies access from the trusted service group to all of the hosts in the bigu.edu domain. Rule 2 permits access for the trusted service group to every other destination.

If Rule 2 were not included, access from the trusted service group would also be denied to every other host. There would be no other rule that explicitly permits this access.

Consider what happens if the rules are reversed (see Table 4-2)

Table 4-2 Destination Access Rules 2

Rule	Name	Destination	Access
1	Trusted(Group)	*	Permit
2	Trusted (Group)	*.bigu.edu	Deny

The hosts in the trusted service group can access any host. Because the firewall reads the rules from top to bottom, and stops when it reaches the first match, it never uses Rule 2.

Considerations for Specifying Destination Addresses

You can specify the destination address by IP address or hostname, or by using the keyword **unknown**. The firewall converts the destination to the same format as the destination access rule and then compares the values. For example, you create a rule that denies access to ftp.bigu.edu. The firewall receives a request with a destination of 192.168.1.33. The firewall uses DNS to convert 192.168.1.33 into a hostname. DNS reports that this IP address maps to ftp.bigu.edu, and the firewall denies the request.

The following table summarizes the behavior of the firewall.

Table 4-3 Firewall Destination Access Behavior

Destination in Packet	Destination in Access Rule	Behavior
IP address	IP address	Compare IP address to IP address.
IP address	hostname	Convert IP address to hostname using DNS reverse lookup. Compare hostname to hostname.
hostname	IP address	Convert hostname to IP address. Compare IP address to IP address.
hostname	hostname	Compare hostname to hostname.

Unknown Keyword

Every system must have an IP address, but it does not necessarily have a hostname. If a system has no hostname, a DNS lookup on its IP address fails. The lack of a registered hostname may be intentional or it may simply be a misconfigured DNS. When there is no hostname for an IP address, the firewall returns unknown as the hostname.

You can create access rules to permit or deny access to these destinations by using the keyword **unknown** as the destination. For example, you create an access rule that denies access to destination **unknown**. The firewall receives a request with a destination of 192.33.112.45. The firewall looks up this IP address, and determines there is no hostname registered in the DNS database for this IP address. The firewall returns **unknown** as the hostname. The firewall then compares the destination in the request (**unknown**) to the destination in the access rule (**unknown**). Because these values match, the firewall denies the request.

Accessing Destination Access Configuration

You can create destination access rules that apply to particular services or to service groups. Similarly, you can configure destination restrictions from these two areas of the Gauntlet Firewall Manager.

To access destination rules configuration:

1. From within the Gauntlet Firewall Manager, select Firewall Rules.
2. Select the Service Groups tab.
The Service Groups window displays.
3. Click Destinations.

or

1. From within the Gauntlet Firewall Manager, select Services.
2. Select a service that supports configuration sets (such as HTTP or TELNET).
3. Click Add to create a new configuration set or Modify to modify an existing configuration set.
The Add or Modify window for the selected service displays.
4. Click Destinations.
The Destination Access window displays.

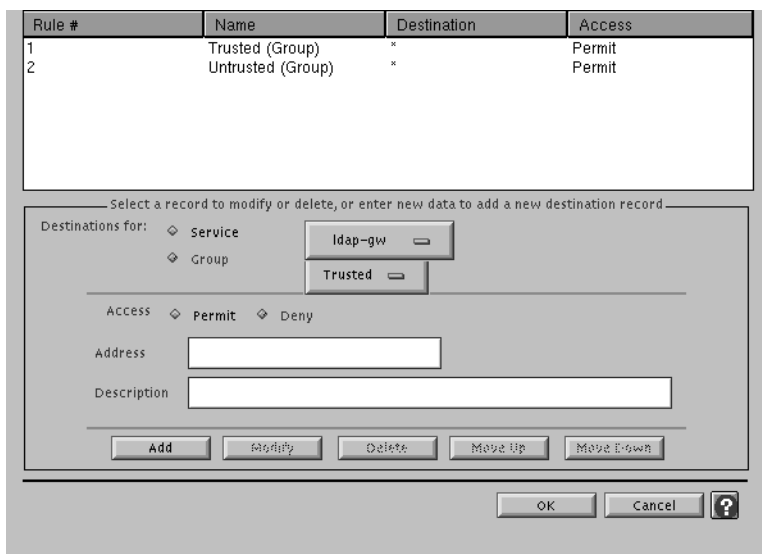


Figure 4-1 Destination Access Window

Configuring Destination Access Rules

Configuring destination access rules are discussed in these sections:

- "Planning Destination Access Rules" on page 45
- "Modifying Destination Access Rules" on page 46
- "Creating Destination Access Rules" on page 45
- "Deleting Destination Access Rules" on page 46
- "Modifying Destination Access Rules" on page 46

Planning Destination Access Rules

When planning destination access rules:

1. Remember that the rules follow the paradigm “That which is not expressly permitted is denied.” Make sure your rules do not deny access unintentionally.
2. Consider whether you wish to create destination restrictions by IP address or by hostname.

Creating Destination Access Rules

To create destination access rules:

1. Indicate whether you are creating this destination restriction for a single service or for a service group.

Service Specifies the service for which this destination access rule applies. Select All to make this rule apply to all services.

Service Group Specifies the service group for which this destination access rule applies.

If you wish to create the same destination access rule for several services or service groups, consider creating one service group to which you can apply the destination access rule. Or, create the same destination access rules for each service or service group.

2. Indicate whether you are creating a permit or deny access rule.

Permit Specifies that you are explicitly permitting access to this site as a destination.

Deny Specifies that you are explicitly denying access to this site as a destination.

3. Enter information about the destination.

Address Enter the IP address or hostname of the system or network to which this rule applies. Specify by IP address, IP address and mask, or hostname. The wildcard * is valid.

Enter **unknown** if you want to permit or deny access to hosts that do not have registered hostnames for their IP addresses.

Description Enter a description for this destination access rule.

4. Click *Add*.
5. Order your new destination access rule, as described in “Changing Order of Precedence” on page 46, so the firewall uses your new rule in the right order.

Modifying Destination Access Rules

To modify destination access rules:

1. Select the destination access rule you want to modify.
2. Modify the information.
3. Click *Modify* to change the rule.
4. Reorder the rule if necessary.

Deleting Destination Access Rules

To delete destination access rules:

1. Select the destination access rule you want to delete.
2. Click *Delete*.
3. Reorder your remaining rules if necessary.

Changing Order of Precedence

Remember that the order in which you place your destination access rules is important. The firewall reads the rules from top to bottom and applies the first one that matches. You generally want to place the most restrictive rules first.

To change the order of precedence:

1. Select the destination access rule you want to move.
2. Click *Move Up* or *Move Down* as many times as necessary to move the rule to the desired position.

Networks and Network Groups

You have many hosts on your internal network. Your security policy may require you to implement different levels of access for groups of these hosts. Rather than creating rules for every single host, it is often easier to think of the hosts as networks. It can also be useful to group those networks into groups of networks. The Gauntlet Firewall allows you to specify networks and network groups.

This chapter discusses the concepts of networks and network groups, and explains how to configure networks and network groups. The chapter consists of the following sections:

- “Understanding Networks and Network Groups” on page 47
- “Accessing Network Configuration” on page 50
- “Configuring Networks” on page 50
- “Accessing Network Group Configuration” on page 53
- “Configuring Network Groups” on page 54

Understanding Networks and Network Groups

The Gauntlet Firewall makes decisions on services and access based on the IP address of the source host. You could try to create rules for every single one of the hosts you manage. In most sites, this would be an extremely tedious process. Or, you can group the systems you manage, and create rules for groups of systems.

A network is simply an object defined by the firewall administrator that represents one or more systems. You denote a network via an IP address or domain name. The following are all valid examples of networks:

- 10.0.1.120
- 10.0.1.*
- 192.5.49.0:255.255.255.0

- dimension.yoyodyne.com
- *.yoyodyne.com

Note: The wildcard * is not valid in the following context: 10.0.8*.

A network group is a collection of one or more networks. Network groups simplify administration. They allow you to create rules for a larger group of systems, instead of having to create individual rules for each system. A network group consists of networks that:

- are logically related because of your security policy (for example, all the networks inside your firewall or all the systems that you use to administer the firewall)
- are physically located in the same place (for example, all the networks located in the Virginia office.)
- have a similar business purpose (for example, all the networks in the Accounting department)
- groups of other networks groups (for example, all network groups in the Maryland office)

Default Network Groups

The initial Gauntlet Firewall configuration defines two network groups: a trusted network group and an untrusted network group.

Trusted Network Group

The trusted network group allows you to group all the hosts that you trust. The trusted network group generally includes:

- the firewall itself (that is, 127.0.0.1)
- hosts and networks inside the firewall (for example, 10.0.1.*, 10.0.2.*, *.yoyodyne.com)
- hosts and networks connected to the firewall via PC Extender

As part of the initial configuration of the firewall, you provide information about your trusted networks, and add them to the trusted networks group. The firewall includes rules that use the trusted service group for the trusted network group.

Untrusted Network Group

The untrusted network group consists of all the hosts that you do not trust. The untrusted network group usually includes every host that you don't explicitly specify as a member of the trusted network group.

The firewall includes rules that use the untrusted service group for the untrusted network group.

When to Add New Network Groups

Define a new network group whenever you need to specify a different set of rules for a different set of networks or hosts. Often, changing business conditions or new requirements require changes to the flow of information to or from an organization.

For example, Yoyodyne is using several of its internal networks for a demonstration for their annual open house. Visitors will have access to the systems on these demonstration networks. The management at Yoyodyne wants to make sure these systems are not used to access inappropriate sites on the Internet. The trusted network group at Yoyodyne allows fairly wide access to the Internet. This is an appropriate time to create a new network group.

To implement the new network group, the firewall administrator:

1. Creates a new network group called Demonstration.
2. Adds the networks that will be used for the demonstration to the Demonstration network group.
3. Creates a Demonstration service group that permits only the services they want visitors to use.
4. Creates a rule that indicates that the Demonstration network group uses the Demonstration service group.
5. Orders the rules so that the firewall reads the Demonstration rule before it reads the trusted rule.

Accessing Network Configuration

To access network configuration:

1. From within the Gauntlet Firewall Manager, select Firewall Rules.
2. Click the Networks tab.

The Networks window displays.

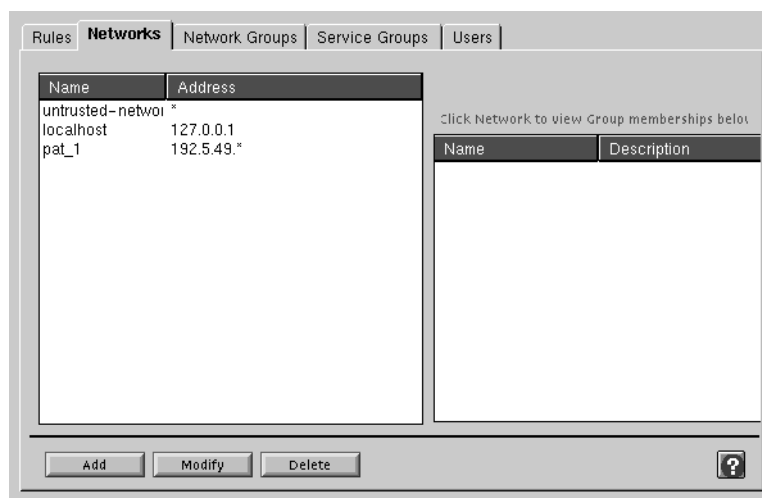


Figure 5-1 Networks Window

Configuring Networks

Configuring a network allows you to refer to the network by name when configuring the Gauntlet Firewall Manager. This section discusses planning, creating, modifying, and deleting networks.

Planning Networks

When planning networks, specify networks by IP address, if at all possible.

Creating Networks

Creating networks identifies various elements in your network.

To create a network:

1. In the Networks window, click *Add*.

The Add Network Definition window displays.

Figure 5-2 Add Network Definition Window

2. Provide the following required information about your network:

Network IP Address Address of the network. Specify individual systems, entire networks, or subnets. Enter by IP address or by IP address and mask. The wildcard * is valid in IP addresses.

- | | |
|-----------|--|
| Interface | Specifies to which interface of the firewall the network is connected.
If all the hosts on this network are inside the firewall, click <i>Inside</i> .
If all the hosts on this network are outside the firewall, click <i>Outside</i> .
If all the hosts on this network are on the service network of the firewall, click <i>ServiceNet</i> .
If the hosts in this network could be on networks inside, outside, or on the service net of the firewall, click <i>Unknown</i> . |
|-----------|--|
3. Provide additional descriptive information about your network.

Group	Click on a network group to make this network a part of a network group.
Description	Description for the network.
Location	Reserved for future use.
MAC Address	MAC address for this host. This setting is valid only if you have specified an individual host.
Reference	Reserved for future use.
 4. Click *OK*.

Modifying Networks

When one or more changes need to be made to an existing network, you can modify the settings for that network.

Note: You cannot modify the IP address of the network. Instead, create a new network and remove the existing network.

To modify a network:

1. In the Networks window, select the network you wish to modify.
2. Click *Modify*.
The Modify Network Definition window displays.
3. Change the settings for the network.
4. Click *OK*.

Deleting Networks

When an existing network is no longer needed, you can delete that network.

To delete a network:

1. In the Networks window, select the network you wish to delete.
2. Click *Delete*.

Accessing Network Group Configuration

To access network group configurations:

1. From within the Gauntlet Firewall Manager, select Firewall Rules.
2. Click the Network Groups tab.

The Network Groups window displays.

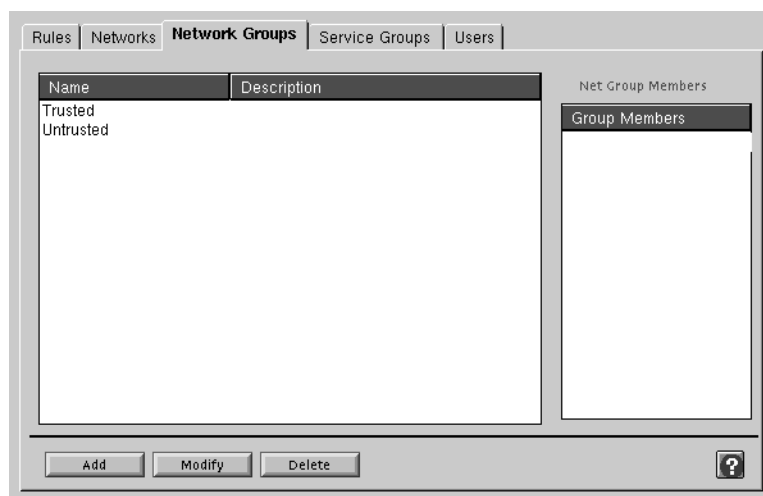


Figure 5-3 Network Groups Window

Configuring Network Groups

This section discusses the following topics:

- “Planning Network Groups” on page 54
- “Creating Network Groups” on page 54
- “Modifying Networks” on page 52
- “Deleting Networks” on page 53

Planning Network Groups

When planning network groups:

- Determine how you wish to group your networks. Consider the types of services you need for various networks.
- Remember not to add the same network to two different network groups.

Creating Network Groups

To create a network group:

1. In the Network Groups window, click *Add*.
The Add Network Group window displays.

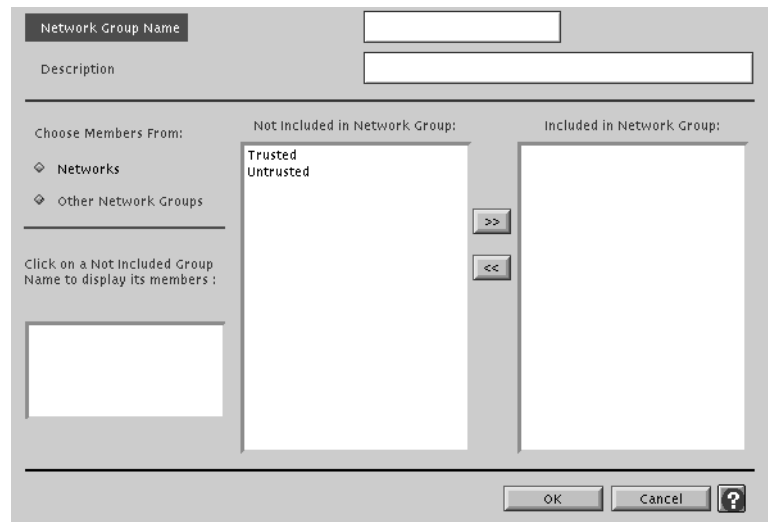


Figure 5-4 Add Network Group Window

2. Provide information about your new network group.
 - Network Group Name Name of the new network group. Names of network groups are case sensitive.
 - Description Description for the network group.
3. Select the networks or network groups you want to include in (make members of) the new network group:
 - Click the network you would like to add from the Not Included list.
 - Click >> to add the network or network group to the Included list.
4. Click OK.

Modifying Network Groups

When one or more changes need to be made to an existing network group, you can modify the settings for that network group.

Note: You cannot modify the name of a network group. Instead, create a new network group and remove the existing network group.

To modify a network group:

1. In the Network Groups window, select the network group you wish to modify.
2. Click *Modify*.
The Modify Network Group window displays.
3. Change the settings for this network group.
4. Click *OK*.

Deleting Network Groups

When an existing network group is no longer needed, you may delete that group.

To delete a network group:

1. In the Network Groups window, select the network group you wish to delete.
2. Click *Delete*.

Users and User Groups

The Gauntlet Firewall can permit or deny access based not just on the host name, but also on the user name. In addition, your security policy may require that users use some form of strong authentication each time they access a particular host or service within their perimeter. To ease the integration of users, strong authentication, and the firewall, the Gauntlet Firewall provides a user authentication management system.

Use of the authentication management system is optional. You must use it any time you have configured the FTP, TELNET, rlogin, POP3, authenticating HTTP, and authenticating circuit proxies to require authentication, which is the default configuration for requests from untrusted networks.

This chapter describes the concepts behind the user authentication management system, and some common administrative tasks in the following sections:

- “Understanding the User Authentication Management System” on page 57
- “Understanding Strong Authentication” on page 60
- “Configuring Users” on page 63
- “Managing Groups” on page 69

Understanding the User Authentication Management System

As part of the security policy, many sites may require some form of strong authentication, which requires users to enter a one-time password or use an authentication token. There are many systems available that can be integrated into a networking environment, each with its own programming and management interface. These are described in more detail in “Understanding Strong Authentication” on page 60 and in Appendix B, “Initializing Strong Authentication Tokens.”

The Gauntlet user authentication management system allows you to easily integrate several different strong authentication systems into your general firewall administration.

You can create, modify, disable, delete, and examine users. The authentication system maintains a database for this information.

How the Firewall Uses Authentication Information

The various proxies use the information in the user authentication management system any time you configure the proxies to require authentication. Using the default Gauntlet service groups and rules, the authentication is required any time a user from an untrusted network tries to access a service inside the perimeter. Recall that untrusted networks are those from which the firewall will accept requests only after authentication by the user.

Remember that using the default service group and rules, the proxies do not authenticate requests from trusted networks. The proxies operate under the assumption that users coming from trusted networks are who they say they are.

Consider the situation of a user, Jane, working at a client site (`blaze.clientsite.com`) who needs information stored on a system at work (`dimension.yoyodyne.com`). When Jane tries to TELNET to `dimension`, which is within the security perimeter, she must first connect to and authenticate at the firewall (`fire-out.yoyodyne.com`).

When `fire-out.yoyodyne.com` receives the information, the TELNET proxy determines that the connection request is from an untrusted network, and that `blaze.clientsite.com` can access inside systems.

The TELNET proxy then prompts Jane for her authentication information (user name and password), which it verifies against the information in the user authentication database. If Jane provided the proper information, and her account is not disabled, the proxy provides a prompt. Jane can then connect to the `dimension` system on the inside network.

How Other Services Use Authentication Information

The `login-sh` program accesses the user authentication server to authenticate users logging into the firewall itself. This login shell authenticates the user before starting the user's normal shell (for example, `csh`, `ksh` or `tcsh`).

The *authmgr* program also uses the user authentication database to authenticate users. This program authenticates the user before allowing the user to do administration from a remote host.

Understanding Users and Groups

This section explains what users and groups mean inside the authentication management system.

Users in the Authentication Management System

User names you create in the user authentication management system are used only for strong authentication.

The user names in the user authentication management system do not need to match user names on the firewall itself because you do not create user accounts on the firewall by default. The exception to this rule is the *login-sh* authentication wrapper program. The *login-sh* program authenticates users before logging them into the firewall. Then, the information in the user authentication management system must match the standard IRIX user information (in */etc/passwd*) for these users.

The user names in the user authentication management system do not need to match any user names on your internal network. For example, John Whorfin might use john as his user name on internal networks. He could use whorfin for strong authentication at the firewall. You may wish to use the same names for the convenience of your users.

Groups in the Authentication Management System

The Gauntlet user authentication management system also makes use of groups. Groups allow you to permit or deny services based on groups of user names, rather than individual user names. For example, you can configure the FTP proxy to permit service to everyone in group sales.

Just as is the case with user names, the groups that you create in the Gauntlet user authentication management system are not the same as the groups you create on the firewall or on the internal network. You can of course use the same names, for easier administration. Each user can be a member of only one group at a time.

Understanding Strong Authentication

The Gauntlet Firewall supports a variety of strong authentication options. The Gauntlet authentication management system understands the types of passwords these systems use, and provides a consistent user interface to these systems.

Note: Your Gauntlet system does not actually include the software or hardware for these products (except as noted below). You must purchase, install, and configure these systems separately.

Currently supported systems are described below.

Access Key II Authentication

This system, from VASCO Data Security, uses a random-challenge password. When the firewall prompts for authentication Access Key II provides a challenge. The user enters the PIN (if one is required) and the challenge into the Access Key II. The Access Key II responds with a password. The user enters this value at the Gauntlet prompt, and the Gauntlet authentication server verifies this value.

You must purchase the Access Key II tokens from VASCO (<http://www.vasco.com>) or their authorized reseller.

See “Access Key II Authentication” on page 331 for more information.

APOP

This system, included with APOP-compliant applications, uses an MD5 secure hash algorithm. The application generates a random challenge and includes it as part of the initial banner.

This option is currently only used by the POP3 proxy.

You must use an APOP-compliant POP3 server and an APOP-compliant POP3 client. Check with the vendors of your POP3 server and client software to determine if their software understands APOP authentication.

CRYPTOCARD RB-1

This system, from CRYPTOCard, uses a random-challenge password. When the firewall prompts for authentication it provides a challenge. The user enters their PIN and the challenge into the CRYPTOCard RB-1. The CRYPTOCard RB-1 encrypts this information to produce a response. The user enters this value at the Gauntlet prompt, and the Gauntlet authentication server verifies this value.

You must purchase the CRYPTOCard RB-1 tokens from CRYPTOCard (<http://www.cryptocard.com>) or their authorized reseller.

Digipass Authentication

This system, from Digipass, uses a time-based password. The Digipass card generates a passcode. When the firewall prompts for authentication, the user selects the appropriate authentication application and enters a personal identification number (PIN) on the card. The card displays a passcode, which the user enters at the Gauntlet prompt. The Gauntlet authentication server verifies this value.

You must purchase the Digipass tokens from Digipass (<http://www.digipass.com>) or their authorized reseller.

See "Digipass Authentication" on page 339 for more information.

SafeWord Authentication Server

This system, from Secure Computing, provides an interface to the SafeWord Authentication Server for Gauntlet authentication. The SafeWord Authentication Server works with a variety of different tokens. The Gauntlet authentication server uses the authentication information registered for a user with the SafeWord Authentication Server.

You must purchase the SafeWord Authentication Server and tokens from Secure Computing (<http://www.securecomputing.com/>) or their authorized reseller.

SecurID Authentication

This system, from Security Dynamics, uses a time-based password. The SecurID card generates a passcode. When the firewall prompts for authentication, the user enters a PIN, if one is required, and the passcode shown on the card. The Gauntlet authentication server verifies this value with the Security Dynamics ACE/Server.

You must purchase the ACE/Server and SecurID tokens from Security Dynamic (<http://www.securitydynamics.com>) or their authorized reseller.

See “SecurID System Authentication” on page 346 for more information.

S/Key Authentication

This system, from Bellcore, uses a one-time password. Users generate a set of passwords based on a “seed” word or phrase. Each time they need to authenticate, they use a different password. When the firewall prompts for authentication, it provides a challenge value. The user enters the appropriate password for that challenge. The Gauntlet authentication server verifies this value.

The Gauntlet Firewall distribution includes a portion of the S/Key package. You can purchase the full S/Key package, which supports both MD4 and MD5 authentication, from Bellcore (<http://www.bellcore.com>). You can also download a freeware version of the full S/Key package that supports only MD4 authentication from Bellcore (<ftp://ftp.bellcore.com/pub/nmh/skey>).

You can also use the Naval Research Lab One-Time Password in Everything (OPIE), which is downward-compatible with Bellcore’s S/Key Version 1 software. The OPIE package is available from the Naval Research Lab (<ftp://ftp.nrl.navy.mil/pub/security/opie>).

See “S/Key System” on page 351 for more information.

RADIUS Authentication

RADIUS stands for Remote Authentication Dial-In User Service, an authentication protocol specified by the IETF. RADIUS authentication can be used with your Gauntlet Firewall in conjunction with a strong authentication method such as Safeword or CRYPTO, or by itself using a plain RADIUS password.

See “RADIUS Authentication” on page 356 for more information.

Reusable Passwords

The Gauntlet Firewall includes a reusable password as part of its authentication system. Reusable passwords are intended for administrator testing only. Every time users need to authenticate; they use the same password.

Caution: Do not use the reusable passwords option for authentication from untrusted networks. Reusable passwords are discouraged. Reusable passwords are vulnerable to password sniffers and are easy to crack. This feature is provided for convenience and audit capability only.

See “Reusable Passwords” on page 360 for more information.

Configuring Users

Configuring users consists of the following potential tasks:

- “Accessing User Configuration” on page 63
- “Creating Users” on page 64
- “Modifying Users” on page 64
- “Enabling Users” on page 68
- “Disabling Users” on page 68
- “Deleting Users” on page 69

Accessing User Configuration

To access user configuration:

1. From within the Gauntlet Firewall Manager, select Firewall Rules.
2. Click the Users tab.

The Users window displays.

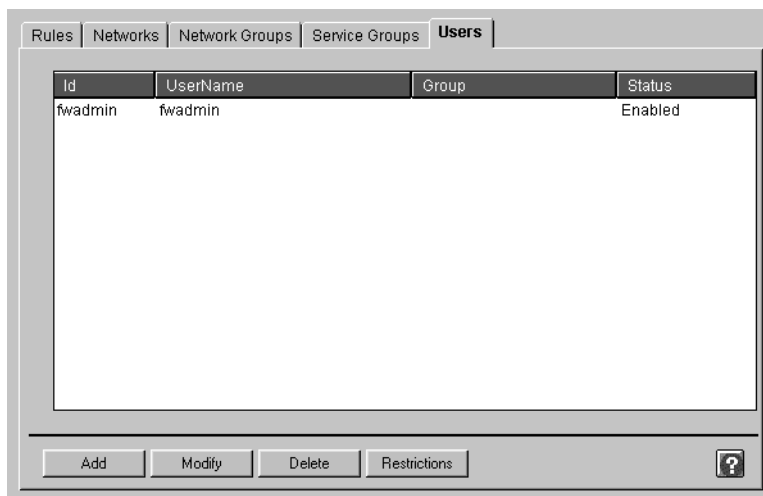


Figure 6-1 Users Window

Creating Users

The process for creating users varies depending on the authentication scheme you are using. Remember that the users that you create in the Gauntlet system are not necessarily the same as the users you create on the firewall or on your internal network. You can of course use the same names, for easier administration. Refer to Appendix B, “Initializing Strong Authentication Tokens,” on page 331 for specific instructions on creating users using your authentication system.

Modifying Users

Modifying users can mean one of the following activities, which are all discussed in this section:

- “Changing User Names” on page 65
- “Changing User IDs” on page 65
- “Changing Authentication Method” on page 66
- “Changing Passwords” on page 66

- “Changing Group Membership” on page 65
- “Allowing Users to Change Their Password” on page 67

Changing User Names

To change a user name:

1. Select the name of the user you wish to modify.
2. Click *Modify*.
The Modify User window displays.
3. Enter a new user name for the user.
4. Click *OK*.

Changing User IDs

You cannot change a user ID. Instead, create a new user ID with the appropriate information and delete the old user ID.

Changing Group Membership

Remember that a user can be a member of only one group.

To change the group of which the user is a member:

1. Select the name of the user you wish to modify.
2. Click *Modify*.
The Modify User window displays.
 - If the group already exists, select the name of the group to which you want to assign the user.
 - If the group does not exist, enter a new group name and click *Add*. Then select that group name.
3. Click *OK*.

Changing Authentication Method

To change authentication methods:

1. Make sure that you have configured your Gauntlet Firewall to work with the authentication system. Refer to Appendix B, "Initializing Strong Authentication Tokens," on page 331 for specific instructions.
2. Select the name of the user you wish to modify.
3. Click *Modify*.
The Modify User window displays.
4. Select the new authentication method.
5. Check the Set Password box to set the user's password.
6. Reenter the same new password in the Verify field.
7. If you would also like to set a POP3 password at this time, check the Set Pop3 Password field and also enter the new POP3 password.
8. Enter the new password in the password field.
9. Click *OK*.

Changing Passwords

Several of the strong authentication systems allow passwords to be set (and reset) by the user:

- CRYPTOCARD
- S/Key
- Reusable password

To change passwords as an administrator:

1. Select the name of the user you wish to modify.
2. Click *Modify*.
The Modify User window displays.
3. Check the Set Password box to set the user's password.
4. Enter the new password in the Password field.

5. Reenter the same new password in the Verify field.
6. Click OK.
7. Provide the user with the new password.

Allowing Users to Change Their Password

Because users are generally not allowed to log directly into the firewall, they need to change their password from another system. The default configuration allows users connecting to the firewall from the inside (trusted) network to change their passwords.

Users can change their passwords through either the TELNET or rlogin proxies. Users should be allowed to change their password only when accessing the firewall from the trusted network.

To change passwords as a user:

1. From a system on the inside network, connect to the firewall using TELNET or rlogin.
2. Use the *password* command.
3. Authenticate to the proxy.
4. Enter the new password.
5. Verify the new password.

The following example shows a sample S/Key password change from the TELNET proxy:

```
dimension-83: telnet fire-in
Trying...
Connected to fire-in.yoyodyne.com
Escape character is '^]'.
tn-gw-> password
Changing passwords
Username: john
Skey Challenge: s/key 644 fi58297 LOAM WOOD BOIL VASE TELL TINY
New Password: #####
Retype New Password: #####
ID john s/key is 664 fi582901
```

Enabling Users

Enabling users allows users who have been disabled to use the system again.

To enable a user:

1. Select the name of the user you wish to modify.
2. Click *Modify*.

The Modify User window displays.

3. To enable the user for multiple logins, select Enabled. To enable the user ID to login once, select Enable One Time Only.

Disabling Users

Disabling users allows you to keep the user information in the system, but does not allow the user to use the system. The user authentication system disables users after a set number (configurable by the administrator) of failed login attempts.

To disable a user:

1. Select the name of the user you wish to modify.
2. Click *Modify*.

The Modify User window displays.

3. Click *Disabled*.

You can also configure the number of consecutive incorrect attempts before the authentication system disables a user's account and you can set the amount of time the account is disabled.

To configure the number of consecutive incorrect attempts, add the `maxbad` attribute to the `netperm` table. Refer to the *Gauntlet Netperm Table Reference Guide* for more information.

To configure the amount of time disabled, add the `badsleep` attribute to the `netperm` table. Refer to the *Gauntlet Netperm Table Reference Guide* for more information.

Deleting Users

Deleting users removes them from the user authentication management system. It does not remove users from your firewall or from your internal network.

To delete a user:

1. Select the name of the user you wish to delete.
2. Click *Delete*.

Managing Groups

As with IRIX systems, the Gauntlet user authentication management system makes use of groups. Groups allow you to permit or deny services based on groups, rather than individual user names. For example, you can configure the TELNET proxy to deny access to a certain destination for everyone in the group `Sales`.

Remember that the groups that you create in the Gauntlet system are not necessarily the same as the groups you create on the firewall or on your internal network. You can, of course, use the same names for easier administration.

Creating Groups

To create a group:

1. Assign a user to a group that did not exist before.

Remember that you may want to make group names the same as existing IRIX groups.

Disabling Groups

You cannot disable entire groups. You must disable usage based on individual users.

Deleting Groups

Currently, to delete a group, you must reassign all users in that group to another group.

User Restrictions

Assume that you have configured your firewall to allow users to access the Internet without restriction. Because of a change in corporate security policy, you must now place some restrictions on the services employees can use and the times during which they can use them. The Gauntlet Firewall includes the ability to restrict certain services by user, time of day, and destination.

This chapter discusses user restrictions concepts and explains how to set up user restrictions in the following sections:

- “Understanding User Restrictions” on page 71
- “How User Restrictions Work” on page 72
- “Accessing User Restriction Configuration” on page 73
- “Configuring User Restriction Rules” on page 75

Understanding User Restrictions

The Gauntlet Firewall allows you to control access by:

- user
- group
- time of day
- destination

The following proxy services support creating user restrictions:

- Circuit
- FTP
- Rlogin
- Rsh
- TELNET

Using the user restrictions, you can control which users can access a particular proxy service. For example, you can deny access to the rlogin or rsh proxy to a particular user. You can also control access by time of day. For example, you can permit access to FTP and TELNET only between 11:00 am and 1:00 pm.

How User Restrictions Work

When the proxy services receive the request from the firewall, they check the user restriction rules. The proxy services read the user restriction rules from top to bottom. The firewall uses the first rule that matches. If the user name in the request matches one of the user restriction rules, and that rule indicates the service is denied, the firewall denies the request and logs the attempt. If the user name matches a rule that explicitly permits a proxy service, the firewall passes the request on to the destination host.

User restriction rules follow the rule “That which is not expressly permitted is denied.” This is an important concept to remember, especially when creating access rules that deny access. For example, suppose you want to deny access to the FTP proxy service for Robert, so you create one user restriction rule that denies this access for him. This rule has the effect of denying access to the FTP proxy to everyone, not just Robert. You have not created a rule that expressly permits access to the FTP proxy for everyone else, so the firewall denies these requests.

To have the firewall use user restriction rules, you must enable authentication. You must also create user IDs in the Gauntlet user authentication system.

Order of Precedence

Order of precedence is extremely important when dealing with user restriction rules. The proxy services read tables from top to bottom. They use the first rule that applies for a particular attribute. If there are multiple rules in the table that could apply for an attribute, the first one found is the one used. Any subsequent conflicting rules are ignored.

Rules that are higher in the list have a higher order of precedence than rules that are lower in the list. In other words, a higher rule will be interpreted by the firewall before rules further down, and if two rules are encountered that match a given situation but contradict each other, the first one encountered will apply. In general, the more specific rules need to be listed first.

For example, consider the following list of user restriction rules in Table 7-1.

Table 7-1 User Restriction Rules Example 1

Rule	Name	Access	Service	Start Time	End Time
1	Robert	Deny	ftp-gw		
2	*	Permit	ftp-gw		

Rule 1 denies access to the *ftp-gw* configuration of the FTP proxy service for Robert. Rule 2 permits access to this service for everyone else.

If rule 2 were not included, access to the *ftp-gw* configuration of the FTP proxy service would also be denied to every other user. There would be no other rule that explicitly permits this access.

Consider what happens if the rules are reversed (see Table 7-2).

Table 7-2 User Restriction Rules Example 2

Rule	Name	Access	Service	Start Time	End Time
1	*	Permit	ftp-gw		
2	Robert	Deny	ftp-gw		

Everyone has access to the *ftp-gw* configuration of the FTP proxy service. Because the firewall reads the rules from top to bottom, and stops when it reaches the first match, it would never use rule 2 in this situation.

Accessing User Restriction Configuration

You can access destination restriction configuration from two areas of the Gauntlet Firewall Manager.

To access user restriction rules configuration:

1. From within the Gauntlet Firewall Manager, select Firewall Rules.
2. Click the Users tab.

The Users window displays.

3. Click *Restrictions*.

The User Restrictions window displays.

Here is an alternative procedure:

1. From within the Gauntlet Firewall Manager, select *Services*.
2. Select the tab for one of the services that supports user restrictions.
3. Click *Add* to create a new configuration set or click *Modify* to modify an existing configuration set.
4. Click *Restrictions*.

The User Restrictions window displays.

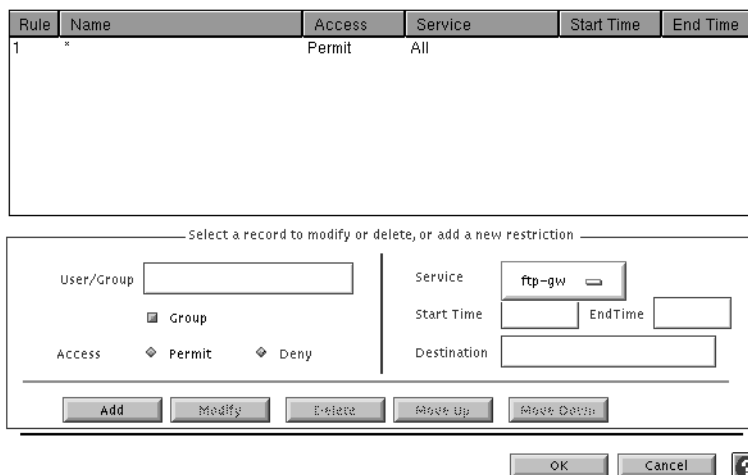


Figure 7-1 User Restrictions Window

Configuring User Restriction Rules

This section discusses configuring user restriction rules in the following sections:

- “Planning User Restriction Rules” on page 75
- “Creating User Restriction Rules” on page 75
- “Modifying User Restriction Rules” on page 76
- “Deleting User Restriction Rules” on page 76
- “Changing Order of Precedence” on page 77

Planning User Restriction Rules

When planning user restriction rules:

- Remember that user restriction rules follow the rule “That which is not expressly permitted is denied.” Once you create a user restriction rule for a service, make sure you are not unintentionally denying that service to other users.
- Be sure to turn on authentication for the service groups or proxies that you want to follow user restriction rules.

Creating User Restriction Rules

To create user restriction rules:

1. In the User Restriction window, enter information about the user for whom this rule applies.

Name	Name of the user or group for whom this rule applies. Use the wildcard * to indicate all users.
------	---
2. Indicate whether you are creating a rule to permit or deny use.

Permit	Specifies that you are explicitly <i>permitting</i> access through this service, during these hours, to this destination.
Deny	Specifies that you are explicitly <i>denying</i> access through this service, during these hours, to this destination.

3. Enter information about the service, time, and destination.

Service	Name of the service for which this rule applies. If you want the rule to affect all services that support user restrictions, select All.
Start Time	Time at which the proxy begins using this rule. Specify in hours and minutes (between 00:00 and 23:59). The wildcard * is valid.
End Time	Time at which the proxy stops using this rule. Specify in hours and minutes (between 00:00 and 23:59). The wildcard * is valid.
Destination	Hosts to which the service can or cannot send requests. Specify individual systems, entire networks, or subnets. Use IP addresses or hostnames. The wildcard * is valid.
4. Click *Add*.
5. Order the new user restriction rule, as described below, so the firewall uses the new rule in the right order.

Modifying User Restriction Rules

To modify user restriction rules:

1. Select the user restriction rule that you want to modify.
2. Modify the information.
3. Click *Modify* to change the rule.
4. Reorder the rules if necessary.

Deleting User Restriction Rules

To delete user restriction rules:

1. Select the user restriction rule you want to delete.
2. Click *Delete*.
3. Reorder the remaining rules if necessary.

Changing Order of Precedence

Remember that the order in which you place user restriction rules is very important. The firewall reads them from top to bottom and applies the first one that matches. You generally want to place the most restrictive rules first.

To change the order of precedence:

1. Select the user restriction rule you want to move.
2. Click *Move Up* or *Move Down* as many times as necessary to move the rule to the desired position.

PART THREE

Configuring and Using Proxy Services

Chapter 8

Managing Proxy Services

Chapter 9

Managing FTP Services

Chapter 10

Managing LDAP Services

Chapter 11

Managing Microsoft SQL Services

Chapter 12

Managing Network Management Services

Chapter 13

Managing Network Management Services

Chapter 14

Managing News Services

Chapter 15

Managing Print Services

Chapter 16

Managing rsh Services

Chapter 17

Managing Sybase Services

Chapter 18

Managing Terminal Services

Chapter 19

Managing WWW and Gopher Services

Chapter 20

Managing X Window Services

Chapter 21

Managing Custom Services

Chapter 22

Managing Custom Services With Authentication

Chapter 23

Managing MediaBase Services

Managing Proxy Services

Although the Gauntlet Firewall provides a wide variety of proxy services, there are a number of common elements to configuring these services.

This chapter explains the concepts of proxy services and provides general information on how to configure proxy services. The chapter consists of the following sections:

- “Understanding Proxy Services” on page 81
- “Configuring Proxy Services” on page 82

Understanding Proxy Services

The proxy services on the Gauntlet Firewall are application-level proxies. The services are called proxies because they relay (or proxy) information from one side of the firewall to the other. This prevents a program on one side of the firewall from talking directly to a system on the other side of the firewall.

They are considered application-level proxies because each proxy handles a different protocol, and thus, a different type of application. For example, the FTP proxy service understands the FTP protocol and handles requests for passing FTP traffic through the firewall. Several of the proxies understand more than one protocol, and can communicate with several different applications. For example, the multimedia proxy understands the protocols used by NetShow, RealAudio, RealVideo, and VDOLive.

The firewall also includes a generic proxy service, called the plug proxy. This proxy service understands the TCP protocol and can work with a variety of different applications.

The descriptions of each proxy service, found in the remaining chapters of this book, provide more information on how the proxy services work.

Configuring Proxy Services

You can configure a variety of parameters for each proxy service. For each proxy, you can enable and disable the proxy service. This allows you to indicate whether or not the proxy should offer a particular service.

Other configurable parameters vary for each proxy service. For some proxies you can set the number of processes the proxy can start (the child limit) while for others you cannot set that number. Refer to the online help for explanations of the parameters. If you are in doubt about changing a parameter, use the default value shown in the Gauntlet Firewall Manager.

The main configuration window for each proxy service in the Gauntlet Firewall Manager shows the parameters that apply to all instances of that proxy.

Creating Multiple Configurations for a Proxy

You may need to have multiple configurations for a proxy service. For example, assume that company security policy requires you to limit the destinations that the hosts on the sales networks can visit on the World Wide Web. You could modify the default configuration in the HTTP proxy to restrict access to these destinations. This, however, affects every host on your trusted network, not just the hosts on the sales network. Instead, you want to have two different configurations; the default configuration and a configuration that is more restrictive than the default.

The firewall allows you to create multiple configuration sets for some proxies. Configuration sets allow you to have multiple configurations for the same proxy. You can then add the configuration sets to different service groups, and create different service group rules to use these service groups.

In our example above, you would leave the default configuration (http-gw) for the HTTP proxy in the trusted service group. You would then create a separate configuration set (http-gw-restrictive) for the HTTP proxy that restricts the destinations that the HTTP proxy can visit. You would create another service group that includes this configuration set. You would then create a rule that forces the hosts on the sales network to use the service group that contains the restrictive configuration set.

In fact, it is configuration sets that allow you to create custom proxy services (plugs) and custom proxy services with authentication (circuit). For the plug proxy, you can create one configuration set that runs on port 17 and listens for Quote-of-the-Day requests and another that listens on port 3572 for requests from internal accounting software.

The following proxies support creating multiple configuration sets:

- circuit
- Gopher
- FTP
- HTTP
- Lp
- plug
- rlogin
- SQL Server
- StreamWorks
- Sybase
- TELNET

Working With Configuration Sets

Adding configuration sets involves planning, then creating the configuration sets. Both activities are discussed in this section.

Planning Configuration Sets

When planning configuration sets:

- Determine whether you need to create a new configuration set or whether you can simply modify the default configuration set.
- Determine which service groups will use this configuration set.

Creating Configuration Sets

When you create a configuration set, the firewall uses the information specified in the default configuration set as a template.

To create a configuration set:

1. From within the Gauntlet Firewall Manager, select Services.
2. Click the tab for one of the services that supports configuration sets.
3. In the main configuration window for the proxy service, provide information that applies to all configuration sets for that proxy.
4. Click *Add*.

The Add Services window for the selected proxy displays.

5. Provide information about the configuration set.

You must provide the following information.

Name Name of the configuration set. The firewall uses this value in a variety of places, including the list of available services for a service group. Use a descriptive name, preferably one that includes the name of the service. This name must be unique among all configuration sets on this firewall.

Description Description for the configuration set. This description helps you track various configuration sets.

6. Configure other settings for the configuration set, such as messages and restrictions. Refer to the online help for descriptions of the options available for each proxy service.
7. Click OK.

Be sure to add your configuration set to one of the service groups. If you do not add the configuration set to a service group, the firewall does not use that set.

Modifying Configuration Sets

Note: You cannot change the name of a configuration set. Instead, create a new configuration set that has the name and properties you want, and delete the old configuration set.

To modify a configuration set:

1. Select the configuration set you wish to modify.
2. Click *Modify*.

The Modify Services window for the selected proxy displays.

3. Change the settings for this configuration set.
4. Click *OK*.

Deleting Configuration Sets

To delete a configuration set:

1. Select the configuration set you wish to delete.
2. Click *Delete*.

Managing FTP Services

Sometimes the easiest way to transfer information from one system to another is to actually transfer the relevant files. The File Transfer Protocol (FTP) is one of several protocols that make this possible. The Gauntlet Firewall includes a proxy that allows secure file transfer between the outside network and the inside network.

This chapter explains the concepts behind the FTP proxy and how it works, how to configure it, and how to use FTP services. One section discusses considerations for running anonymous FTP servers. The chapter consists of the following sections:

- “Understanding the FTP Proxy” on page 87
- “How the FTP Proxy Works” on page 88
- “Accessing FTP Proxy Configuration” on page 89
- “Configuring the Firewall for FTP Services” on page 89
- “Using FTP Services” on page 91
- “Running an Anonymous FTP Server” on page 94

Understanding the FTP Proxy

The FTP proxy is an application-level proxy that provides configurable access control, authentication, and logging mechanisms.

The FTP proxy, which runs on the firewall, passes FTP requests through the firewall, using rules you supply. You can configure the FTP proxy to allow file transfer activity based on:

- source IP address
- source hostname
- destination IP address
- destination hostname
- FTP commands (for example, STOR and RETR)

Using these options, you can configure your firewall to allow specific hosts on outside networks to transfer files to and from inside hosts. Employees working at specific customer sites can access files on their workstations. Similarly, you can configure your firewall to permit users on the inside network to copy files (using the FTP daemon RETR command) from hosts on the outside network, but not place files (using the FTP daemon STOR command) on these outside hosts.

The FTP proxy allows administrators to require users to authenticate before transferring files. The FTP proxy logs all successful and unsuccessful file transfer attempts, and the number of bytes transferred.

The FTP proxy's access controls allow you to have more control over the files entering and leaving your system than you would by using the standard IRIX FTP daemon. The logging capabilities are also more extensive.

How the FTP Proxy Works

The firewall runs the network access control daemon (*netacl*) as a daemon listening for requests on the standard FTP port (TCP port 21). Whenever the daemon receives an FTP request on this port, the *netacl* daemon checks its configuration information and determines whether the initiating host has permission to use FTP. If the host has permission, the *netacl* daemon starts the standard FTP server (*ftpd*) or the FTP proxy (*ftp-gw*). If the host does not have permission, the daemon displays an error message.

The default trusted service group and rules allow all inside hosts to initiate FTP sessions and transfer files without authenticating. The inside host passes FTP requests to the firewall, which starts the *netacl* daemon. The *netacl* daemon checks its permissions, and determines that the inside host can use FTP. The *netacl* daemon starts *ftp-gw*. The proxy logs the transaction and passes the request to the outside host. *ftp-gw* remains active until either side terminates the connection. The default untrusted service group and rules also allow outside hosts to initiate FTP sessions. They must, however, authenticate before accessing inside hosts.

The default configuration does not allow either inside or outside hosts to FTP directly to the firewall itself. If you configure your Gauntlet Firewall to allow FTP to the firewall, hosts connect to the firewall with an FTP request. The firewall starts the *netacl* daemon. The *netacl* daemon checks its permissions, and determines that outside hosts can use FTP to the firewall itself. The *netacl* daemon starts the standard FTP daemon (in a chrooted environment).

This configuration using *netacl* allows a fair amount of flexibility in configuring FTP services. Users inside the perimeter can continue to interact with outside hosts, generally without authentication. Users outside the perimeter can interact with inside hosts, generally with authentication.

Accessing FTP Proxy Configuration

To access the FTP proxy configuration:

1. From within the Gauntlet Firewall Manager, select Services.
2. Click the FTP tab.

The FTP window displays.

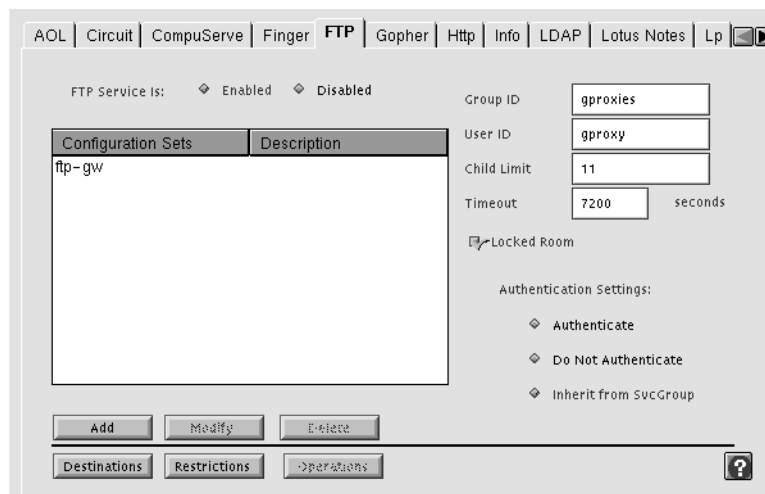


Figure 9-1 FTP Window

Configuring the Firewall for FTP Services

Configuring the Gauntlet Firewall for FTP services involves planning, configuring the FTP proxy to enforce company security policy, enabling the proxy, and creating user accounts for users who will need to authenticate.

Planning FTP Proxy Settings

When planning FTP proxy settings, determine policies for:

- Requiring authentication.
- Allowing specific FTP commands (for example, RETR and STOR).
- Permitting or denying specific sources and destinations.

Configuring FTP Proxy Settings

Configure the FTP proxy to enforce company security policies.

To configure FTP proxy settings, you can provide optional information about time-out values and other configuration settings for the FTP proxy. Refer to the online help for specific information about the available settings. Refer to Chapter 30, “Managing Content Scanning,” on page 285 for information on configuring the content scanning features for the FTP proxy.

Enabling FTP Proxy Services

To enable the FTP proxy service:

1. In the FTP configuration tab, click Enabled.
2. Add the FTP configuration to the service groups you want to use the FTP proxy.
3. Before exiting the Gauntlet Firewall Manager, Save and Apply your changes.

The firewall enables the FTP proxy.

Creating Authentication User Entries

Use the authentication management system to create authentication user entries for any users who must authenticate when using FTP services. See Chapter 6, “Users and User Groups,” on page 57 for more information.

Verifying Your Setup

Verify your configuration by transferring files to an inside host from an outside host. For example, connect to your favorite FTP site and download their *README* file. See the section below for instructions.

Using FTP Services

The idea behind the FTP proxy is that most users working on the trusted networks behind the firewall will not see a change in their daily FTP activities. The default configuration allows users on trusted networks to FTP to untrusted networks without authenticating. Users on the trusted networks do not need to change their FTP procedures.

Using Authentication

If you have configured any FTP activities to require authentication, users will need to follow different procedures to use FTP.

To FTP using authentication:

1. FTP to the firewall itself.
2. Authenticate to the proxy.
3. Connect to the desired FTP server.
4. Continue as before.

A common security policy for the FTP proxy is to authenticate all requests from untrusted networks to or through the firewall. The example below shows a sample FTP session from an untrusted network to a trusted network, using S/Key authentication at the firewall:

```
blaze.clientsite.com-27: ftp fire-out.yoyodyne.com  
Connected to fire-out.yoyodyne.com  
220-Proxy first requires authentication  
220 fire-out.yoyodyne.com FTP proxy (Version 4.0a) ready.  
Name (fire-out.yoyodyne.com:clancy): clancy  
331 Skey Challenge: s/key 653 fi19289  
Password:password does not display
```

```
230 User authenticated to proxy
ftp> user clancy@dimension
331- (-----GATEWAY CONNECTED TO dimension-----)
331- (220 dimension FTP server ready.)
331 Password required for clancy.
Password: #####
230 User clancy logged in.
ftp>
```

In this example, Clancy, working at a client site (blaze.clientsite.com), needs FTP access to a system behind the firewall (dimension.yoyodyne.com). He first FTPs to the outside address of the firewall for Yoyodyne (fire-out.yoyodyne.com). The FTP proxy on fire-out prompts him to authenticate. Clancy provides his authentication user ID (clancy). When the proxy prompts, he enters the response to the authentication challenge, which does not display. The proxy authenticates clancy.

Clancy indicates the host he needs to access and his user name for that host (clancy@dimension). The FTP proxy connects Clancy to dimension and prompts him for his password on dimension. Clancy enters his password for dimension. The FTP server on dimension verifies Clancy's user name and password, and logs him in. Clancy can now transfer files.

Using Authentication With Some GUI FTP Tools

The FTP proxy can require you to authenticate twice. Some GUI FTP tools for Microsoft Windows and the Macintosh require you to specify the user name and password in a dialog box. These tools assume that once you supply this information, you are connected. The FTP proxy displays the challenge and response information for authentication in FTP comments.

Some Microsoft Windows and Macintosh FTP tools do not display FTP comments. Unless the user sees the comment, they will have a really difficult time trying to guess the current challenge. You can still use these FTP tools with S/Key authentication, by combining the authentication and FTP host information.

To authenticate using some GUI tools:

1. For the hostname, supply the name of the firewall.
2. For the user name, supply the firewall authentication user name, the FTP host user name, and the name of the FTP host in this form:

authentication-username@ftp-host-username@ftp-host

3. For the password, supply the authentication response and FTP host password:

authentication-response@ftp-host-password

You may need to TELNET to the firewall to see what the next challenge is.

The example below shows the information a user enters in their FTP tool when going from an untrusted network to a trusted network, using S/Key authentication for the firewall:

```
host:      fire-out.yoyodyne.com
username:  clancy@clancy@dimension
password:  elk elba iris odd skim lee@#####
```

Because you cannot tell what the next challenge will be when using an authentication system that uses random challenges (such as the AssureNet Pathways SecureNet Key), you may not be able to use these instructions with some GUI FTP tools.

Running an Anonymous FTP Server

By its very nature, an anonymous FTP server requires easy access by the public. If you place the anonymous FTP server behind the firewall, you are allowing an additional type of access within your security perimeter. If you place the FTP server on the firewall itself, you are allowing additional access to your firewall.

Gauntlet for IRIX allows you to run the standard IRIX FTP server (*ftpd*) in an isolated (chrooted) environment as an anonymous FTP server (but you give up the ability to allow authenticated users from untrusted networks to use *ftp-gw* to access trusted networks).

The best solution is generally to place your anonymous FTP server on a system outside the perimeter. Follow good security practices for this system:

- Turn off all other services.
- Create the minimum number of user accounts.
- Use strong authentication.
- Patch the operating system and applications.
- Use checksums to watch for file changes.
- Back up frequently.

Managing LDAP Services

The employees of your company want to engage in secure communications with other people. The mail, file encryption, and other types of applications they are using use certificates as part of the authentication and encryption process. Before communicating with someone, they need the appropriate certificate. The Lightweight Directory Access Protocol (LDAP) is a commonly used protocol for providing this sort of information. The Gauntlet Firewall includes a proxy that allows connections between LDAP clients and servers.

This chapter explores the concepts behind the LDAP proxy and explains how it works, how to configure the proxy, and how to use LDAP services. The chapter consists of these sections:

- “Understanding the LDAP Proxy” on page 95
- “How the LDAP Proxy Works” on page 96
- “Configuring LDAP Clients” on page 97
- “Configuring LDAP Proxy Settings” on page 97
- “Planning the LDAP Proxy” on page 97
- “Enabling LDAP Services” on page 99

Understanding the LDAP Proxy

The LDAP proxy is an application-level proxy that provides configurable access control and logging mechanisms. The LDAP proxy, which runs on the firewall, passes LDAP requests through the firewall (at the application level), using rules you supply.

You can configure the proxy to allow connections based on:

- source IP address
- source hostname

- source port
- destination IP address
- destination hostname
- destination port

Using these options, you can configure the firewall to allow LDAP clients on certain trusted hosts to access an LDAP server on an untrusted host. Employees working behind the firewall can access LDAP servers at customer sites.

You can configure the LDAP proxy to allow LDAP clients on untrusted hosts to access LDAP servers on the trusted networks. According to most security policies, it is not a good idea to allow unlimited access to the inside network. If you have to offer LDAP services to other hosts on the Internet, consider running the LDAP server on the outside network.

The proxies log all successful and unsuccessful connection attempts, and the amount of data transferred. These access controls allow you to have much more control over the connections to and from your system than you would without a firewall. The logging capabilities are also much more extensive.

How the LDAP Proxy Works

The firewall runs the LDAP proxy as a service listening for requests on the standard LDAP port. Whenever the firewall receives an LDAP request on this port, the LDAP proxy checks its configuration information and determines whether the initiating host has permission to initiate this type of request. If the host does not have permission, the LDAP proxy logs the unsuccessful connection.

If the host has permission, the proxy logs the transaction and passes the request to the destination host. The LDAP proxy remains active until either side closes the connection or the connection times out.

The recommended configuration allows trusted hosts to access LDAP servers on untrusted networks. The recommended configuration does not allow untrusted hosts to access LDAP servers on trusted networks.

Planning the LDAP Proxy

When planning the LDAP proxy:

- Determine which LDAP servers users need to access. Determine whether you want to limit access to a particular server or not. Obtain host name or IP address information for each server.
- For each server, determine the port(s) on which the server accepts connections.
- Determine which internal hosts can use these services.

Configuring LDAP Clients

You can configure LDAP clients with or without transparency.

Configuring LDAP Clients With Transparency

If you are using transparency on the firewall (the default configuration) and you have installed the LDAP proxy on the default port (TCP port 389), you do not have to change the way that inside hosts access LDAP servers on the outside network.

Configuring LDAP Clients Without Transparency

If you are not using transparency, configure the LDAP client to know about the LDAP proxy. Consult the documentation included with your mail reader or other LDAP client for instructions.

Configuring LDAP Proxy Settings

To configure LDAP settings:

1. From within the Gauntlet Firewall Manager, select Services.
2. Click the LDAP tab.

The LDAP window displays.

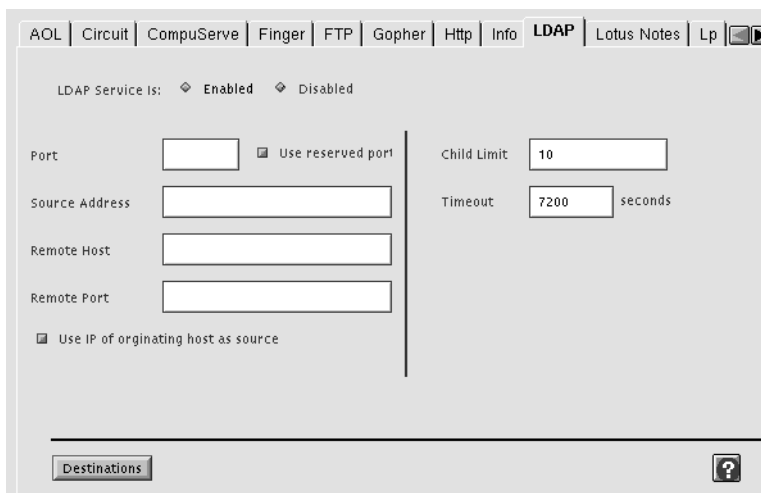


Figure 10-1 LDAP Window

3. Configure the LDAP proxy settings.

LDAP Service Is	Select Enabled or Disabled.
Port	The default port number will work if you are using a standard LDAP port. If your LDAP services are using a non-standard port, enter the port number.
Use a reserved Port	Check if you are using a reserved port, one with a port number less than 1023.
Source Address	Enter the IP addresses of hosts from which connections can originate. Specify single hosts, entire networks, or subnets. Specify by IP address or host name. The wildcard * is valid in hostnames.
Remote Host	Enter the IP addresses of the host to which the LDAP proxy connects. Specify single hosts, entire networks, or subnets. Specify by IP address or host name. The wildcard * is valid.
Remote Port	Enter the port on the remote host to which the LDAP proxy connects.
Use IP of originating host source	Check to use the IP address of the originating host as the source address. Leave blank otherwise.

Child Limit	Enter the maximum number of child processes the LDAP proxy allows to run at a given time.
Timeout	Enter the number of minutes the connection can be idle before it is disconnected.

Enabling LDAP Services

To enable the LDAP proxy:

1. In the LDAP window make sure LDAP service is enabled.
2. Add the LDAP configuration to the service groups you want to use the LDAP proxy.
3. Before exiting the Gauntlet Firewall Manager, save and apply your changes.

The firewall enables the LDAP proxy.

Managing Microsoft SQL Services

Database services are essential in most organizations. As with other services you offer, you want to securely configure database access. Microsoft SQL is a relational database management system in use in many organizations. The Gauntlet Firewall includes a proxy that securely allows connections between Microsoft SQL clients on the inside network and servers on the outside network.

This chapter discusses the concepts behind the SQL Server proxy and explains how it works, how to configure it, and how to use Microsoft SQL services. The chapter consists of these sections:

- “Understanding the SQL Server Proxy” on page 101
- “How the SQL Server Proxy Works” on page 102
- “Accessing SQL Server Proxy Configuration” on page 103
- “Configuring the Firewall for Microsoft SQL Services” on page 104
- “Enabling SQL Server Proxy Services” on page 107
- “Configuring Microsoft SQL Clients” on page 108
- “Verifying Your Setup” on page 108

Understanding the SQL Server Proxy

The SQL Server proxy is an application-level proxy that provides configurable access control, authentication, and logging mechanisms. The SQL Server proxy, which runs on the firewall, passes Microsoft SQL requests through the firewall (at the application level), using rules you supply. You can configure instances of the SQL Server proxy to service:

- Microsoft SQL client-to-server communications
- Microsoft SQL server-to-server communications

For each instance of the SQL Server proxy, you can configure the proxy to allow connections based on:

- source IP address
- source hostname
- source port
- destination IP address
- destination hostname
- destination port

Using these options, you can configure the firewall to allow Microsoft SQL clients on certain trusted hosts to access a Microsoft SQL server on an untrusted host. Employees working behind the firewall can access Microsoft SQL databases at customer sites. You can also configure the firewall to allow Microsoft SQL servers on opposite sides of the firewall to communicate. A Microsoft SQL replication server can communicate with another Microsoft SQL replication server on the other side of an intranet firewall.

You can configure the SQL Server proxy to allow Microsoft SQL clients on untrusted hosts to access Microsoft SQL servers on your trusted networks. According to most security policies, including the Gauntlet Firewall default, allowing untrusted hosts such access is not a good idea. If you must allow this sort of service, consider using client-side password encryption. Consider limiting the databases and data to which users have access, because all of the data is transferred unencrypted.

The proxies log all successful and unsuccessful connection attempts and the amount of data transferred.

These access controls allow you to have much more control over the connections to and from your system than without a firewall. The logging capabilities are also much more extensive.

How the SQL Server Proxy Works

The firewall runs different instances of the SQL Server proxy (*mssql-gw*) as daemons on different ports for different Microsoft SQL applications. Whenever the firewall receives a Microsoft SQL request on one of these ports, the SQL Server proxy checks its configuration information and determines whether the initiating host has permission to

initiate this type of request. If the host does not have permission, the Microsoft SQL daemon logs the connection attempt and displays an error message.

If the host has permission, the proxy logs the transaction and passes the request to the destination host. The SQL Server proxy remains active until either side closes the connection.

The default service groups do not allow either inside or outside hosts to use the SQL Server proxy. The recommended configuration allows trusted hosts to access Microsoft SQL servers on untrusted networks. The recommended configuration does not allow untrusted hosts to access Microsoft SQL servers on trusted networks.

While the SQL Server proxy does perform checks to make sure the packets appear to be Microsoft SQL packets, someone could spoof this protocol. The SQL Server proxy does not perform any user authentication. You are relying on the authentication mechanisms of the Microsoft SQL server to control access to your Microsoft SQL server and its data.

Accessing SQL Server Proxy Configuration

To access the SQL Server proxy configuration:

1. From within the Gauntlet Firewall Manager, click *Services*.
2. Click the SQL Server tab.

The SQL Server window displays.

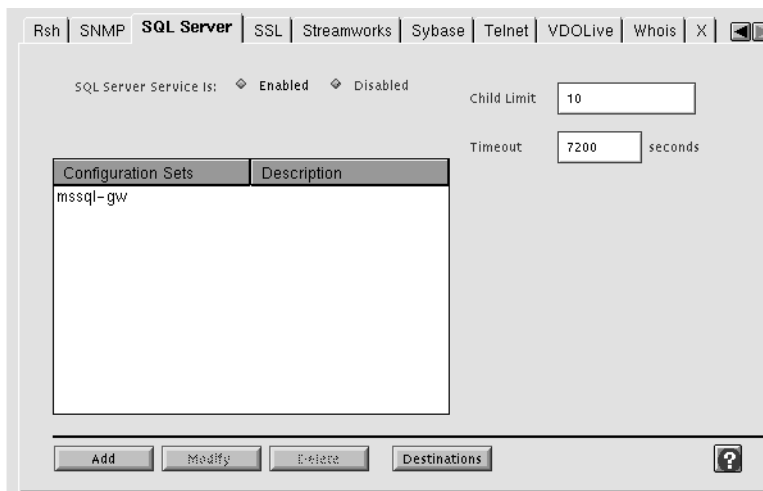


Figure 11-1 SQL Server Window

Configuring the Firewall for Microsoft SQL Services

Configuring the Gauntlet Firewall involves planning, configuring the proxies to enforce company security policy, and enabling the proxy.

Planning SQL Server Proxy Settings

When planning the Microsoft SQL service proxy:

- Determine which Microsoft SQL servers users need to access. Determine whether you want to limit access to a particular server or not. Obtain host name or IP address information for each server.
- For each server, determine the port(s) on which the server accepts connections.
- Determine which external hosts can use these services.
- Determine which internal hosts can use these services.

Configuring SQL Server Proxy Settings

Configure the SQL Server proxy to enforce company security policies.

To configure SQL Server proxy settings:

1. In the SQL Server window, configure the SQL Server proxy settings.

SQL Server Service is Select Enabled or Disabled.

Child Limit Maximum number of child processes the SQL Server proxy allows to run at a given time.

Timeout Number of minutes the connection can be idle before it is disconnected.

2. Add SQL Server configuration sets as appropriate.

To add an SQL Server proxy configuration set:

1. In the SQL Server window, click *Add*.

The Add SQL Server Services window displays.

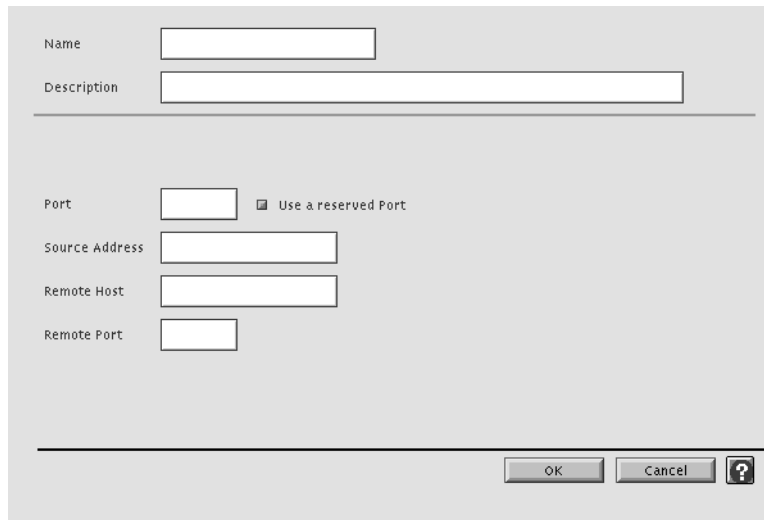


Figure 11-2 Add SQL Server Services Window

2. Provide information about the hosts on which you are running the Microsoft SQL service.

Name	Name for this SQL Server configuration set.
Description	Description for this SQL Server configuration set.
Port	Port number on which the proxy runs. The default is 1612.
Use a reserved Port	Check if you are using a reserved port, one with a port number less than 1023.
Source Address	IP addresses of hosts from which connections can originate. Specify single hosts, entire networks, or subnets. Specify by IP address or host name. The wildcard * is valid in hostnames.
Remote Host	IP addresses of the host to which the SQL Server proxy connects. Specify single hosts, entire networks, or subnets. Specify by IP address or host name. The wildcard * is valid.
Remote Port	Port on the remote host to which the SQL Server proxy connects.

Use IP of originating host source Check to use the IP address of the originating host as the source address. Leave blank otherwise.

3. Click *OK*.

The SQL Server window re-displays.

To modify an existing configuration set:

1. In the SQL Server window, select the configuration set you wish to modify.
2. Click *Modify*.

The Modify SQL Server Services window displays.

3. Make the desired modifications to the configuration set.
4. Click *OK*.

The SQL Server window re-displays.

To delete a configuration set:

1. In the SQL Server window, select the configuration set you wish to delete.
2. Click *Delete*.

The configuration set disappears from the list of configuration sets.

Enabling SQL Server Proxy Services

To enable the SQL Server proxy service:

1. In the SQL Server window, make sure Enabled is selected.
2. Add the SQL Server proxy configuration to the service groups you want to use the Microsoft SQL proxy.
3. Before exiting the Gauntlet Firewall Manager, save and apply your changes.

The firewall enables the SQL Server proxy.

Configuring Microsoft SQL Clients

Add or modify the interfaces file on the client to provide information about the Microsoft SQL server:

To configure Microsoft SQL clients:

1. Specify the port number you selected for the SQL Server proxy.
2. If you are using transparency (the default configuration), specify the hostname as the hostname of the actual system running the Microsoft SQL server. If you are not using transparency, specify the hostname as the IP address of the firewall.

If you are using server-to-server communications, configure all servers as clients. Consult the Microsoft SQL administration documentation for further information on configuring clients for accessing servers.

Verifying Your Setup

To verify your setup, use the Microsoft SQL client on a trusted host to run a simple query against the Microsoft SQL server on the untrusted host. Watch the logs on the firewall for error messages.

Managing Multimedia Services

More and more people are listening to audio and watching video files found at sites on the Internet. As with other protocols, access to these files is not without risk. They require logging and access control, as with other services.

There are a number of protocols that allow people to listen to audio material and view video material. The Gauntlet Firewall includes multimedia proxies that securely handle requests for the following formats:

- NetShow
- RealPlayer (RealAudio and RealVideo)
- StreamWorks
- VDOLive

The Gauntlet Firewall also includes a MediaBase proxy, which is described in chapter 23 of this manual. This chapter discusses the concepts behind the multimedia proxies and explains how they work and how to configure and use them. The chapter contains the following sections:

- “Understanding the Multimedia Proxy” on page 110
- “Accessing Multimedia Proxy Configuration” on page 111
- “Configuring the Firewall for Multimedia Services” on page 112
- “Using the NetShow Proxy” on page 114
- “Using the RealPlayer Proxy” on page 115
- “Using the StreamWorks Proxy” on page 117
- “Using the VDOLive Proxy” on page 118

Understanding the Multimedia Proxy

The Gauntlet multimedia proxies are application level proxies that provide configurable access control. The proxy, which runs on the firewall, passes client requests through the firewall, using rules you supply. You can configure the multimedia proxies to allow connections based on:

- source hostname
- source IP address
- destination hostname
- destination IP address
- service port number

Using these options, you can configure the firewall to allow clients on the inside network to access audio or video servers on the outside network. You can also limit the sites users can access from systems on the inside network. The multimedia proxies also log all successful and unsuccessful connection attempts and the amount of data transferred.

You cannot configure the multimedia proxies to allow access to servers on the inside network.

These access controls allow you to have more control over the connections to and from your system than without a firewall. The logging capabilities are also more extensive.

How the Multimedia Proxy Works

The firewall runs different instances of the multimedia proxy on the appropriate ports for NetShow (TCP port 1755), RealAudio and RealVideo (TCP port 7070), and VDOLive (TCP port 7000). The firewall also runs the StreamWorks proxy as a daemon listening for requests on the StreamWorks XDMA port (UDP port 1558).

When the firewall receives requests for services on this port, the proxy checks its configuration information and determines whether the initiating host has permission to use the requested service. If the host has permission, the proxy logs the transaction and passes the request to the outside host. The proxy remains active until either side closes the connection.

The default service group and rules allow all inside hosts to use the multimedia proxies without authenticating. Multimedia services are included in the trusted service group, allowing users on the trusted network to view and access multimedia content from outside hosts. The default configuration does not allow outside hosts to connect to multimedia servers inside the perimeter.

This prohibits running multimedia servers on the firewall itself. Because the proxies are running on the default ports for these services on the firewall, all requests to these ports access the proxy. There is no way to start the server daemons needed for these multimedia requests.

Accessing Multimedia Proxy Configuration

To access the NetShow proxy configuration:

1. From within the Gauntlet Firewall Manager, select Services.
2. Click the NetShow tab.

The NetShow window displays.

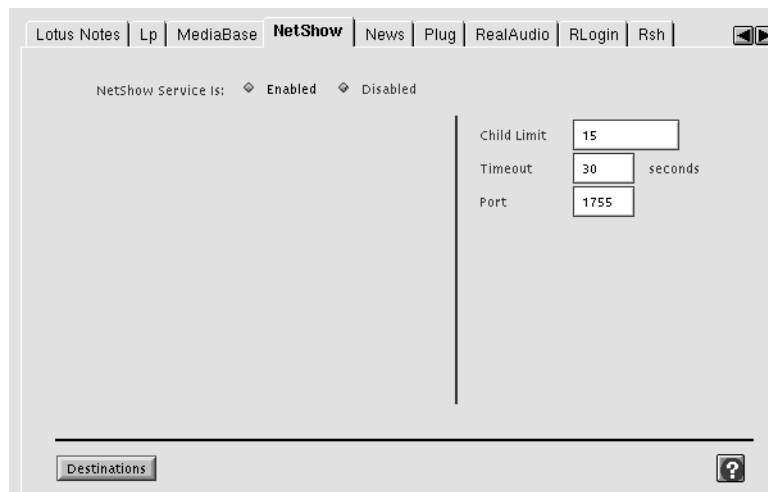


Figure 12-1 NetShow Window

To access the proxy configuration for RealAudio and RealVideo (RealPlayer):

1. From within the Gauntlet Firewall Manager, select Services.
2. Click the RealAudio tab.

The RealPlayer window displays.

To access the StreamWorks proxy configuration:

1. From within the Gauntlet Firewall Manager, select Services.
2. Click the StreamWorks tab.

The StreamWorks window displays.

To access the VDOLive proxy configuration:

1. From within the Gauntlet Firewall Manager, select Services.
2. Click the VDOLive tab.

The VDOLive window displays.

Configuring the Firewall for Multimedia Services

Configuring the Gauntlet Firewall involves planning, configuring the proxies to enforce company security policy, and enabling the proxy.

Planning the Firewall for Multimedia Services

When planning the Firewall for multimedia services:

- Determine which of the multimedia services you wish to allow.
- Determine if individual access controls will be applied to specific hosts or subnets. Obtain the hostname or IP address information of the allowed or restricted clients. If you choose to do this, plan to create specific service groups and rules for these activities. Alternatively, unrestricted access may be allowed.
- Determine if individual access controls to external servers will be applied, limiting those external sources users may access. Obtain hostname or IP address information for each server. If you choose to do this, plan to create destination access rules to deny these sites. Alternatively, unrestricted external server access may be allowed.

- Determine if you need to use non-default ports on the firewall for handling requests.
- Determine how many users of each proxy service are allowed to go through the firewall at the same time.

Configuring Multimedia Proxy Services

Configure the multimedia proxies to enforce company security policies.

To configure Multimedia proxy settings, you can provide optional information about configuration settings for the multimedia proxy. Refer to the online help for specific information about the available settings.

Enabling Multimedia Proxy Services

Even though the multimedia proxy services understand several different protocols, you do not need to enable all of them. For example, you can enable the RealPlayer proxy, and leave the NetShow and VDOLive proxies disabled.

To enable the proxy service:

1. In the appropriate multimedia configuration tab, click Enabled.
2. Add the proxy configuration to the service groups that you want to allow use of the proxy.
3. Before exiting the Gauntlet Firewall Manager, save and apply your changes.
The firewall enables the proxy.

Verifying Your Setup

Verify the installation by using the player to listen to audio files or view video files from hosts on the outside network. See the section below for instructions:

- “Using the NetShow Proxy” on page 114
- “Using the RealPlayer Proxy” on page 115
- “Using the StreamWorks Proxy” on page 117
- “Using the VDOLive Proxy” on page 118

Using the NetShow Proxy

Most users and most sites do not need to change the way they access NetShow files after installing the proxy.

Using the NetShow Proxy With Transparency

If the firewall uses transparency (the default configuration), you do not need to change the way you access NetShow files. As you did before, simply access the site and listen to audio files and view video files.

Using the NetShow Proxy Without Transparency

If the firewall does not use transparency, you can still access NetShow files. The NetShow player does not support NetShow proxies. However, it does support HTTP proxies. If you are not using transparency, you can access NetShow files by passing these files through the HTTP proxy.

To configure the NetShow player:

1. From within the NetShow player, select File.
2. Select Properties.
3. Select the Advanced tab.

The Advanced tab displays.

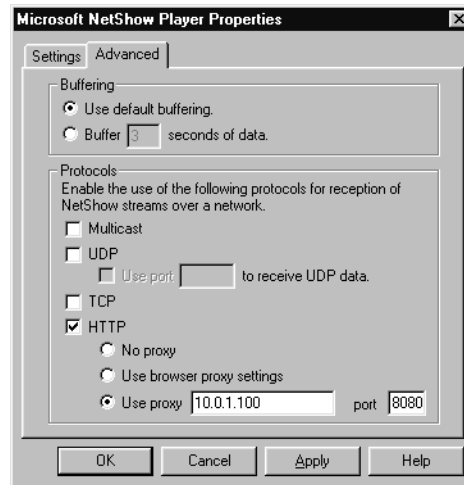


Figure 12-2 Microsoft NetShow Player Advanced Properties Window

4. Enter information about the HTTP proxy:

Protocol Select HTTP so that the NetShow player uses HTTP to transfer NetShow files.

Use Proxy Host name or IP address of the inside interface of the firewall.

Port Port number on which HTTP proxy is running (usually 8080).

5. Click *OK*.

Using the RealPlayer Proxy

Most users and most sites do not need to change the way they access RealPlayer files after installing the proxy.

Using the RealPlayer Proxy With Transparency

If you are using transparency on the firewall (the default configuration) and you have installed the RealPlayer proxy on the default player port (TCP port 7070), you do not need to change the way you access RealPlayer (RealAudio and RealVideo) files. As you did before, simply access the site and listen to audio files and view video files.

Using the Proxy for Real Audio and Real Video Proxy Without Transparency

If you are not using transparency, you need to configure RealPlayer to know about the proxy.

To configure the RealPlayer:

1. From within the RealPlayer, click View.
2. Select Preferences.
3. Click *Proxy*.

The Proxy tab displays.

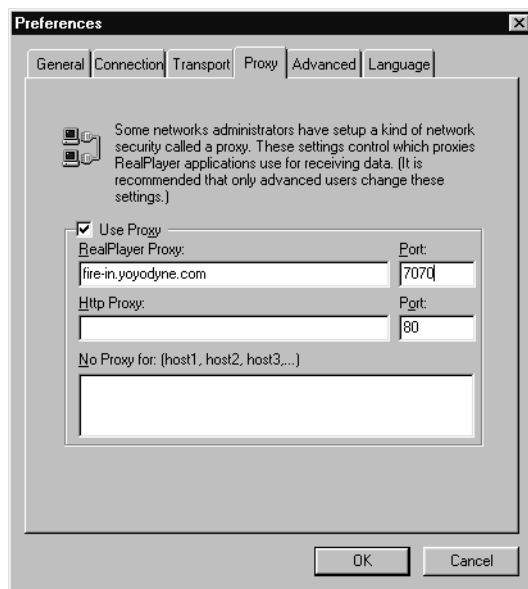


Figure 12-3 RealPlayer Proxy Preferences Window

4. Check the Use Proxy box.

5. Enter information about the RealPlayer proxy.

RealPlayer Proxy Host name or IP address of the inside interface of the firewall.

Port Port number on which the proxy is running. This is usually 7070.

6. Click OK.

Using the StreamWorks Proxy

The StreamWorks proxy does not support transparency; therefore, you need to configure the StreamWorks player to know about the proxy.

To configure the StreamWorks player:

1. From within the StreamWorks player, select Setting.
2. Select Network.

The StreamWorks Player Network Settings window displays.

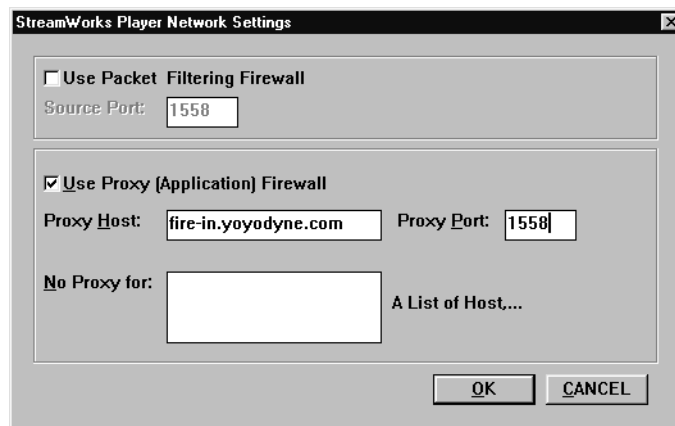


Figure 12-4 StreamWorks Player Network Settings Window

3. Check Use Proxy (Application) Firewall.

4. Enter information about the StreamWorks proxy:
 - Proxy Host Host name or IP address of the inside interface of the firewall.
 - Port Port number on which the proxy is running. This is generally 1558.
5. Click *OK*.

Now, when you point your Web browser or StreamWorks player at a StreamWorks file, it will use the proxy.

Using the VDOLive Proxy

Most users and most sites do not need to change the way they access VDOLive files after installing the proxy.

If you are using transparency on the firewall (the default configuration) and you have installed the VDOLive proxy on the default player port (TCP port 7000), you do not need to change the way you access VDOLive files. Access and use the VDOLive files as you did before.

If you are not using transparency, you need to configure the VDOLive player to know about the proxy.

To configure the VDOLive player:

1. From within the player, click *Setup*.
2. Click *Settings*.

The *Setup* tab displays.

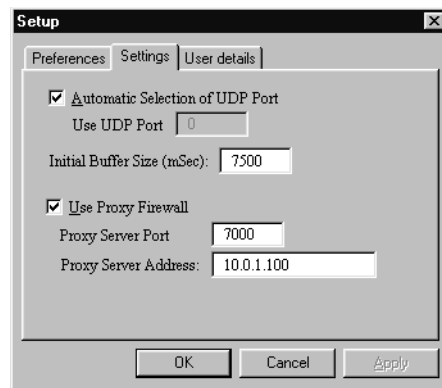


Figure 12-5 VDOLive Setup Settings Window

3. Check Use Proxy Firewall.
4. Enter information about the VDOLive proxy:
 - Proxy Server Port Port number on which the proxy is running; usually 7000.
 - Proxy Server Address Host name or IP address of the inside interface of the firewall.
5. Click *OK*.

Managing Network Management Services

Network management and monitoring are crucial in today's increasingly complex and heterogeneous network environment. Swift detection and response to a failing mission-critical networked resource can prevent considerable financial losses. Most of today's network management platforms use the industry-standard Simple Network Management Protocol (SNMP) to communicate with resources being managed. The Gauntlet Firewall includes an SNMP proxy that allows the network management station to communicate securely with the managed resources across the firewall.

The Gauntlet Firewall also includes SNMP agent software. Refer to Chapter 32, "Managing the Network Management Agent," on page 305 for more information about the SNMP agent.

This chapter discusses the concepts behind the SNMP proxy and explains how it works, how to configure it, and how to start it. The chapter consists of these sections:

- "Understanding the SNMP Proxy" on page 121
- "How the SNMP Proxy Works" on page 122
- "Accessing SNMP Proxy Configuration" on page 123
- "Configuring the Firewall for SNMP Services" on page 124
- "Configuring SNMP Agents" on page 127

Understanding the SNMP Proxy

The SNMP proxy is an application-level proxy that provides configurable access control and logging of SNMP traffic. The SNMP proxy passes network management station (manager) requests to the managed resources (agents), and accepts trap requests from agents using rules you supply.

Using the SNMP configuration rules, you can configure the firewall to allow a manager to query specific agents, and accept traps from selected agents. You can configure the proxy to allow connections based on:

- agent
- operation

The proxies log all successful and unsuccessful connection attempts and the amount of data transferred.

These access controls allow you to have more control over the connections to and from your system than you had without a firewall. The logging capabilities are also much more extensive.

How the SNMP Proxy Works

The firewall runs the SNMP proxy (*snmp-gw*) as a daemon listening for requests on the default SNMP port (UDP port 161) and agent trap port (UDP port 162). These ports are configurable.

SNMP Requests

When the firewall receives SNMP messages on the SNMP port, the proxy assumes the request is from the SNMP network manager. The proxy compares the name of the requesting host with the name configured on the firewall as the network manager. If the request came from a different host, the SNMP proxy will log the request and drop the packet. If the name of the requesting host matches the name of the firewall manager, the proxy will check its configuration to see if the requesting manager is allowed to use the operation to the requested agent.

If the request is for an operation or agent that is not permitted, the SNMP proxy will log the request and drop the packet. If the request is for an operation and agent that are permitted, the proxy will forward the request to the SNMP agent.

When the agent responds, the proxy verifies that the response came from the agent and that the packet is a valid SNMP response packet. The proxy then forwards the response back to the SNMP manager.

Because the SNMP proxy forwards the manager's request to the agent (after permission checks) and doesn't proxy the request on behalf of the SNMP manager, the SNMP request can work only when the firewall is using transparency. Therefore, the SNMP proxy works only when the SNMP manager is on the inside network and agents are on the outside or service network with the default firewall configuration.

SNMP Trap Requests

When the firewall receives an SNMP trap on the agent trap port, the proxy checks its configuration information to determine whether the agent is allowed to send traps through the firewall. If the agent is not allowed to send traps through the firewall, the firewall logs the trap request and drops the packet. If the agent is allowed to send traps through the firewall, the firewall sends the trap onto the SNMP network manager.

If the proxy is operating transparently, the firewall sends the trap to the SNMP network manager specified in the trap. If the proxy is not operating transparently, the firewall sends the trap to the SNMP network manager specified on the firewall. In the default configuration, the SNMP manager is on the trusted network and can transparently access the agents, which are on the untrusted networks.

Accessing SNMP Proxy Configuration

To access the SNMP proxy configuration:

1. From within the Gauntlet Firewall Manager, select Services.
2. Click the SNMP tab.

The SNMP window displays.

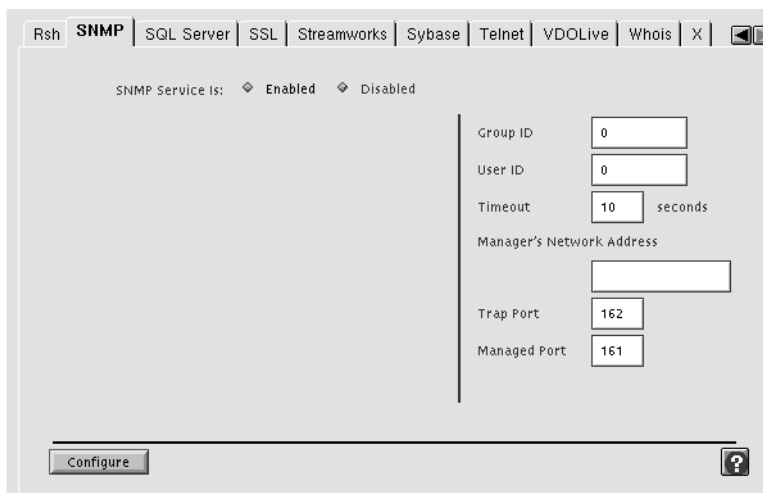


Figure 13-1 SNMP Window

Configuring the Firewall for SNMP Services

Configuring the Gauntlet Firewall involves planning, configuring proxy settings, and enabling the proxy.

Planning the Firewall for SNMP Services

When planning the firewall for SNMP services:

- Identify the network manager that will use the proxy.
- Determine the hostname or IP address of the resources (agents) that the network manager will manage across the firewall and the SNMP operation restrictions on each agent.
- Determine the agent response time-out that the network manager will use. The firewall uses a default value of 10 seconds. Consider setting the time-out to the same value used by the SNMP manager.
- Determine the SNMP Port and agent trap port that will be used.

Configuring SNMP Proxy Settings

Configure the SNMP proxy to enforce company security policies.

To configure SNMP proxy settings:

1. In the SNMP window, provide information about the proxy and the network manager that can use the proxy.

Manager's IP address or host name of the system that is the network Network Address manager. Wildcards * are not valid.

2. Click Configure.

The Add SNMP Agents window displays.

Select to modify or delete, or enter new SNMP Agents below:

Agent Host/IP	Access	Get	Set	Trap
---------------	--------	-----	-----	------

Currently selected or new Agent information

HostName:

IP Address:

Access: Permit Deny

Operations allowed on this Agent:

- Get Operations
- Set Operations
- Receive Traps from Agent

Add Modify Delete

OK Cancel ?

Figure 13-2 Add SNMP Agents Window

3. Provide information about the agents the network manager can contact and the types of operations the manager can perform.
 - Hostname Host name of an agent. Specify by host or network. The wildcard * is valid. If you enter a host name, you do not need to enter an IP address.
 - IP Address EIP address of an agent. Specify by host, subnet, or network. The wildcard * is valid. If you enter an IP address, you do not need to enter a host name.
 - Access Specify whether these rules permit or deny access to and from this agent. Remember, “That which is not explicitly permitted is denied.” Once you create a deny rule for one agent, you must create explicit permit rules for other agents.
 - Operations Types of operations that the network manager can exchange with this agent.
4. Click *Add*.
5. Click *OK*.

Note: Some network management software (such as HP OpenView and CA Unicenter TNG) use the *ping* program to verify that an agent is reachable before they send the actual SNMP request. By default, the Gauntlet Firewall does not allow *ping* traffic, or any ICMP traffic through the firewall. If your network management software uses the *ping* program, you must add packet screening rules to the firewall. These rules must allow ICMP between the network manager and the agent.

Enabling SNMP Proxy Services

To enable the SNMP proxy service:

1. In the SNMP window, click Enabled.
2. Add the SNMP proxy configuration to the service groups you want to use the SNMP proxy.
3. Before exiting the Gauntlet Firewall Manager, save and apply your changes.
The firewall enables the SNMP proxy.

Configuring SNMP Agents

If the agent cannot transparently reach the network manager, you must configure the agent so that it accepts requests from the firewall. This is the default when the agent is outside the firewall and the network manager is inside the firewall.

Managing News Services

Usenet news continues to be one of the most widely used Internet applications. Many sites rely on Usenet news for information on the latest technology. Although the Network News Transfer Protocol (NNTP) does little in comparison to other network protocols, you must configure it carefully to protect internal news groups that may contain sensitive proprietary information.

The plug proxy included with the Gauntlet Firewall allows administrators to tunnel NNTP-based news feeds through their firewall. The NNTP connections come from known sites (as opposed to the multitude of sites that may connect via SMTP to deliver mail). NNTP is also a very straightforward protocol. For these reasons, it can be proxied using the generic plug proxy.

The following sections discuss the concepts behind the News proxy and explain how it works and how to configure the proxy for NNTP-based news:

- “Understanding the News Proxy” on page 130
- “How the News Proxy Works” on page 130
- “Configuring the Firewall for News Services” on page 131
- “Configuring Your News Server” on page 134
- “Using News” on page 135

Understanding the News Proxy

The Gauntlet plug proxy is a TCP gateway that provides configurable access control and logging mechanisms. The plug proxy, which runs on the firewall, passes NNTP or other application requests through the firewall, using rules you supply. It essentially tunnels information from a port on the firewall to a specific port on another system.

You can configure the News proxy to allow connections based on:

- source IP address
- source hostname
- source port
- destination IP address
- destination hostname
- destination port

Using these options, you can configure the firewall to allow your service provider's host on the outside to connect to the firewall and pass news via NNTP to your news system on the inside network. You can also configure the firewall to allow a set of news clients on your internal system to connect to news servers on hosts outside your firewall.

The News proxy logs all successful and unsuccessful connection attempts and the amount of data transferred.

These access controls allow you to have more control over the connections to and from your system than without a firewall. The logging capabilities are also more extensive.

How the News Proxy Works

The firewall runs the News proxy as a daemon listening for requests on the standard NNTP port (TCP port 119). When the News proxy receives a request, the proxy checks its configuration information and determines whether the initiating host has permission to initiate this type of request. If the host has permission, the proxy passes the connection to the specified port on the specified system. This News proxy remains active until either side terminates the connection.

The default configuration for the Gauntlet Firewall allows requests to and from one internal news server and one external news server. The firewall itself cannot run an NNTP news server, or any other service that you are passing through the plug proxy, because the plug proxy is using the standard port for these services.

Hosts on both the inside and outside think the firewall is servicing requests. The external news server thinks it is feeding news to the firewall, and the internal news server thinks that it is receiving news from the firewall. The firewall is simply acting as the tunnel, via the plug proxy.

Another common configuration allows request from internal news clients to multiple news server. Again, the firewall is simply acting as a tunnel.

Configuring the Firewall for News Services

Configuring the Gauntlet Firewall for news services involves planning, configuring the firewall, configuring the proxy to enforce company security policy, and enabling the proxy.

Planning News Settings

When planning news settings:

- Do not use the firewall as a news server.
- Restrict or disallow automatic group creation and deletion.
- Allow external NNTP connections from known servers only.

Configuring News Settings

To configure News Feed settings:

1. From within the Gauntlet Firewall Manager, select Environment.
2. Click the News tab.

The News window displays.

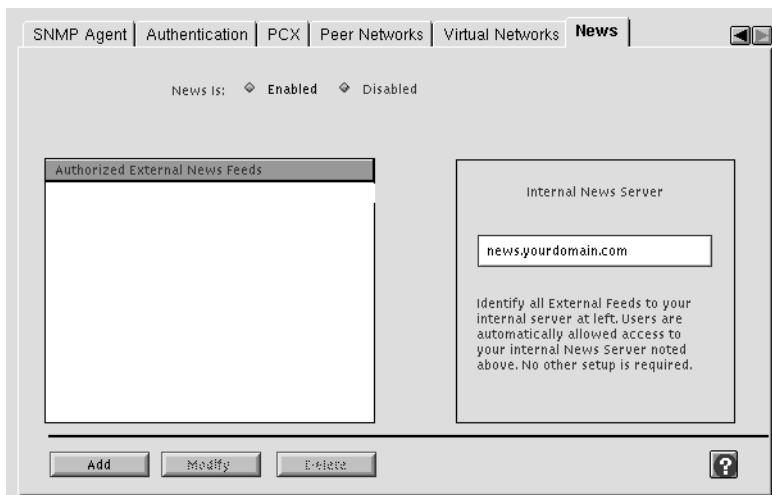


Figure 14-1 News Window

3. Provide information about your internal news server.
Internal News Server IP address or host name for your internal NNTP news server.
4. Click *Add*.
The News Feed Server Identification Screen displays.




Figure 14-2 News Feed Server Identification Screen

5. Provide information about your external news feeds.
External News Feed IP address or host name for the system that provides your external NNTP news server.
6. Click *OK*.

To configure News Reader settings:

1. From within the Gauntlet Firewall Manager, click Services.
2. Click the News tab.

The News window displays.

3. Indicate whether news readers can connect to any external news server or only to specific servers.
4. If you wish to limit the external news servers, click *Add*.

The External News Server window displays.

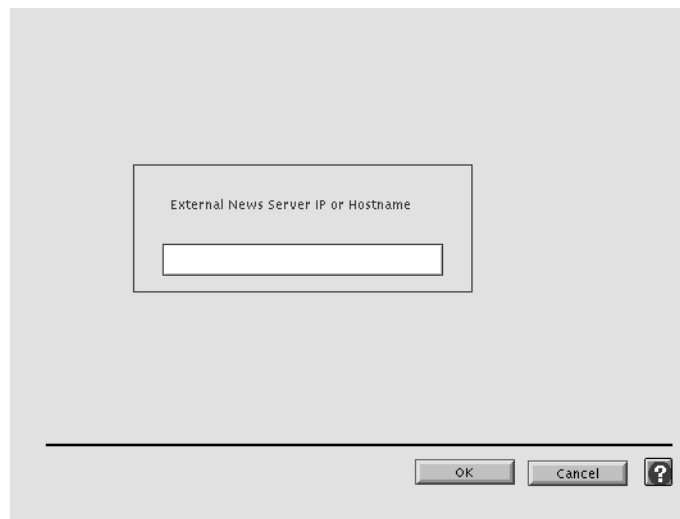


Figure 14-3 External News Server Window

5. Provide information about allowed external news servers.
News Server IP address or host name of an external NNTP news server.
6. Click *OK*.

Enabling News Proxy Services

Enabling the News reader proxy service also enables the News feed proxy service. If you want to use only one of these services, do not configure settings for the other service.

Because the firewall follows the rule “That which is not expressly permitted is denied,” you are not leaving the firewall vulnerable to attack if you only configure one type of News service.

To enable the News proxy service:

1. On either of the News configuration tabs, click Enabled.
2. Add the News proxy configuration to the service groups that you want to use the News proxy.
3. Before exiting the Gauntlet Firewall Manager, Save and Apply your changes.

The firewall enables the News proxy.

Informing Your News Feed

If you are using the News feed configuration of the proxy, inform your external news feed provider (often your Internet service provider) that tall NNTP news should be sent to the firewall (instead of the internal news server). Be sure to provide them with the outside IP address for your firewall.

Configuring Your News Server

If you are using the News feed configuration of the proxy, configure your internal news server software to transfer and receive articles from the firewall, rather than your external news server. Specify the inside IP address for the firewall.

Using News

The firewall and the plug proxy for NNTP traffic are transparent to the user. If you are using the News feed configuration, users should continue to point their news readers (*rn*, *trn*) or other news-aware tools (Netscape Navigator, Microsoft Internet Explorer) towards your internal news server.

If you are using the News Reader configuration, users should continue to point their news readers or other news-aware tools towards the appropriate news server on the outside of the firewall.

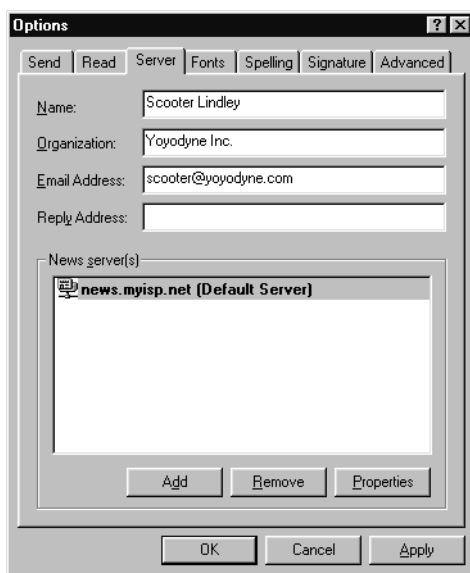


Figure 14-4 News Reader Window

Managing Print Services

Printing continues to be a widely used feature of most computer networks. In some circumstances, users need to print information using printers connected to other systems on other networks. Users behind a firewall might want to print to printers on systems on the outside, or behind other firewalls. Others might want to be able to print from a remote system, for example from a mobile PC to a printer behind a firewall. The Gauntlet Firewall includes an *lp* proxy that securely handles the transfer of print requests.

This chapter discusses the concepts behind the *lp* proxy and explains how *lp* works, how to configure the proxy, and how to use *lp* services. The chapter consists of these sections:

- “Understanding the *lp* Proxy” on page 137
- “How the *lp* Proxy Works” on page 138
- “Configuring the Print Client” on page 140
- “Configuring the Print Server” on page 140
- “Accessing *lp* Proxy Configuration” on page 141
- “Configuring the Firewall for *lp* Services” on page 142
- “Using *lp* Services” on page 144

Understanding the *lp* Proxy

The *lp* proxy is an application-level gateway that provides configurable access control and logging mechanisms. The *lp* proxy, which runs on the firewall, passes *lp* requests through the firewall using rules you supply. You can configure the *lp* proxy to allow file transfer activity based on:

- source IP address
- source hostname
- destination IP address

- destination hostname
- *lp* commands (for example, number and priority)
- printer queue

Using these options, you can configure your firewall to allow specific hosts on the inside network to print files on outside hosts. Employees working behind the firewall can send print jobs to printers at customer sites. Similarly, traveling employees can send print jobs to printers at corporate headquarters inside the defense perimeter. You can deny access to some *lp* commands, allowing users to print, but not allowing them to restart or remove print jobs.

The *lp* proxy logs all successful and unsuccessful file transfer attempts and the number of bytes transferred. These access controls allow you to have much more control over the files entering and leaving your system than using the standard IRIX *lp* program. The logging capabilities are also much more extensive.

How the *lp* Proxy Works

The firewall runs the *lp* proxy (*lp-gw*) as a daemon listening for requests on the standard printer port (TCP port 515). When the firewall receives requests for services on this port, the *lp* proxy checks its configuration information and determines whether the initiating host has permission to use *lp*. If the host has permission, the proxy logs the transaction and passes the request to the printer server. The *lp-gw* proxy remains active until either side closes the connection.

The default configuration allows inside hosts to use *lp*. Users on inside hosts can continue to print to outside hosts as they did before the firewall was put into place. The default configuration does not allow outside hosts to connect to inside hosts for printing.

This configuration prohibits running an *lp* server on the firewall itself. Because the *lp* proxy is running on the standard *lp* port on the firewall, there is no way to start the *lp* daemon needed to service *lp* requests. Thus, you cannot print from the firewall itself.

However, a common configuration is to allow outside hosts to print to inside printers. Consider Yoyodyne's network, as shown in Figure 15-1. The Web server is on the outside network. When working on the web, the Webmaster needs to print files. It is possible to print files to the printer inside the firewall if the system outside the firewall, the firewall, and the system controlling the printer inside the firewall are all properly configured.

To configure this service, the administrator at Yoyodyne creates a remote print queue (called `inside_fw`) on the Web server that indicates that files to the print queue inside should be sent to the outside address of the firewall (204.255.154.1). On dimension, the system inside the firewall that controls the printer, the administrator creates a print queue (called `from_fw`) that sends requests to the printer (10.0.1.40). The administrator also configures dimension to accept requests from the inside address of the firewall (10.0.1.100). On the firewall, the administrator configures the `lp` proxy to take requests sent to the queue `inside_fw` and to send the requests to the queue `from_fw` on dimension.

Now, the Webmaster sends print jobs to the print queue `inside_fw`. The firewall acts as a relay and passes the print job to the print queue `from_fw` on dimension. The print queue `from_fw` sends the file on to the printer, which prints the job.

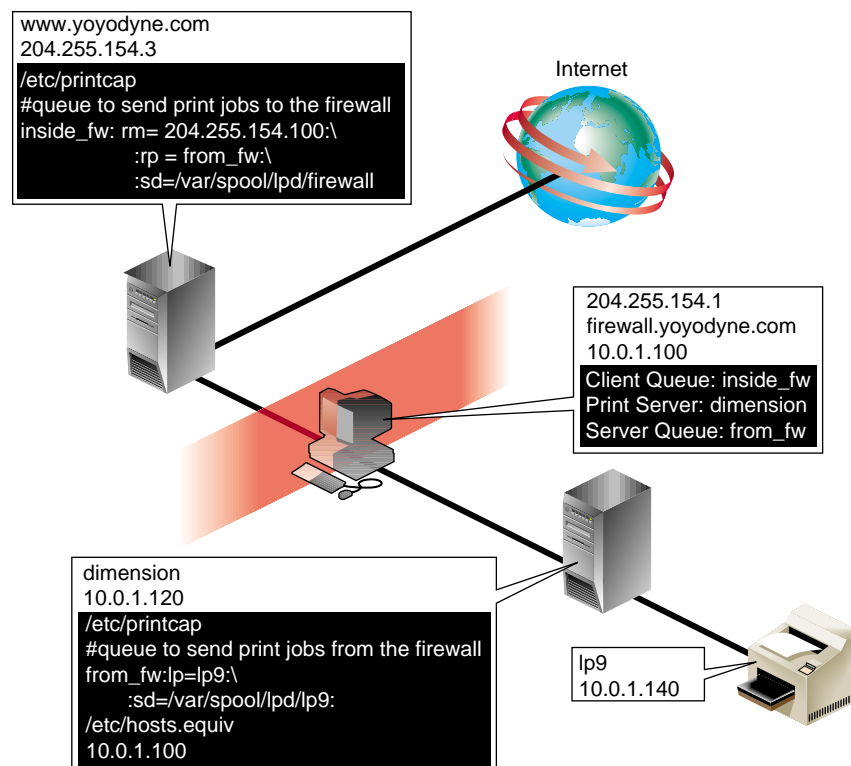


Figure 15-1 Example Firewall lp Configuration

Configuring the Print Client

Configuring the print client involves creating a remote print queue to send print jobs to the firewall. You create a print queue using Print Manager.

Note: Consult your IRIX system documentation for additional information.

To create a remote print queue using Print Manager, follow these steps:

1. Define the remote print queue using the Printer Manager in the Toolchest.
 - If you are not using transparency (as when printing from outside the firewall to inside the firewall), specify the host name of the firewall as the remote host.
 - If you are using transparency (the default configuration when printing from inside the firewall to outside the firewall), specify as the remote host the hostname of the system where the printer is connected.
2. Instruct users to print to the client queue name (`inside_fw`) to print to the remote printer queue.

Configuring the Print Server

Configuring the print server involves instructing the host running the print server to accept jobs from the firewall.

To configure the print server, follow these steps:

1. On the host running the print server, add the firewall to the lists of hosts that can print to the desired printer using *lpadmin*.
 - If the print server is inside the firewall, specify the IP address of the inside interface of the firewall:
`10.0.1.100`
 - If the print server is outside the firewall, specify the IP address of the outside interface of the firewall:
`204.255.154.100`

2. Create a print queue to send requests to the printer. You can create a separate print queue to handle requests from the firewall, or you can continue to use one of your standard print queues. Be sure that the print queue name matches the name you specified as the remote printer name on the client (*from_fw*).
3. Consult your IRIX system documentation for other standard steps in creating and starting print servers. Ensure that you create a print queue that actually will send print jobs to the printer.

Accessing Ip Proxy Configuration

To access the *lp* proxy configuration:

1. From within the Gauntlet Firewall Manager, select Services.
2. Click the Lp tab.

The Lp window displays.

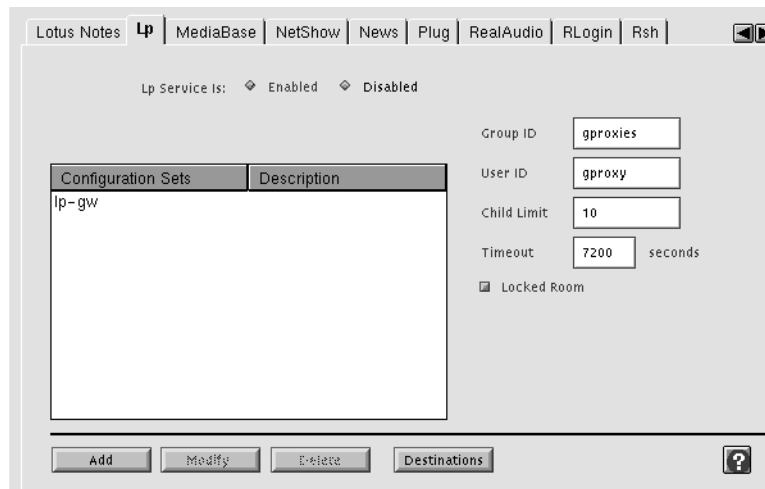


Figure 15-2 Lp Window

Configuring the Firewall for Ip Services

Configuring the Gauntlet Firewall involves planning, configuring the *lp* proxy to enforce your security policy, and enabling the proxy.

Planning Ip Proxy Settings

When planning *lp* proxy settings:

1. Determine which internal users and hosts can use these services.
2. Determine which external users and hosts can use these services.

Configuring Ip Proxy Settings

Configure the *lp* proxy to enforce your security policies.

To configure *lp* proxy settings:

1. In the Lp window, select the lp-gw configuration to modify the default settings.
2. Click *Modify*.

The Modify Printer Services window displays.

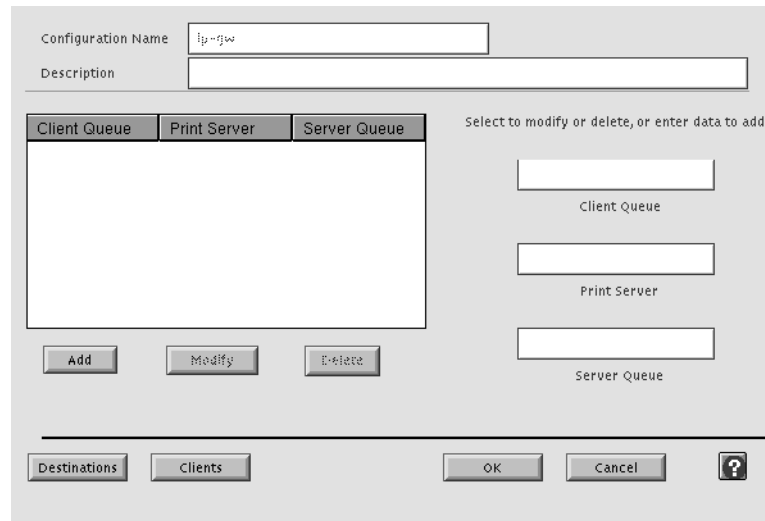


Figure 15-3 Modify Printer Services Window

3. Provide information about the printer queue to which the *lp* proxy sends requests.

Client Queue Name of the remote print queue that you defined on the print client. For example, Yoyodyne enters *inside_fw*.

Print Server IP address of the system running the *lp* daemon that will handle the print request. Specify by IP address or host name. If you are using transparency, this value is optional. For example, Yoyodyne enters "dimension."

Server Queue Name of the print queue on the print server. For example, Yoyodyne enters "from_fw."

4. Click *Add* to add this printer information.
5. Click *OK*.

Enabling the *lp* Proxy

To enable the proxy:

1. In the *Lp* window, click *Enable*.
2. Add the *lp* proxy configuration to the service groups that you want to use the *lp* proxy.
3. Before exiting the Gauntlet Firewall Manager, Save and Apply your changes.

The firewall enables the *lp* proxy.

Using *lp* Services

The firewall and the *lp* proxy are transparent to users. Ask them to print to the desired print queue as before.

Managing rsh Services

Administration and support activities can be easier when you can just execute a shell on a remote system. The *rsh* service allows users to do this. The *rsh* program is not without risks: it runs programs on another system and requires some privileges to log in. The Gauntlet Firewall includes a proxy that securely handles the execution of *rsh* requests from systems inside the network to systems outside the network.

The following sections explain the concepts behind the *rsh* proxy and how it works, how to configure the proxy, and how to use *rsh* services:

- “Understanding the rsh Proxy” on page 145
- “How the rsh Proxy Works” on page 146
- “Accessing rsh Proxy Configuration” on page 147
- “Configuring the Firewall for rsh Services” on page 147
- “Using rsh Services” on page 149

Understanding the rsh Proxy

The *rsh* proxy is an application-level gateway that provides configurable access control, authentication, and logging mechanisms. The *rsh* proxy, which runs on the firewall, passes *rsh* requests through the firewall, using rules you supply. You can configure the *rsh* proxy to allow remote shell activity based on:

- source IP address
- source hostname
- destination IP address
- destination hostname

Using these options, you can configure your firewall to allow specific hosts on the inside network to start remote shells on outside hosts. Employees working behind the firewall can start remote shells on outside hosts at a customer site. The *rsh* proxy logs all successful and unsuccessful remote shell attempts, and the number of bytes transferred.

These access controls allow you to have much more control over the *rsh* requests entering and leaving your system than using the standard IRIX *rsh* program. The logging capabilities are also much more extensive.

How the rsh Proxy Works

In this default configuration, the firewall runs the *rsh* proxy (*rsh-gw*) as a daemon listening for requests on the standard *rsh* port (TCP port 514). Whenever the system receives an *rsh* request on this port, the *rsh* proxy checks its configuration information and determines whether the initiating host has permission to use *rsh*. If the host has permission, the proxy logs the transaction and passes the request to the outside host. The *rsh-gw* remains active until either side closes the connection.

The default trusted service group does not include the *rsh* proxy. If you add the *rsh* proxy to the trusted service group and enable the proxy, users on inside hosts can continue to use *rsh* as they did before the firewall was put into place. The default untrusted service group and rules does not allow outside hosts to use *rsh* to hosts inside the perimeter.

The default configuration using just the *rsh* proxy prohibits running an *rsh* server on the firewall itself. Because the *rsh* proxy is running on the standard *rsh* port on the firewall all *rsh* requests start the proxy. There is no way to start the *rsh* daemon needed to service *rsh* requests.

Accessing rsh Proxy Configuration

To access the *rsh* proxy configuration:

1. From within the Gauntlet Firewall Manager, select Services.
2. Click the Rsh tab.

The Rsh window displays.

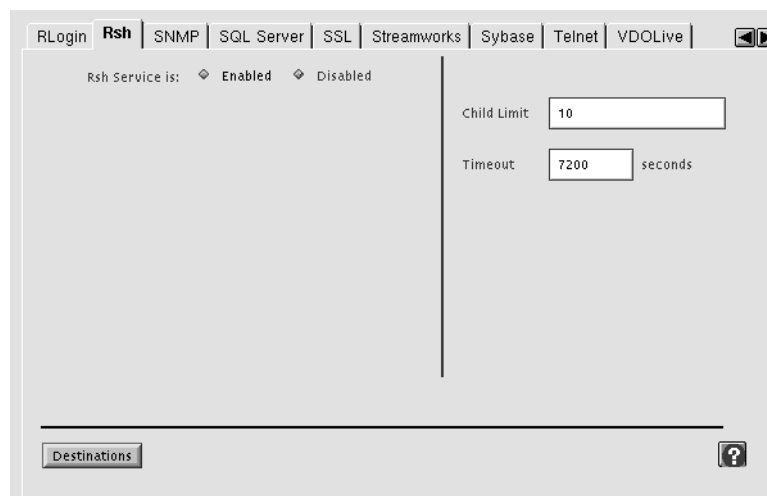


Figure 16-1 Rsh Window

Configuring the Firewall for rsh Services

Configuring the Gauntlet Firewall involves planning, configuring the *rsh* proxy to enforce your security policy, and enabling the proxy.

Planning rsh Proxy Settings

When planning *rsh* proxy settings, determine whether you wish to allow inside hosts to *rsh* through the firewall to outside hosts.

Configuring rsh Proxy Settings

Configure the *rsh* proxy to enforce your security policies.

To configure *rsh* proxy settings, you can provide optional information about time-out values and other configuration settings for the *rsh* proxy.

Enabling rsh Proxy Services

To enable the *rsh* proxy service:

1. In the Rsh window, click Enabled.
2. Add the *rsh* configuration to the service groups that you want to use the *rsh* proxy.
3. Before exiting the Gauntlet Firewall Manager, save and apply your changes.

The firewall enables the *rsh* proxy.

Verifying Your Setup

Verify your configuration by accessing a system outside the perimeter from a system inside the perimeter.

Using rsh Services

Following some initial configuration, the firewall and the *rsh* proxy are transparent to the user. Users can continue to use *rsh* to connect to outside hosts as they did before.

Configuring the Remote System

Before using *rsh*, users must configure the remote host to accept *rsh* requests from the firewall.

To configure the remote host, have each user edit their *.rhosts* file on the remote system and add their user name and the name of the firewall to their *.rhosts* file on the remote system:

```
firewall [user]
```

- *firewall* is the name (including domain if necessary) of the firewall. This name should be the name of the interface on the firewall closest to the remote system.
- *user* is their user name within the domain from which the request comes. The user does not actually need to have an account on the firewall itself. The *rsh* request simply appears to be coming from the firewall.

For example, Penny, who works at Yoyodyne, needs to execute something remotely using her account at Big University. She adds this line to the *.rhosts* file in her account at Big University:

```
fire-out.yoyodyne.com penny
```

Managing Sybase Services

Database services are essential in most organizations. As with other services you offer, you want to configure database access securely. Sybase is a relational database management system used by many organizations. The Gauntlet Firewall includes a proxy that securely allows connections between Sybase clients on the inside network and servers on the outside network.

This chapter discusses the concepts behind the Sybase proxy and explains how it works, how to configure the proxy, and how to use Sybase services in the following sections:

- “Understanding the Sybase Proxy” on page 151
- “How the Sybase Proxy Works” on page 152
- “Accessing Sybase Proxy Configuration” on page 153
- “Configuring the Firewall for Sybase Services” on page 154
- “Configuring Sybase Clients” on page 156
- “Verifying Your Setup” on page 156

Understanding the Sybase Proxy

The Sybase proxy is an application-level proxy that provides configurable access control, authentication and logging mechanisms. The Sybase proxy, which runs on the firewall, passes Sybase requests through the firewall (at the application level), using rules you supply. You can configure instances of the Sybase proxy to service:

- Sybase client-to-server communications
- Sybase server-to-server communications

For each version of the Sybase proxy, you can configure the proxy to allow connections based on:

- source IP address

- source hostname
- source port
- destination IP address
- destination hostname
- destination port

Using these options, you can configure your firewall to allow Sybase clients on certain trusted hosts to access a Sybase server on an untrusted host. Employees working behind the firewall can access Sybase databases at customer sites. You can also configure your firewall to allow Sybase servers on opposite sides of the firewall to communicate. A Sybase replication server can communicate with another Sybase replication server on the other side of an intranet firewall.

You can configure the Sybase proxy to allow Sybase clients on untrusted hosts to access Sybase servers on your trusted networks. According to most security policies, including the Gauntlet Firewall default, it is not a good idea to allow such access to untrusted hosts. If you must allow this sort of service, consider using client-side password encryption. Consider limiting the databases and data to which users have access, because all of the data is transferred unencrypted.

The proxies log all successful and unsuccessful connection attempts and the amount of data transferred.

These access controls allow you to have more control over the connections to and from your system than you have without a firewall. The logging capabilities are also more extensive.

How the Sybase Proxy Works

The firewall runs different instances of the Sybase proxy (*syb-gw*) as daemons on different ports for different Sybase applications. Whenever the firewall receives a Sybase request on one of these ports, the Sybase proxy checks its configuration information and determines whether the initiating host has permission to initiate this type of request. If the host does not have permission, the Sybase daemon logs the connection attempt and displays an error message.

If the host has permission, the proxy logs the transaction and passes the request to the destination host. The Sybase proxy remains active until either side closes the connection.

The default service groups do not allow either inside or outside hosts to use the Sybase proxy. The recommended configuration allows trusted hosts to access Sybase servers on untrusted networks. The recommended configuration does not allow untrusted hosts to access Sybase servers on trusted networks. While the Sybase proxy does perform checks to ensure that the packets appear to be Sybase packets, someone could spoof this protocol. The Sybase proxy does not perform any user authentication. You are relying on the authentication mechanisms of the Sybase server to control access to your Sybase server and its data.

Accessing Sybase Proxy Configuration

To access the Sybase proxy configuration:

1. From within the Gauntlet Firewall Manager, select Services.
2. Click the Sybase tab.

The Sybase window displays.

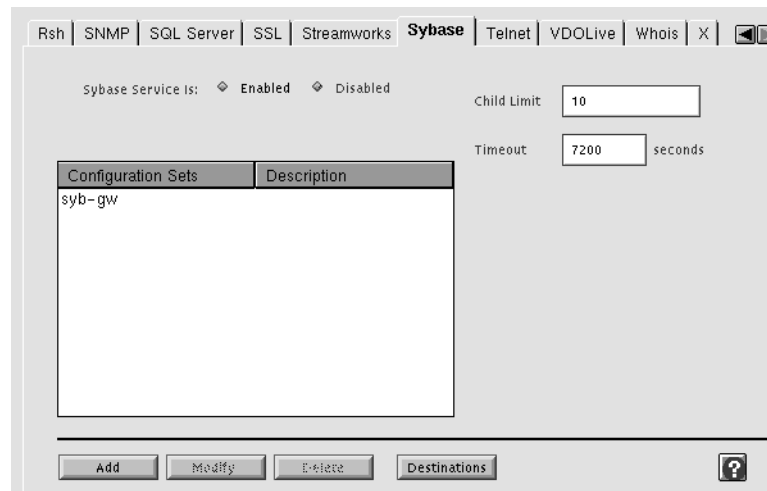


Figure 17-1 Sybase Window

Configuring the Firewall for Sybase Services

Configuring the Gauntlet Firewall involves planning, configuring the proxies to enforce your security policy, and enabling the proxy.

Planning Sybase Proxy Settings

When planning Sybase proxy settings:

1. Determine which Sybase servers users need to access. Determine whether you want to limit access to particular a server or not. Obtain hostname or IP address information for each server.
2. For each server, determine the port(s) on which the server accepts connections.
3. Determine which external hosts can use these services.
4. Determine which internal hosts can use these services.

Configuring Sybase Proxy Settings

Configure the Sybase proxy to enforce your security policies.

To configure Sybase proxy settings:

1. Select the syb-gw configuration to modify the default settings.
2. Click *Modify*.

The Modify Sybase Services window displays.

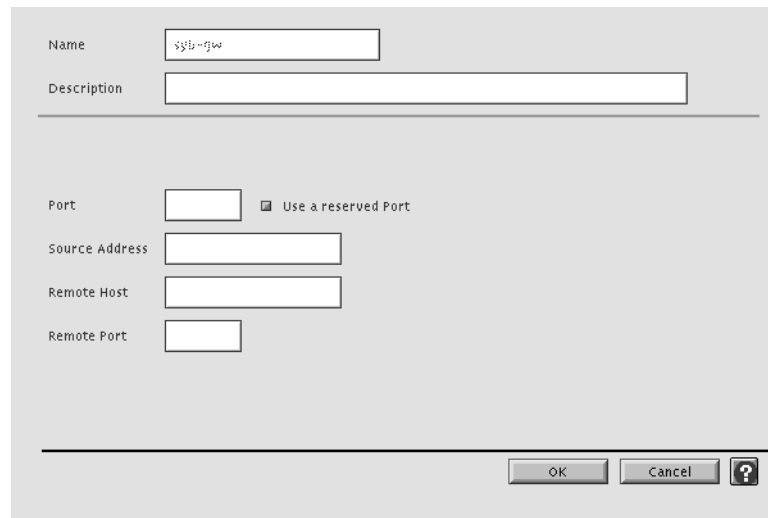


Figure 17-2 Modify Sybase Services Window

3. Provide information about the port on which you are running the Sybase service.
 Port Enter the port number on which the proxy runs.
4. Provide other optional information about the source and hosts for the Sybase proxy.
5. Click OK.

Enabling Sybase Proxy Services

To enable the Sybase proxy service:

1. On the Sybase window, click Enabled.
2. Add the Sybase proxy configuration to the service groups that you want to use the Sybase proxy.
3. Before exiting the Gauntlet Firewall Manager, save and apply your changes.
 The firewall enables the Sybase proxy.

Configuring Sybase Clients

Add or modify the interfaces file on the client (using a tool like *sybinit* or *SQLEdit*) to provide information about the Sybase server.

To configure Sybase clients:

1. Specify the port number you selected for the Sybase proxy.
2. If you are using transparency (the default configuration), specify the hostname as the hostname of the actual system running the Sybase server. If you are not using transparency, specify the hostname as the IP address of the firewall.

If you are using server-to-server communications, configure your servers as clients. Consult your Sybase administration documentation for further information on configuring clients for accessing servers.

Verifying Your Setup

Use your Sybase client on a trusted host to run a simple query against the Sybase server on the untrusted host. Watch the logs on the firewall for error messages.

Managing Terminal Services

Terminal service access to other computers can be a vital part of many network activities. The TELNET and rlogin protocols are used for making these terminal connections, but they are not without risk. The Gauntlet Firewall includes proxies for both the TELNET and rlogin protocols, which securely handle terminal services between the inside and outside networks.

This chapter discusses the concepts behind the TELNET and rlogin proxies and explains how they work, how to configure the proxies, and how to use terminal services. The chapter consists of these sections:

- “Understanding the TELNET and rlogin Proxies” on page 158
- “How the TELNET and rlogin Proxies Work” on page 159
- “Accessing TELNET and rlogin Proxy Configuration” on page 160
- “Configuring the Firewall for Terminal Services” on page 161
- “Using Terminal Services” on page 163

Understanding the TELNET and rlogin Proxies

The TELNET and rlogin proxies are application-level proxies that provide configurable access control, authentication, and logging mechanisms. The TELNET and rlogin proxies, which run on the firewall, pass TELNET and rlogin requests through the firewall, using rules you supply. The TELNET proxy also passes TN3270 requests through the firewall. You can configure the proxies to allow connections based on:

- source IP address
- source hostname
- destination IP address
- destination hostname

Using these options, you can configure your firewall to allow specific hosts on outside networks to connect to inside hosts or vice versa. Employees working at customer sites can access their workstations inside the perimeter.

The proxies allow administrators to require users to authenticate before connecting. The proxies log all successful and unsuccessful connection attempts and the amount of data transferred.

These access controls allow you to have much more control over the connections to and from your system than using the standard IRIX TELNET and rlogin programs. The logging capabilities are also much more extensive.

Note: You can use the TELNET proxy without the rlogin proxy, or rlogin without TELNET. You can configure different policies for hosts and authentication as well.

How the TELNET and rlogin Proxies Work

The firewall runs the network access control daemon (*netacl*) as a daemon listening for requests on the standard TELNET port (TCP port 23). Whenever the firewall receives a TELNET request on this port, the *netacl* daemon checks its configuration information and determines whether the initiating host has permission to use TELNET. If the host has permission, the *netacl* daemon starts the standard TELNET program (*telnetd*) or the TELNET proxy (*tn-gw*) depending upon the originating host. If the host does not have permission, the daemon displays an error message. Similarly, the *netacl* daemon running on the standard login port (TCP port 513) starts either the rlogin program (*rlogind*) or the rlogin proxy (*rlogin-gw*).

The default trusted service group and rules allow all inside hosts to initiate TELNET or rlogin sessions without authenticating. The inside host passes TELNET requests to the firewall, which starts the *netacl* daemon. The *netacl* daemon checks its permissions, and determines that the inside host can use TELNET. The *netacl* daemon starts the proxy. The proxy logs the transaction and passes the request to the outside host. The proxy remains active until either side closes the connection. The default untrusted service group allows outside hosts to initiate TELNET or rlogin sessions after authenticating. The outside host passes TELNET requests to the firewall, which starts the *netacl* daemon. The *netacl* daemon checks its permissions, and determines that the outside host can use TELNET. The *netacl* daemon starts the proxy. The proxy prompts the user for authentication. If it is successful, the proxy prompts the user for the inside host, logs the transaction, and passes the request to the inside host. The proxy remains active until either side closes the connection.

Note that users are not logging into the firewall directly. While users use the proxy on the firewall for authentication, the proxy simply passes the user's TELNET or rlogin session on to the appropriate host.

A configuration using *netacl* allows administrators on either inside or outside hosts to initiate TELNET requests to the firewall. The *netacl* daemon checks its permissions and determines that the host can use TELNET. The *netacl* daemon starts the proxy. The proxy prompts the user for authentication. If it succeeds, the proxy prompts the user for the host and logs the transaction. When users indicate they want to connect to the firewall itself, the *netacl* daemon reviews the source and starts the actual TELNET daemon.

Accessing TELNET and rlogin Proxy Configuration

To access the TELNET proxy configuration:

1. From within the Gauntlet Firewall Manager, select Services.
2. Click the TELNET tab.

The TELNET window displays.

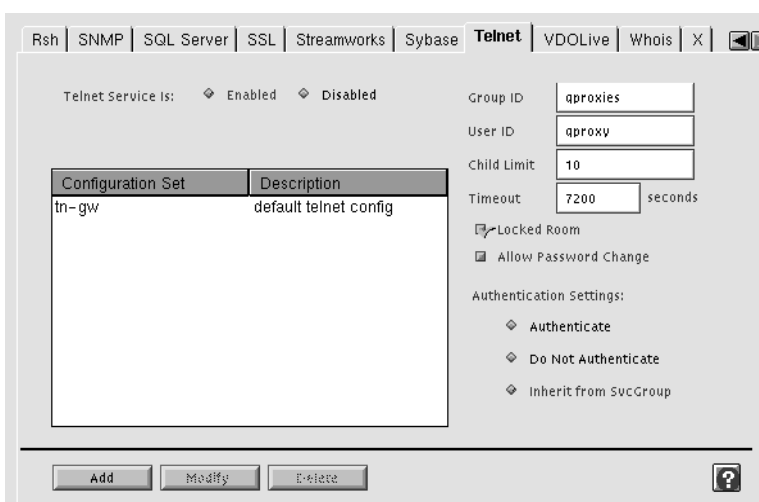


Figure 18-1 TELNET Window

To access the rlogin proxy configuration:

1. From within the Gauntlet Firewall Manager, select Services.
2. Click the Rlogin tab.

The Rlogin window displays.

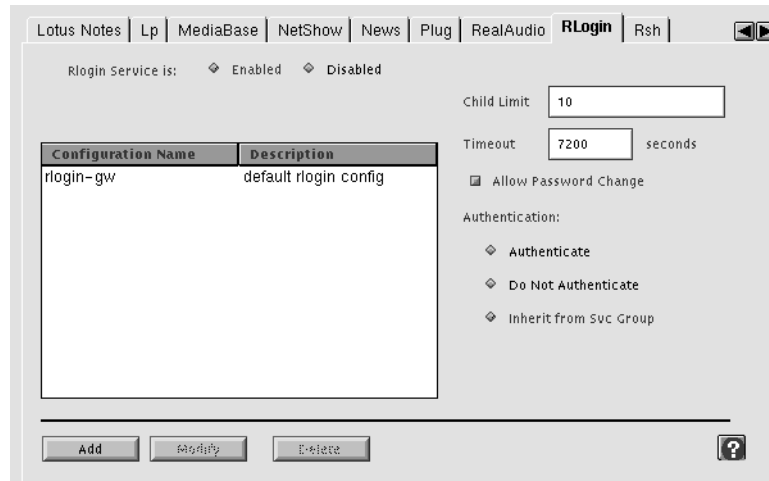


Figure 18-2 rlogin Window

Configuring the Firewall for Terminal Services

Configuring the Gauntlet Firewall involves planning, configuring the proxies to enforce your security policy, enabling the proxy, and creating user accounts for users who will need to authenticate.

Planning TELNET and rlogin Proxy Settings

When planning TELNET and rlogin proxy settings:

1. Determine whether you wish to allow TELNET and TN3270 connections through the firewall.
2. Determine whether you wish to allow rlogin connections through the firewall.
3. Determine your policies for authentication.

Configuring TELNET and rlogin Proxy Settings

Configure the TELNET or rlogin proxy to enforce your security policies.

To configure TELNET or rlogin proxy settings, provide optional information about time-out values and other configuration settings for the TELNET and rlogin proxy if desired. Refer to the online help for specific information about the available settings.

Note: The settings for the TELNET proxy (*tn-gw*) affect both TELNET and TN3270 access through the firewall. If you need different settings for these two types of access, consider creating a separate configuration for TN3270 access.

Enabling TELNET and rlogin Proxy Services

To enable the TELNET or rlogin proxy service:

1. On the TELNET or rlogin windows, select Enabled.
2. Add the TELNET or rlogin proxy configuration to the service groups you want to use the proxies.
3. Before exiting the Gauntlet Firewall Manager, Save and Apply your changes.
The firewall enables the TELNET or rlogin proxies.

Creating Authentication User Entries

Use the authentication management system to create authentication user entries for any users who authenticate when using TELNET and rlogin services. Refer to Chapter 6, "Users and User Groups," on page 57 for more information.

Verifying Your Setup

Verify your configuration by connecting to an inside host from an outside host. See the next section, "Using Terminal Services" on page 163, for instructions.

Using Terminal Services

This section discusses how to use different configurations of TELNET, rlogin, and TN3270.

TELNET, rlogin, and TN3270 Without Authentication

You can configure the proxies so that they are transparent to your users. Common policies (including the Gauntlet Firewall default policies) configure the proxies so that users working on the trusted networks behind the firewall will not see a change in their daily TELNET and rlogin activities.

For example, Scooter, working on his workstation at Yoyodyne headquarters, wants to continue to connect to his account at Big University using TELNET. Because the Yoyodyne firewall is setup with transparency and does not require authentication, he simply uses TELNET to access the remote host directly. He does not need to explicitly mention the firewall at all.




Figure 18-3 TELNET Connect Window

TELNET and rlogin with Authentication

If you have configured terminal services to require authentication, users will need to follow different procedures to use TELNET or rlogin.

To TELNET using authentication:

1. Use TELNET to connect to the firewall itself.
2. Authenticate to the proxy.
3. Connect to the desired host.
4. Continue as before.

A common security policy for the TELNET proxy is to authenticate all requests from untrusted networks to or through the firewall. The example below shows a sample TELNET session from an untrusted network to a trusted network, using S/Key authentication at the firewall:

```
blaze.clientsite.com-28: telnet fire-out.yoyodyne.com
Trying 204.255.154.100...
Connected to fire-out.yoyodyne.com
Escape character is '^]'.
Username: scooter
Skey Challenge: s/key 651 fi19289 SAFE DUB RISK CUE YARD NIL
Login Accepted
fire-out.yoyodyne.com telnet proxy (Version 4.0a) ready:
tn-gw> c dimension
Trying 10.0.1.120 port 23...
Connected to dimension.yoyodyne.com
BSDI BSD/OS 2.1 (dimension) (ttyp5)
login: scooter
Password: #####
Welcome to dimension.yoyodyne.com
```

In this example, Scooter, working at a client site (blaze.clientsite.com), needs TELNET access to a system behind the firewall (dimension.yoyodyne.com). He first connects to the outside interface of the firewall (fire-out.yoyodyne.com) using TELNET. The TELNET proxy on fire-out prompts him to authenticate. Scooter provides his authentication user ID (scooter). When the proxy prompts, he enters the response to the authentication challenge. The proxy authenticates scooter.

Scooter indicates the host he needs to access (dimension). The TELNET proxy connects Scooter to dimension, and the TELNET daemon running on that system. The TELNET daemon on dimension prompts Scooter for his user name and password on dimension. The TELNET daemon on dimension verifies Scooter's user name and password, then logs him in.

TN3270 With Authentication

If you have configured terminal services to require authentication, users need to follow different procedures to use TN3270.

To use TN3270 with authentication:

1. TN3270 to the firewall itself, disabling true TN3270 support for the initial handshake.
2. Authenticate to the proxy.
3. Connect to the desired 3270 host.
4. Continue as before.

The corporate security policy that requires authentication before using TELNET from untrusted hosts to trusted hosts also applies to using TN3270. Generally, the only difference is in starting the TN3270 client:

```
blaze-55: x3270 -model 2 -efont 3270-12 a: fire-out.yoyodyne.com
```

Managing WWW and Gopher Services

There is a wealth of information stored on systems connected to the Internet. Because of this, your users probably argue that they really need access to the World Wide Web to do their jobs. The graphical interfaces of browsers and Web pages make it easy to access this information, but, along with this ease can come problems. Web services allow for the transfer of a wide variety of file types and for running a number of different programs. This complexity means a greater potential for security problems. Web services are essentially generic file transfer mechanisms and require logging and access control consistent with FTP and terminal services.

The HTTP proxy and authenticating HTTP proxy included with the Gauntlet Firewall securely handles requests for information via hypertext, Gopher, and file transfer. The proxy supports hypertext transfer via the HTTP, SHTTP, and SSL protocols; Gopher transfer via Gopher and Gopher+ protocols; and file transfer via FTP.

This chapter discusses the concepts behind the HTTP proxy and authenticating HTTP proxies and explains how they work; how to configure the proxies for Web, Gopher, and file transfer services; and how to configure these services to run through the firewall. In addition, it includes information on running HTTP servers. The chapter consists of the following sections:

- “Understanding the HTTP and Gopher Proxies” on page 168
- “How the HTTP, Gopher, and SSL Proxies Work” on page 168
- “Accessing HTTP, SSL, and Gopher Proxy Configuration” on page 171
- “Configuring the Firewall for Web and Gopher Services” on page 173
- “Using Web Services” on page 175
- “Using Gopher Services” on page 180
- “Running a Web Server” on page 180

Understanding the HTTP and Gopher Proxies

The HTTP proxy is an application-level proxy that provides configurable access control and logging mechanisms. The HTTP proxy, which runs on the firewall, passes HTTP, SHTTP, SSL, and Gopher requests, and FTP URLs and selectors through the firewall, using rules you supply. You can configure the proxy to allow connections based on:

- source IP address
- source hostname
- destination IP address
- destination hostname

Using these options, you can configure your firewall to allow clients on the inside network to access Web and Gopher sites on the outside network. You can also limit the Web sites your employees can access from systems on the inside network. The proxies log all successful and unsuccessful connection attempts and the amount of data transferred.

The authenticating HTTP proxy works in conjunction with the HTTP proxy to authenticate users. Using the authenticating HTTP proxy, you can configure the proxy to allow connections based on user name. You can require all users to use strong or weak authentication before accessing information on the outside network.

You can configure the HTTP proxy to allow outside hosts to access Web and Gopher servers behind your firewall on inside networks. According to most security policies (including the Gauntlet Firewall default), this access is not a good idea. By design, these services require easy access by people all over the Internet.

How the HTTP, Gopher, and SSL Proxies Work

The firewall runs the HTTP proxy (*http-gw*) as a daemon listening for requests on TCP port 8080. When the firewall receives requests for services (via HTTP, SHTTP, SSL, Gopher, Gopher+, or FTP) on this port, the proxy looks at the request and places it in one of several categories. The proxy then checks the appropriate configuration information and determines whether the initiating host has permission to use the desired service to the desired destination. If the host does not have permission, the proxy logs the connection and displays an error message.

If the host has permission, the HTTP proxy passes the request on to the desired host using the standard port (or the port specified in the request). As the outside host returns data to the firewall, the firewall translates the data into the form the client expects and returns the data to the client. The proxy remains active until either side terminates the connection.

The default configuration for HTTP requests allows all inside hosts to access any Web sites. In this scenario, the Web browser (configured to know about the HTTP proxy) on the inside host passes a request with a URL for a particular Web page to the firewall on port 8080. The firewall calls the HTTP proxy. The proxy examines the request and determines that it is a basic request for HTTP service. The proxy checks the source and destination ports in its configuration information. It then sends the request on to the Web server specified in the URL. When it receives the requested data, it passes the data back to the requesting Web browser.

Changing the HTTP Proxy Port

If you have not configured or cannot configure the Web browser to know about the HTTP proxy, you can change your HTTP configuration so that the firewall calls the HTTP proxy for requests on port 80. To do this, you will need to log in to the firewall as root and use your favorite text editor to edit the file `/usr/local/etc/mgmt/rc/S110http`. Look for the line containing the port number:

```
PORT = 8080
```

and change the port number from 8080 to 80. In order for your changes to take effect, you will need to restart the http proxy with the following commands:

```
/usr/local/etc/mgmt/rc/S110http stop;  
/usr/local/etc/mgmt/rc/S110http start
```

Another option is to run a second or even a third HTTP proxy on an alternate port (such as port 80).

Authenticated HTTP

If you want to authenticate users before allowing them to access information, the firewall runs the authenticating HTTP proxy (*ahhttp-gw*) as a daemon listening for requests on TCP port 8080. When the firewall receives requests for service on this port, it performs the

normal configuration checks to make sure the initiating host has permission to use the desired service to the desired destination.

If the host has permission, the authenticating HTTP proxy prompts the user to authenticate. It verifies the information with the Gauntlet authentication database. If the user provided proper authentication, the authenticating HTTP proxy passes processing over to the HTTP proxy.

The proxy remains active as long as a persistent connection between the source and destination remains. Each time the connection breaks (due to inactivity, pressing the stop button, selecting a link before the initial page finishes loading, or any other reason), the authenticating HTTP proxy reauthenticates you. If you are using reusable passwords, your browser remembers this information and reauthenticates on your behalf. If you are using strong authentication, you must reauthenticate each time the connection breaks.

Gopher and FTP Services

If the request is for Gopher services (from a Web or Gopher client), the firewall calls a second copy of the HTTP proxy, running as *http-gw* on TCP port 70.

If the request is for FTP services (from a Web client), the firewall still calls the HTTP proxy and uses the HTTP rules if you have your FTP proxy set to the HTTP proxy. If you have not set an FTP proxy in your Web browser, the FTP proxy (*ftp-gw*) handles requests for FTP service.

SHTTP and SSL Services

If the request is for some sort of secure HTTP transaction using either the SHTTP protocol (on TCP port 8080) or SSL protocol (on TCP port 443), the proxy performs the appropriate hand-off with the secure server at the other end of the connection.

If you have not configured or can not configure the Web browser to know about the HTTP proxy as the security proxy, the firewall calls the SSL plug proxy for all requests on port 443.

Accessing HTTP, SSL, and Gopher Proxy Configuration

To access the HTTP proxy configuration:

1. From within the Gauntlet Firewall Manager, select Services.
2. Click the HTTP tab.

The HTTP window displays.

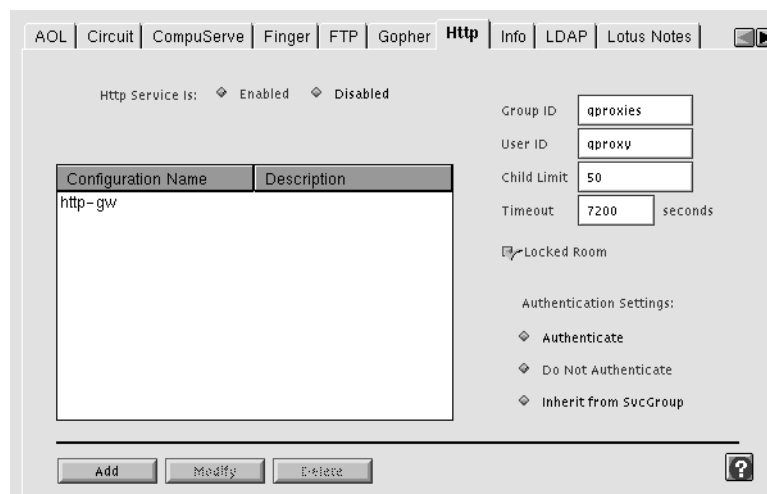


Figure 19-1 HTTP Window

To access the SSL proxy configuration:

1. From within the Gauntlet Firewall Manager, select Services.
2. Click the SSL tab.

The SSL window displays.

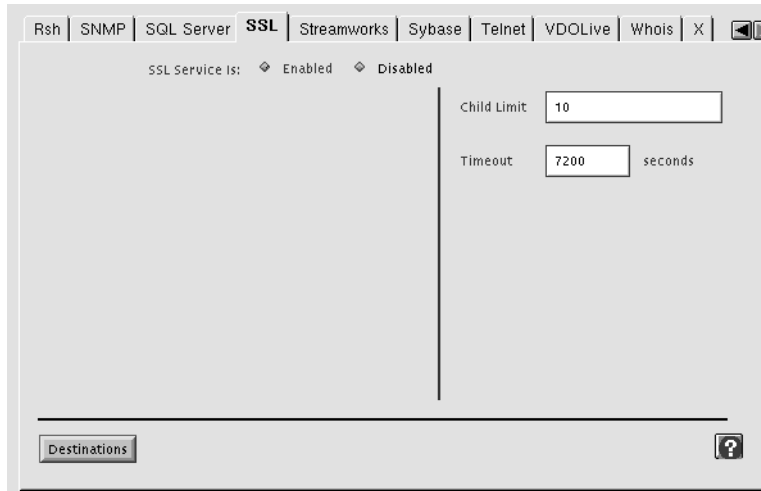


Figure 19-2 SSL Window

To access the Gopher proxy configuration:

1. From within the Gauntlet Firewall Manager, select Services.
2. Click the Gopher tab.

The Gopher window displays.

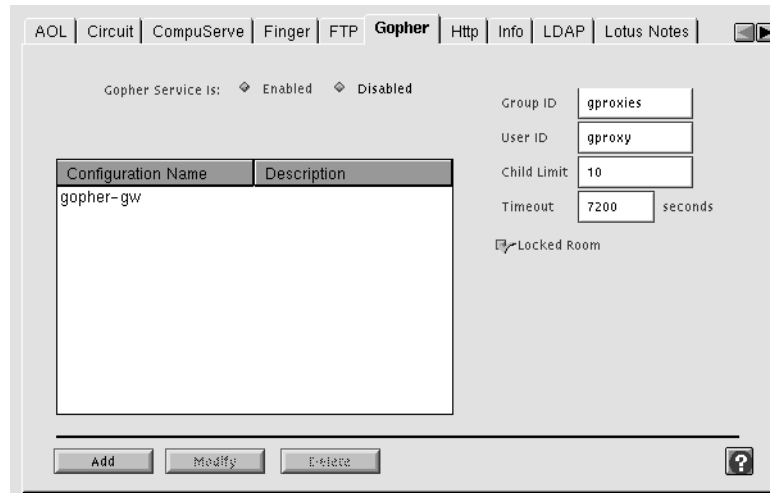


Figure 19-3 Gopher Window

Configuring the Firewall for Web and Gopher Services

Configuring the Gauntlet Firewall involves planning, indicating which daemons the system will run, configuring the proxies to enforce your security policy, turning on the proxies, creating authentication user entries, and rebooting your firewall.

Planning Web and Gopher Service Proxy Settings

When planning proxy settings for Web and Gopher services:

1. Determine which services you will allow.
2. Determine your policies for source and destination sites.
3. Determine whether you wish to require authentication.

Configuring Web and Gopher Service Proxy Settings

Configure the proxies to enforce your security policies.

Configuring HTTP Proxy Settings

To configure HTTP proxy settings:

1. You do not need to configure any settings to use the HTTP proxy service. If you wish, you can use some of the optional settings, or create configuration sets. Refer to the online help for specific information about the available settings. Some commonly used settings include:
 - Enabling Java, JavaScript, or ActiveX blocking, available in the configuration set screen in the Gauntlet Firewall Manager. Refer to the online help for information on this feature.
 - Content Scanning, available in the configuration set screen in the Gauntlet Firewall Manager. Refer to Chapter 30, "Managing Content Scanning," on page 285 for more information.

Configuring Authenticated HTTP Proxy Settings

To configure Authenticated HTTP proxy settings:

1. On the HTTP configuration tab, click Authenticate. If you are placing the HTTP proxy into a service group that requires authentication, click Inherit from SrvGrp.
2. If you wish, you can use some of the optional settings. Refer to the online help for specific information about the available settings.

Configuring SSL Proxy Settings

You do not need to configure any settings to use the SSL proxy service. If you wish, you can use some of the optional settings. Refer to the online help for specific information about the available settings.

Configuring Gopher Proxy Settings

You do not need to configure any settings to use the Gopher proxy service. If you wish, you can use some of the optional settings, or create configuration sets. Refer to the online help for specific information about the available settings.

Enabling Proxy Services

To enable the HTTP, SSL, or Gopher proxy service:

1. On the HTTP, SSL, or Gopher configuration tab, click Enabled.
2. Add the HTTP, SSL, or Gopher proxy configuration to the service groups that you want to use the HTTP, SSL, or Gopher proxy.
3. Before exiting the Gauntlet Firewall Manager, save and apply your changes.

The firewall enables the HTTP, SSL, or Gopher proxy.

Creating User Authentication Entries

If you are using authenticated HTTP, you must create entries in the user authentication system. Refer to Chapter 6, “Users and User Groups,” on page 57 for more information.

Verifying Your Setup

Verify your setup by connecting to some of your favorite Web, Gopher, and FTP sites. Connect to secure Web sites as well.

Using Web Services

Depending upon your configuration, your users may need to modify their activities to access sites using their Web browsers or Gopher tools.

Non-Transparent Access

Under the default configuration the HTTP proxy runs on TCP port 8080, not the standard HTTP port (80). In order to handle Web requests, users must configure their browsers to know about the proxies, as explained under “Using Proxy-Aware Browsers” on page 176. Once you have configured your Web browser, the firewall and the HTTP and SSL proxy are transparent. When using a browser that does not support proxies with a firewall that is not configured for transparency, users need to modify their activities, as explained under “Using Non-Proxy-Aware Browsers” on page 179.

Authenticated HTTP

If you have enabled authentication for the HTTP proxy, users must use a proxy-aware browser. It must support persistent connections if you wish to use strong authentication. Once you have configured their Web browser, they are aware of the proxy because they must authenticate to access outside sites.

Transparent Access

If you wish configure transparent web access on your firewall, you will need to change the default HTTP proxy configuration so that the proxy runs on the standard HTTP port (80). Refer to “Changing the HTTP Proxy Port” on page 169 for instructions on configuring the HTTP proxy to run on port 80. The SSL proxy handles all requests on the standard SSL port (443),so users do not need to modify their activities for secure transactions.

Using Proxy-Aware Browsers

Many Web browsers, such as Netscape Navigator and Microsoft Internet Explorer, are aware of application proxies for different types of Web services. Once you configure these browsers, the browser sends the request to the appropriate proxy.

If you are using the authenticating HTTP proxy, make sure the browser supports proxy authentication and persistent connections.

Configuring Web Browsers

The steps vary depending upon the browser, operating system, and version. Some allow you to indicate the information using a dialog box from a preferences menu, while others require you to edit a configuration file, and others use environment variables.

To configure the browser:

1. Open the settings window for the Web browser you are using.

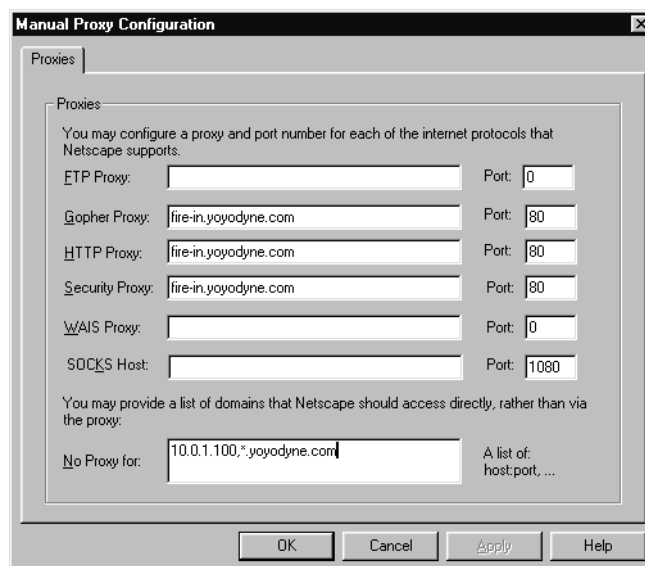


Figure 19-4 Browser Settings Window

2. Provide information about the various proxies.

FTP	Leave name and port for the FTP proxy blank. This allows the FTP proxy to do the processing.
Gopher	Enter name of the inside interface of the firewall and port 8080.
HTTP	Enter name of the inside interface of the firewall and port 8080.
Security	Enter name of the inside interface of the firewall and port 8080.
No Proxy For	Enter names of hosts for which you do not want to access the HTTP proxy in the No Proxy section. These are generally hosts on your trusted networks. These include: <ul style="list-style-type: none"> – inside IP address of your firewall (if you plan to use the graphical user interface to configure your firewall) – hostnames of any internal or corporate HTTP servers – localhost (127.0.0.1)

Note: If you use the IP address instead of the hostname in any of these settings, you must use the IP address of the inside interface of the firewall.

Accessing Web Services without Authentication

Once configured, the proxy is transparent to the user. Users can continue to access the Web as they did before. If you have configured the proxies to block certain types of services (for example, no Gopher services) or to block certain destinations (for example, no educational—.edu—sites), users will see your denial messages.

Accessing Web Services with Authentication

Once configured, users are aware of the proxy. In a particular session, the proxy prompts for authentication the first time you attempt to access a site on the outside network.

To use Web service using weak authentication:

1. Open a URL.
2. Authenticate to the proxy.
3. Continue as before.

If you are using weak authentication, enter your user name and password when your browser prompts you. The proxy remembers this information and reauthenticates you if the connection breaks.

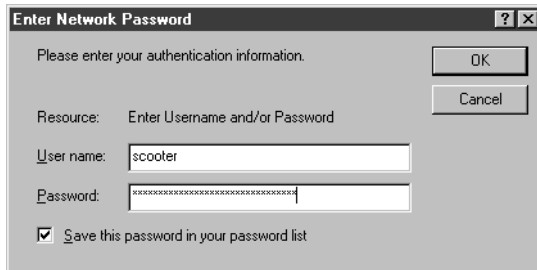


Figure 19-5 Network Password Window

If you are using strong authentication, enter your user name when your browser prompts you. The proxy uses your user name to determine the type of authentication you are using. It prompts you a second time with the appropriate challenge. Enter your user name and your response. Be prepared to reauthenticate each time your connection breaks.




Figure 19-6 User Password Window

Using Non-Proxy-Aware Browsers

Some older Web browsers are not aware of proxies. When using these browsers, you must explicitly send your requests through the firewall.

Configuring Web Browsers

Configuration steps vary depending upon the browser, operating system, and version.

To configure the browser, set up the default home page as the name of the firewall, using the inside address:

```
http://fire-in.yoyodyne.com:8080
```

Accessing Web Services

For regular use of a Web browser, if you cannot create a default home page, prefix each URL you enter with the name of the firewall. For example,

```
http://www.clientsite.com
```

becomes

```
http://fire-in:8080/http://www.clientsite.com
```

where *fire-in* is the hostname of the inside interface of the firewall (fire-in.yoyodyne.com). You must also prepend all saved URLs in bookmarks and hotlists.

Using Gopher Services

Unless you have configured the HTTP proxy to be transparent for Gopher access, users may need to rewrite each Gopher address. If a user has a set of bookmarks for Gopher servers that were created before you installed the firewall, the user may need to modify the bookmark information to include the name of the firewall. For example:

```
name      Big University Gopher Server
host      gopher.bigu.edu
port      70
type      1
path
```

becomes

```
name      Big University Gopher Server
host      fire-in.yoyodyne.com:8080
port      70
path      gopher://gopher.bigu.edu:70/11/
```

Running a Web Server

By its very nature, a Web server requires easy access by the public. If you place the Web server behind the firewall, you are allowing an additional type of access within your security perimeter. If you place the Web server on the firewall itself, you are allowing additional access to your firewall.

The best solution is generally to place your Web server on a system outside the perimeter. Follow good security practices for this system: turn off all other services, create the minimum number of user accounts, use strong authentication, patch your operating system and applications, use checksums to watch for file changes, and back up frequently.

You can also use the Info Server included with the Gauntlet Firewall as a Web server on the firewall itself. Refer to Chapter 29, "Managing Web and Gopher Servers," on page 275 for more information.

Managing X Window Services

The X Window System provides many features and functions that allow systems to share input and output devices. A user running the X Window System on one system can display the results of a graphical program on another system running an X Window client. This flexibility is also the source of a number of well-known security problems.

When you allow access to your display, you are essentially allowing access to your screen, mouse and keyboard. Most sites do not want to provide this sort of free access to their systems, but administrators recognize that these services can be useful. The X11 proxy included with the Gauntlet Firewall allows administrators to selectively allow X11 services through their firewall.

This chapter explores the concepts behind the X11 proxy and explains how it works, how to configure the proxy, and how to use X11 services through the firewall. The chapter consists of the following sections:

- “Understanding the X11 Proxy” on page 181
- “How the X11 Proxy Works” on page 182
- “Accessing X11 Proxy Configuration” on page 183
- “Configuring the Firewall for X11 Services” on page 183
- “Using X11 Services” on page 185

Understanding the X11 Proxy

The X11 proxy is an application-level proxy that provides configurable access control. The proxy, which runs on the firewall, passes X11 display requests through the firewall, using rules you supply. You can configure the proxy to allow display requests based on the display name and user name.

Using these rules, you can configure your firewall to allow only certain systems on the inside network to display information from systems on an outside network. Employees working on the inside network can configure their system to display information from a

program on a client's system on the outside network. Or, an employee working at a client site can display output from a program running on a system inside the company's firewall. Similarly, you can configure your firewall to permit only certain users to use the X11 proxy.

The X11 proxy also requires that users confirm each new request for a connection to their display. Because of the lack of strong authentication systems for X11, this reconfirmation provides an additional opportunity to confirm that you really want to accept the connection.

Because the X11 proxy works in conjunction with the TELNET and Rlogin proxies, you can still configure access based on the source or destination hostname or IP address. The strong authentication feature is also available. The TELNET and Rlogin proxies also log X requests and connections.

How the X11 Proxy Works

Unlike some of the other Gauntlet proxies, the firewall does not start the X11 proxy when it receives display requests. Instead, users must explicitly start the X11 proxy from either the TELNET or Rlogin proxy. The firewall logs and denies all requests for services from hosts other than the firewall on the standard X port (TCP port 6000).

A user connects to the firewall, which runs the TELNET proxy. After checking permissions and authenticating users, the TELNET proxy (*tn-gw*) displays a prompt for the user. At the prompt, the user indicates she or he wishes to allow X displays across the firewall. The TELNET proxy starts the X11 proxy (*x-gw*) on port 6010 or higher. The X11 proxy checks its configuration information and determines whether the initiating user has permission to use X11 services related to the desired display.

If the user has permission, the proxy creates a "virtual display" on the firewall for the requesting client. When the outside X client requests access to the user's display, the virtual display server passes a query window request to the X server on the display system. This X server displays the query window on the real display, prompting the user to confirm the request. After the user confirms the request, the real X server receives the display information from the virtual X server. The proxy remains active until either end closes the connection.

The default service groups do not include the X11 proxy. The firewall itself can run an X server because the X server runs on the standard X port (TCP port 6000) and the X11 proxy runs on other ports (TCP ports 6010 or higher).

Accessing X11 Proxy Configuration

To access the X11 proxy configuration:

1. From within the Gauntlet Firewall Manager, select Services.
2. Click the X tab.

The X window displays.

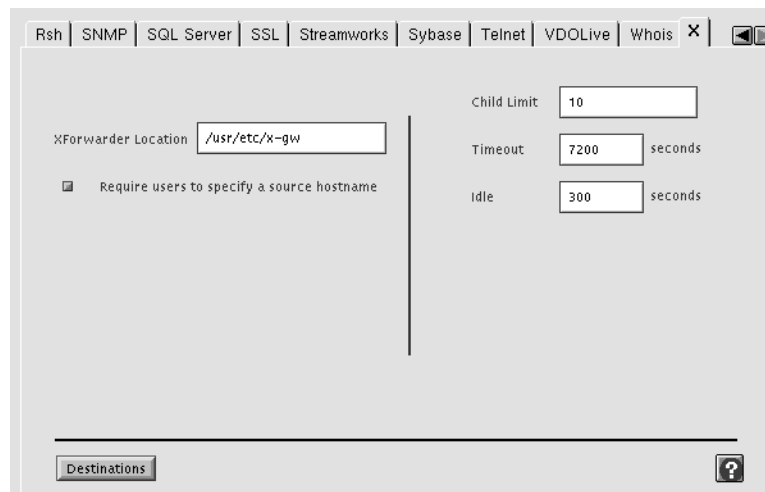


Figure 20-1 X Proxy Configuration Window

Configuring the Firewall for X11 Services

Configuring the Gauntlet Firewall involves planning and configuring the proxy to enforce your security policy.

Planning X11 Proxy Settings

When planning X11 proxy settings:

1. Determine whether you wish to allow X11 display connections through the firewall.
2. Determine which users and which displays can issue and receive display requests.
3. Ensure that your settings for X11 services and TELNET and Rlogin are compatible.

Configuring X11 Proxy Settings

Configure the X11 proxy to enforce your security policies:

1. Click the TELNET or Rlogin tab.
The window for the selected proxy displays.
2. Select the configuration set for which you wish to allow X11 services.
3. Click *Modify*.

The Modify window for the selected proxy displays.

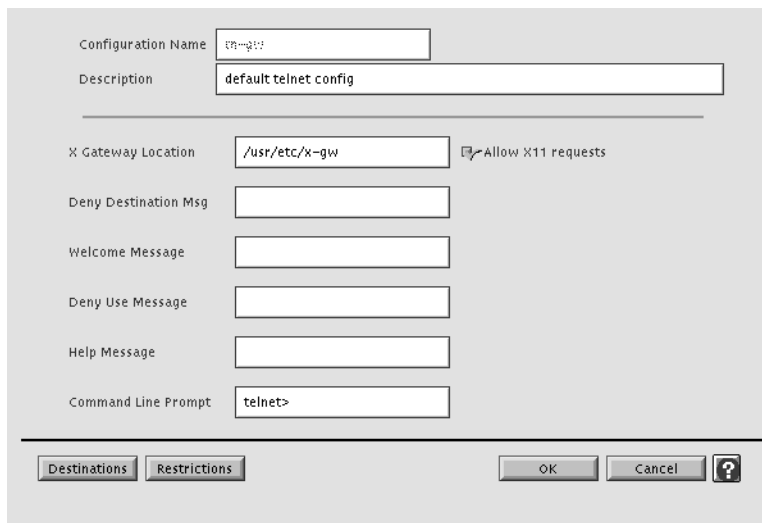


Figure 20-2 Modify TELNET Proxy Configuration Window

4. Click Allow X11 Requests to indicate that you want to allow X11 requests for this configuration set.
5. Click OK.

Enabling X11 Proxy Services

To enable the X11 proxy service:

1. In the X11 window, click Enabled.
2. Add the X11 proxy configuration to the service groups that you want to use the X11 proxy.
3. Make sure you have added the TELNET or Rlogin configuration set to the same service group.
4. Before exiting the Gauntlet Firewall Manager, save and apply your changes.
The firewall enables the X11 proxy.

Verifying Your X11 Proxy Setup

Use TELNET to connect to a system outside the perimeter and display an X11 client on your system inside the perimeter. See "Using X11 Services" for instructions.

Using X11 Services

Users need to follow slightly different procedures to use X11 services through a firewall. The minimal time needed for these additional steps outweighs the time and money to recover after someone hijacks your display. The procedure is the same from either side of the firewall.

To use X11 services:

1. Allow the appropriate interface of the firewall to access your display (remember, it is the firewall you permit to access your display, not the client)
2. TELNET or Rlogin to the firewall.
3. Authenticate to the proxy if necessary.

4. Start the X proxy.
5. TELNET or Rlogin to the desired host.
6. Inform the client of the host and display information that the proxy provides.
7. Start the X client application.
8. Confirm the display request on the real display.

The following example shows a user working on the inside network who needs to display information from a program running on a system on an outside network. The user starts the X11 proxy and establishes a TELNET connection with the outside host:

```
dimension-27: xhost +fire-in
dimension-28: telnet fire-in
Trying 204.255.154.100...
Connected to fire-in.yoyodyne.com
Escape character is '^]'.
Fire-out.yoyodyne.com telnet proxy (Version 4.1) ready:
tn-gw> x
tn-gw> Display port is fire-in.yoyodyne.com:10
tn-gw> c blaze.clientsite.com
Connecting to blaze.clientsite.com .... connected
BSDI BSD/OS 2.1 Kernel #0: Wed Mar 27 20:22:33 MST 1996
login: crawhide
Password: #####
Please wait...checking for disk quotas
You have mail.
blaze.clientsite.com-1:
```

Cindy Rawhide, working at her system (dimension) on the inside network, needs to run an X program on a client system (blaze.clientsite.com) on an outside network, and display the results on her display. She performs these actions:

1. Cindy gives the firewall access to display.
2. She uses TELNETs to connect to the inside interface of the firewall for Yoyodyne (fire-in.yoyodyne.com).

The security policy for her site does not require authentication for inside requests, so the firewall connects her to the TELNET proxy.

3. Cindy indicates that she wants to start an X proxy.

The firewall displays an X status window on Cindy's display, showing the port.

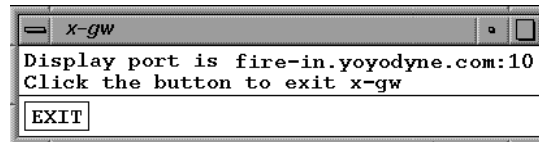


Figure 20-3 X Status Window

4. Cindy uses TELNET to connect to the client system (blaze.clientsite.com).

The TELNET daemon on blaze prompts Cindy for her user name (crawhide) and password on blaze. The TELNET daemon on blaze verifies Cindy's user name and password, and logs her in.

5. Cindy provides the X display information to the client system (blaze) and starts the client application. She uses the hostname of the outside interface of the firewall and the port information that the X proxy provided when she started the X proxy:

```
blaze.clientsite.com-1: setenv DISPLAY fire-out.yoyodyne.com:10
blaze.clientsite.com-2: xclock &
blaze.clientsite.com-3:
```

6. Cindy confirms the display request on her system.




Figure 20-4 X Connection Confirmation Window

7. Cindy views the results on her screen inside the firewall.

Managing Custom Services

Many sites rely on applications such as America Online, CompuServe, Lotus Notes, or custom applications written specifically for their company. Each of these services uses a proprietary protocol. To support these services, you would need a multitude of application-specific proxies. Instead, administrators can use the Plug proxy to tunnel these applications through the firewall because the protocols they use are TCP-based. Other common programs, such as *whois* and *webster*, run over TCP. You can also tunnel these TCP-based services through the firewall with the Plug proxy.

Caution: Allowing proprietary protocols through your firewall can potentially have serious consequences. Because the protocols are proprietary, the firewall and the proxy have no information on the data or requests the applications are sending. It is also unknown how safe the application itself is. Always perform a risk assessment before using the plug proxy for proprietary protocols.

The Plug proxy does not support UDP-based services. UDP is not a connection-oriented protocol. Because there is no connection, there are no sequence numbers. This makes it much easier for someone to create a UDP packet that appears valid but contains fabricated source and destination information.

This chapter explains the concepts behind the Plug proxy and how it works, how to configure the proxy for other services, and how to configure these services to run through the firewall. The chapter consists of the following sections:

- “Understanding the Plug Proxy” on page 190
- “How the Plug Proxy Works” on page 191
- “Accessing Plug Proxy Service Configuration” on page 192
- “Configuring the Firewall for Plug Proxy Services” on page 192
- “Configuring Your Service” on page 195

Understanding the Plug Proxy

The Plug proxy is a TCP gateway that provides configurable access control and logging mechanisms. The Plug proxy, which runs on the firewall, passes application requests through the firewall, using rules you supply. It essentially tunnels information from a port on the firewall to a specific port on another system.

The firewall includes instances of the Plug proxy for:

- America Online
- CompuServe
- *finger*
- LDAP
- Lotus Notes
- NNTP News Client
- NNTP News Server
- *whois*

The Plug proxy is protocol neutral, so you can tunnel a variety of other applications. Weigh the risks carefully for each application. For each version of the Plug proxy, you can configure the proxy to allow connections based on:

- source IP address
- source hostname
- source port
- destination IP address
- destination hostname
- destination port

Using these options, you can configure your firewall to allow your travel department to use their custom reservations system through the firewall. Clients on the inside network can communicate with servers on the outside network.

The proxies log all successful and unsuccessful connection attempts and the amount of data transferred.

These access controls allow you to have more control over the connections to and from your system than you have without a firewall. The logging capabilities are also more extensive. However, you may be allowing proprietary protocols into your network, which can be dangerous.

How the Plug Proxy Works

The firewall runs different instances of the Plug proxy (*plug-gw*) as daemons on different ports for different applications, based on information you supply. This information indicates which services the firewall should run on which ports. For example, you can configure the firewall to run an instance of the Plug proxy on TCP port 5190 to handle America Online requests.

When the Plug proxy receives a request on one of these ports, it checks its configuration information and determines whether the initiating host has permission to initiate this type of request. If the host has permission, the Plug proxy passes the connection on to the specified port on the specified system. This instance of the Plug proxy remains active until either side terminates the connection.

Hosts on both the inside and outside think the firewall is servicing requests. The firewall is simply acting as the tunnel, via the Plug proxy.

Accessing Plug Proxy Service Configuration

To access the Plug proxy configuration:

1. From within the Gauntlet Firewall Manager, select Services.
2. Click the Plug tab.

The Plug window displays.

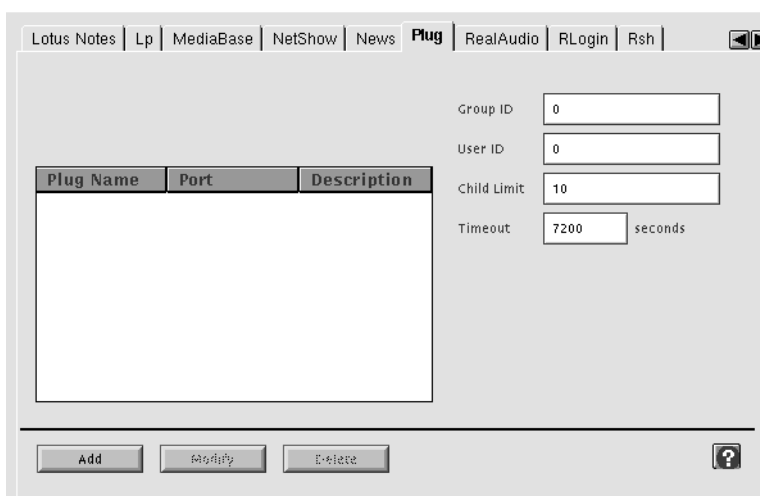


Figure 21-1 Plug Configuration Window

Configuring the Firewall for Plug Proxy Services

Configuring the Gauntlet Firewall involves planning, indicating which daemons the system will run, configuring the proxies to enforce your security policy, and enabling your proxy.

This section uses the Quote of the Day (*qotd*) service as an example. Of course, you must carefully determine if the benefits of something like a quote of the day service outweigh the risks of allowing that type of service within your defense perimeter.

Planning

Caution: Allowing proprietary protocols through your firewall can potentially have serious consequences. Because the protocols are proprietary, the firewall and the proxy have no information on the data or requests the applications are sending. It is also unknown how safe the application itself is. Always perform a risk assessment before using the plug proxy for proprietary protocols.

When planning plug proxy settings:

1. Determine which protocols and which applications you wish to proxy through your firewall.
2. Verify that the protocol is TCP-based.
3. Determine what port these services use. Verify that the service uses the same port for sending and receiving.
4. Determine which external hosts can use these services.
5. Determine which internal hosts can use these services.

Configuring Plug Proxy Settings

As you configure Plug proxy settings, you are actually creating configuration sets. You simply create a different configuration set for each proxy.

To configure a Plug proxy:

1. On the Plug window, click Add.
The Add Plug Services window displays.

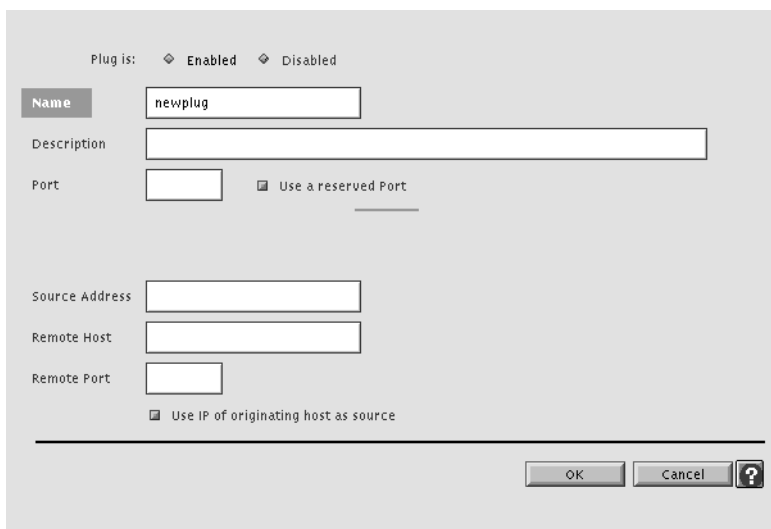


Figure 21-2 Add Plug Services Window

2. Provide information about your service.

Name	Name for the Plug proxy. The firewall uses this name in displaying the list of available services.
Description	Description for this Plug proxy.
Port	TCP port on which this proxy runs. Enter the number of the TCP port.
Use Reserved Port	Specifies that the proxy uses a reserved port number when connecting.
Source Address	Hosts from which connections can originate. Specifies single hosts, entire networks, or subnets. Specify by IP address or host name. The wildcard * is valid.
Remote Host	Hosts to which the Plug proxy connects. Specifies single hosts, entire networks, or subnets. Specify by IP address or hostname. The wildcard * is valid. If you are using transparency, this setting is optional.
Remote Port	Port on which the Plug proxy connects on the remote host. If this option is not specified, the firewall uses the Port value you specified above.

Use IP of Originating Host As Port	Specifies that the Plug proxy uses the IP address of the originating host as the source address of the packet when sending the request on to the destination host. If this option is not specified, the firewall uses its IP address as the source address of the packet, causing all packets to look like they originated on the firewall. You must be using officially registered, routable addresses on your trusted networks in order to use this option.
------------------------------------	--

Enabling Plug Proxy Services

To enable the Plug proxy service:

1. On the Plug window, click Enabled.
2. Add the appropriate configuration set to the service groups that you want to use the this configuration set of the Plug proxy.
3. Before exiting the Gauntlet Firewall Manager, save and apply your changes.

The firewall enables the Plug proxy.

Configuring Your Service

You may need to configure your service and application to connect to the firewall instead of directly to the server. Consult the documentation included with your plugged service for information on possible configurations.

Managing Custom Services With Authentication

Many sites rely on groupware or other interconnected applications, such as accounting packages and database applications. Each of these services uses a proprietary protocol, which would require its own application-specific proxy. The Plug proxy might be an ideal candidate for many of these applications. However, administrators also want to control who can access the service (by user name), which the plug proxy cannot do. Administrators can use the circuit proxy instead to allow certain users to tunnel these proprietary applications through the firewall.

Caution: Allowing proprietary protocols through your firewall can potentially have serious consequences. Because the protocols are proprietary, the firewall and the proxy have no information on the data or requests the applications are sending. It is also unknown how safe the application itself is. Always perform a risk assessment before using the circuit proxy for proprietary protocols.

This chapter discusses the concepts behind the circuit proxy and explains how it works and how to configure and use the circuit proxy. The chapter consists of these sections:

- “Understanding the Circuit Proxy” on page 198
- “How the Circuit Proxy Works” on page 199
- “Accessing Circuit Proxy Configuration” on page 200
- “Configuring the Firewall for Circuit Proxy Services” on page 201
- “Using the Circuit Proxy” on page 204

Understanding the Circuit Proxy

The circuit proxy is an authenticated TCP gateway that provides configurable access control and logging mechanisms. The proxy, which runs on the firewall, authenticates users and passes TCP-based application requests through the firewall, using rules you supply. It essentially tunnels information from a port on the firewall to a specific port on another system, after authenticating the user.

You can configure the circuit proxy to service:

- database applications
- financial applications
- groupware

This is not an exhaustive list. The circuit proxy is protocol neutral, so you can tunnel a variety of other TCP-based applications. Weigh the risks carefully for each application.

You can configure the circuit proxy to allow connections based on:

- user name
- source hostname
- source IP address
- source port
- destination hostname
- destination IP address
- destination port

Using these options, you can configure your firewall to allow certain users to access a database server on a system outside the defense perimeter. Employees working outside the perimeter can access important services inside the perimeter.

The strong authentication features of the circuit proxy require users to authenticate before connecting. The circuit proxy also logs all successful and unsuccessful connection attempts, and the amount of data transferred.

These access controls allow you to have more control over the connections to and from your system than you have without a firewall. The logging capabilities are also more extensive.

How the Circuit Proxy Works

The firewall runs the circuit proxy (*ck-gw*) as a daemon on a user-specified port (generally on a TCP port above 1024). The user initiates the connection using TELNET to the port where the circuit proxy is listening (which is a different port than the port on which the service runs). When the proxy receives a request on this port, it checks its configuration information and determines whether the initiating host has permission to initiate this type of connection. If the host has permission, the circuit proxy authenticates the user with the authentication server specified in the configuration information.

If the authentication is successful, the proxy uses its configuration information to create a menu listing the available services for this user. The user selects the service they want to start. The proxy then waits at the port specified for this service for a connection from the user's system.

The user then starts the client application, which connects to the firewall on the service's port. The proxy accepts the connection and displays a confirmation request in the user's original TELNET window. After the user confirms the request for the connection, the circuit proxy starts a child process to handle the service request. The child process creates a connection from the client to the application server on the other side of the defense perimeter. The proxy then passes requests back and forth between the application client and server. The child process of the circuit proxy remains active until either side terminates the connection. The original TELNET window also remains active until either side terminates the connection.

Accessing Circuit Proxy Configuration

To access the circuit proxy configuration:

1. From within the Gauntlet Firewall Manager, select Services.
2. Click the Circuit tab.

The Circuit window displays.

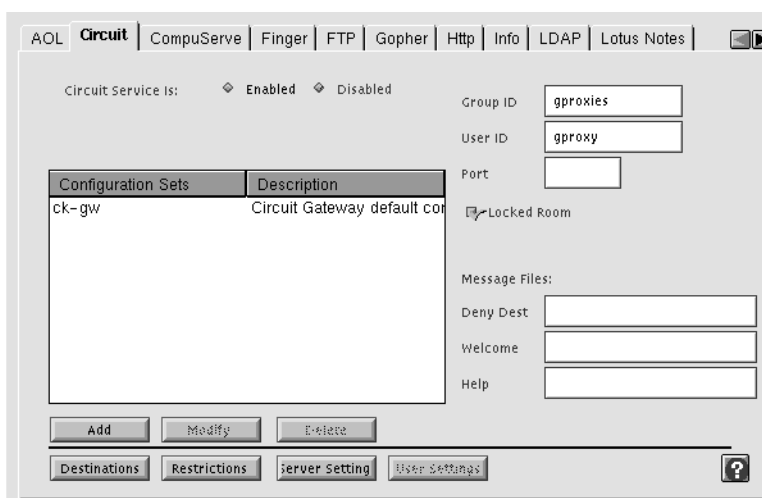


Figure 22-1 Circuit Configuration Window

Configuring the Firewall for Circuit Proxy Services

Configuring the Gauntlet Firewall involves planning, indicating which daemons the system will run, configuring the proxies to enforce your security policy, and enabling the proxy.

Planning

Caution: Note again that allowing proprietary protocols through your firewall can potentially have serious consequences. Because the protocols are proprietary, the firewall and the proxy have no information on the data or requests the applications are sending. It is also unknown how safe the application itself is. Always perform a risk assessment before using the circuit proxy for proprietary protocols.

1. Verify that the protocol is TCP-based by consulting your protocol documentation and trying it with the plug proxy.
2. Verify that the protocol uses one port for all server connections.
3. Verify that the protocol opens only one connection for communicating between the client application and server.
4. Determine which external hosts can use these services.
5. Determine which internal hosts can use these services.

Configuring Circuit Proxy Settings

To configure circuit proxy settings:

1. In the Circuit window, provide information about the port you are using.
Port Enter the port number on which the circuit proxy runs.
2. Click Server Settings.
The Circuit Service Definition window displays.

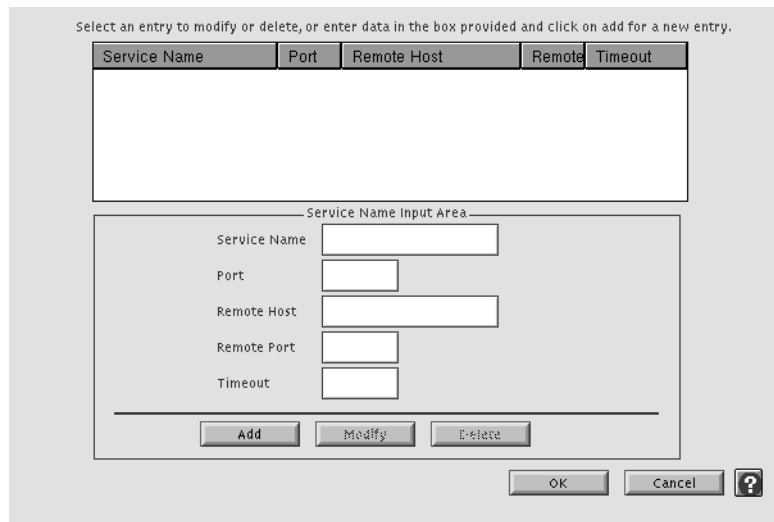


Figure 22-2 Circuit Service Definition Window

3. Provide information about the services you are offering.
Service Name Name of the available service. Used by the proxy to create the menu of available services.
Port Port on the remote host to which the circuit proxy connects. Specify by port number.
Remote Host Name of the remote host to which the circuit proxy connects. Specify an individual system. Use IP addresses or hostnames. This option is required if you are not using transparency.
4. Click *Add* to add this service to the list of available services.

5. Click *OK* to return to the Circuit window.
6. Click *Add* to create a new circuit proxy configuration set.
The Add Circuit Gateway Configuration window displays.

Figure 22-3 Add Circuit Gateway Configuration Window

7. Provide information about your configuration.

Configuration Name	Name for the circuit proxy configuration set.
Description	Description for this circuit proxy configuration set.
8. Click *OK*.

Enabling Circuit Proxy Services

To enable the circuit proxy service:

1. On the Circuit window, click *Enabled*.
2. Add the appropriate configuration set to the service groups that you want to use the this configuration set of the circuit proxy.
3. Before exiting the Gauntlet Firewall Manager, Save and Apply your changes.

The firewall enables the circuit proxy.

Verifying Your Setup

Verify your installation by using your application through the circuit proxy. Watch the logs on the firewall for error messages.

Using the Circuit Proxy

Users follow slightly different procedures to use their application through the circuit proxy:

1. Use TELNET to connect to the circuit proxy.
2. Authenticate to the circuit proxy, if required.
3. Select the desired service.
4. Start your client application.
5. Confirm the client application connection.
6. Use your application.

The example below shows a user working on the trusted network inside the defense perimeter. The company has a policy to authenticate the use of some outside services. He is accessing an Oracle database on a system outside the perimeter at a client site.

First, the user connects to the port on the firewall where the circuit proxy is running. He authenticates using S/Key password:

```
dimension-59: telnet fire-in ck-gw
Trying 10.0.1.100...
Connected to fire-in.yoyodyne.com
Escape character is '^]'.
Username: hikita
Key challenge: s/key 502 fi34762 SILK SCAR DES DON JOEY RUNT
Login Accepted
fire-in.yoyodyne.com ck-gw proxy (Version 4.1) ready:
ck-gw->
```

In this example, Robert Hikita, working at a system (dimension) inside the perimeter needs to access an Oracle database on a system outside the perimeter. He follows these steps:

1. He first uses TELNET to connect to the port (ck-gw) on the firewall for Yoyodyne (fire-in.yoyodyne.com) on which the circuit proxy is running.

The circuit proxy prompts Robert for his authentication user ID, which he provides (hikita).

2. When the proxy responds with a challenge, Robert enters his S/Key response.

The proxy authenticates him using the appropriate authentication server and provides him with a circuit proxy prompt.

3. Robert uses the services command to view a menu of available services. He indicates he wants to connect to the Oracle service (**c oracle**).

```
ck-gw->services
Valid services are:
oracle
finance
reservations
ck-gw-> oracle
waiting for oracle client to be started (type 'q<return>' to
abort)...
```

4. Robert then starts his client application. Because Yoyodyne is using transparency (the default configuration), he indicates that the database server is on the remote host (db.clientsite.com). If Yoyodyne was not using transparency, Robert would tell the client that the database server was the inside address of the firewall (fire-in.yoyodyne.com), allowing the firewall to connect to the database server on his behalf.

5. Robert returns to the original TELNET window in which he connected to the circuit proxy. He notes that the circuit proxy has received a request for service. He confirms the request. He leaves this TELNET window active while he works so that the circuit proxy remains active.

```
waiting for oracle client to be started (type 'q<return>' to abort).
oracle client started
okay to proceed (answer yes only if you started a oracle client)? y
ck-gw->
```

The proxy connects to the remote application server, and begins passing information between the client and server

6. When Robert no longer needs to use the application, he closes the application and the original TELNET window.

Managing MediaBase Services

MediaBase is a collection of multimedia and hypertext that allows users to select and play videos using their Web browser. The Gauntlet Firewall includes a MediaBase proxy that securely handles outside user requests to view video data on a MediaBase server inside the firewall. This proxy also allows users inside the firewall to access MediaBase servers on outside networks.

This chapter explains the concepts behind the MediaBase proxy and how it works.

Note: For additional information on setting up the Gauntlet firewall for MediaBase, see “Configuring a MediaBase Proxy for the Gauntlet Internet Firewall” in the *WebFORCE MediaBase Administrator’s Guide*.

Understanding the MediaBase Proxy

The Gauntlet MediaBase proxy is an application level proxy that provides configurable access control. The proxy, which runs on the firewall, passes MediaBase client and server requests through the firewall, using rules that you supply. You can configure the MediaBase proxy to allow connections based on:

- source host name
- source IP address
- destination host name
- destination IP address

Using these options, you can configure the firewall to allow MediaBase clients on the inside network to access MediaBase servers on the outside network. You can also limit the MediaBase sites your users can access from machines on the inside network.

Used together, these access controls and log files give you much more control over the MediaBase connections to and from your system than you would have without the firewall.

How It Works

The firewall runs the MediaBase proxy (*mbase-gw*) as a daemon listening for requests on a series of ports: ports 6301, 6309, 6310, 6312, and 6313 handle control information; ports 6320 through 6323 and 6340 handle data information. When the firewall receives requests for those ports, the MediaBase proxy checks its configuration information (in the *netperm-table* file) and determines whether the initiating client has permission to use MediaBase. If the client has permission, the proxy logs the transaction and passes the request to the appropriate host.

The *mbase-gw* daemon is always active. This daemon requires that MediaBase players also be configured to use a proxy.

The default policy allows clients inside the network to connect to MediaBase servers; it does not allow outside clients such access, however. Because the firewall runs the MediaBase proxy on all MediaBase ports, all requests from outside clients access the MediaBase proxy rather than the server. This configuration prohibits running a MediaBase server on the firewall itself—there is no way to start a MediaBase server to accept such requests.

Configuring the Firewall to Use the MediaBase Proxy

Configuring the Gauntlet firewall involves planning, indicating which servers may be accessed, and configuring the MediaBase proxy to enforce your policy.

Planning

When planning MediaBase proxy settings:

- Determine which MediaBase servers your users need to access. Obtain hostname or IP address for each server.
- For each user, determine whether you want to limit access to a particular server.
- Determine which external hosts can use these services.
- Determine which internal hosts can use these services.

Configuring MediaBase Proxy Settings

To configure MediaBase proxy settings:

1. From within the Gauntlet Firewall Manager, select Services.
2. Click the MediaBase tab.

The MediaBase window displays.

3. Configure the MediaBase proxy settings:

Source Address	IP addresses of hosts from which connections can originate. Specifies single hosts, entire networks, or subnets. Specify by IP address or hostname. The wildcard * is valid in hostnames.
MediaBase Server	IP addresses of the host to which the proxy connects. Specifies single hosts, entire networks, or subnets. Specify by IP address or hostname. The wildcard * is valid.

Enabling MediaBase Services

To enable the MediaBase proxy:

1. In the MediaBase window, make sure MediaBase service is enabled.
2. Add the MediaBase configuration to the service groups you want to use the MediaBase proxy.
3. Before exiting the Gauntlet Firewall Manager, save and apply your changes.

The firewall enables the MediaBase proxy.

Using the MediaBase Proxy

Users must set up the client-side configuration files to enable the MediaBase client to communicate with a MediaBase firewall proxy.

Verifying Your Setup

Verify your installation by using your MediaBase Player to connect to MediaBase servers on the outside network.

PART FOUR

Managing the Firewall Environment

Chapter 24

Mail Services

Chapter 25

Managing Packet Screening

Mail Services

For many people, electronic mail is an integral tool for conducting business. Exchanging electronic mail (e-mail) is often the reason sites decide they need to connect to the Internet. Such connections are not without risks, however.

The main protocol for transferring e-mail around the Internet is the Simple Mail Transport Protocol (SMTP). The transfer requests are handled by a message transfer agent, such as the *sendmail* program used on many IRIX systems. The *sendmail* program is large and requires many privileges. The Gauntlet design of reductionism does not allow the use of *sendmail* as a critical security component of the Gauntlet Firewall. The Gauntlet Firewall includes a two-part proxy that securely handles the transfer of SMTP mail between the inside and outside networks.

Employees and companies are expanding the places in which and the types of systems on which they need to read their electronic mail. For a variety of reasons, it is not convenient to run a full mail transfer system using SMTP on these systems. The Post Office Protocol Version 3 (POP3) is one of the protocols that allow a workstation to access a mail server. The POP3 proxy included with the Gauntlet Firewall allows administrators to selectively allow outside hosts to exchange mail with a POP3 mail server through the firewall. The POP3 server should use APOP for authenticating the user.

This chapter discusses the concepts behind the SMTP and POP3 proxies and explains how they work, how to configure the proxy for mail transfer, and how to configure these services to run through the firewall. The chapter consists of the following sections:

- “Understanding the SMTP Proxy” on page 214
- “Accessing SMTP Proxy Configuration” on page 215
- “Configuring the Firewall for SMTP Services” on page 216
- “Configuring Other Settings” on page 218
- “Using Mail” on page 219
- “Understanding the POP3 Proxy” on page 219
- “Accessing POP3 Proxy Configuration” on page 220

- “Configuring the Firewall for POP3 Services” on page 221
- “Configuring Your Internal POP3 Mail Server” on page 223
- “Using POP3 to Exchange Mail” on page 223

Understanding the SMTP Proxy

The proxy for SMTP is actually two different processes: a client (*smap*) and daemon (*smapd*). Together, they provide configurable access control and logging mechanisms. The processes, which run on the firewall, transfer mail between internal and external mail servers, based on rules you supply. You can also configure the message transfer agent that the firewall uses to deliver the messages to other hosts.

The proxies also prevent versions of *sendmail* on the inside network from talking with versions of *sendmail* on the outside network. The proxies log all successful and unsuccessful mail connections, and the number of bytes transferred.

How the SMTP Proxy Works

The firewall runs the client proxy (*smap*) as a daemon listening for requests on the standard SMTP port (TCP port 25). When the firewall receives requests for SMTP services on this port, the SMAP client collects the mail from the sender, logs the message, and places the mail in a temporary directory. Periodically, based on a configurable value (by default, every 60 seconds), the daemon (*smapd*) wakes up and checks if there is any new mail. The *smapd* daemon checks the headers of the mail for formatting problems. It then calls the configured message transfer agent (usually *sendmail* in delivery mode) for final delivery.

Both the *smap* client and the *smapd* daemon run using a user ID you specify, such as *uucp*. Rather than running as a root process as *sendmail* often does, the *smap* and *smapd* processes run with as few or as many privileges as you assign. In addition, both programs change their root directory to the transfer directory you specify.

A common configuration is to have one mail hub for the inside network. In this scenario, outside networks know (via DNS) that they should send all mail for the domains (*yoyodyne.com*) on the inside networks to the firewall (*fire-out.yoyodyne.com*) itself for processing. An outside host informs the firewall that it has mail. The firewall calls the

smap client to handle the request. The *smap* client collects the mail from the outside host and writes it to a directory (*/var/spool/smap*) on the firewall.

At some interval (configurable by the system administrator), the *smapd* daemon awakens and looks for new mail on the firewall. It parses the mail headers and calls *sendmail* to deliver the messages to the SMTP mail hub on the trusted network. *sendmail* checks its configuration information, which indicates that it should deliver all internal mail to the internal mail hub (mail.yoyodyne.com). The *sendmail* program on the firewall transfers the mail to the SMTP service on the mail hub.

Accessing SMTP Proxy Configuration

To access the SMTP proxy configuration:

1. From within the Gauntlet Firewall Manager, select Environment.
2. Click the Mail tab.

The Mail window displays.

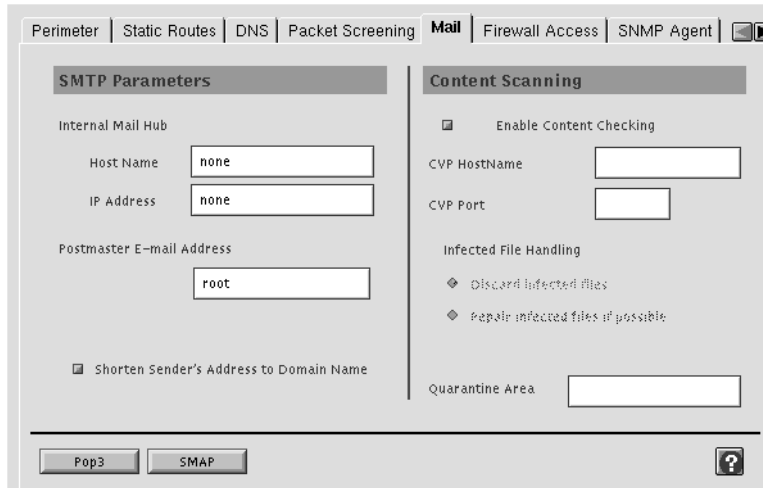


Figure 24-1 Mail Window

Configuring the Firewall for SMTP Services

Configuring the Gauntlet Firewall involves planning, configuring the firewall, configuring the proxies to enforce your security policy, and enabling the proxy.

Planning SMTP Proxy Services

When planning SMTP proxy services:

1. Understand your existing mail configuration: hosts, hubs, and so on.
2. Plan early to make your DNS changes for mail records. This may require contacting an outside organization providing DNS service, such as an internet service provider (ISP).

Note: This step is very important.

Configuring SMTP Proxy Services

If you wish to allow SMTP traffic through the firewall, you must configure the firewall. The administrative interface writes this information to the *sendmail.cf* file, which contains information the *sendmail* program uses to deliver mail to internal systems.

To configure SMTP proxy settings:

1. In the Mail window, provide information about your SMTP mail settings.

Hostname	Name of the host to which the firewall forwards all e-mail for your domain. Enter a fully qualified host name to forward e-mail to a host inside your firewall for processing. Enter None if you want the firewall to process and deliver the mail itself. Make sure the <i>/etc/aliases</i> file on the firewall contains information about internal systems.
IP Address	IP address of the internal mail hub.
Postmaster E-Mail Address	Address to which users and systems can send e-mail about mail. Name of a user, group, or alias who regularly reads this e-mail.
Shorten Sender's Address To Domain Name	Specifies whether the firewall rewrites outgoing e-mail addresses to remove system names from certain lines (such as From: and cc:) in the mail headers. Select enabled if you want your firewall to modify e-mail addresses in outgoing mail from penny@dimension.yoyodyne.com to penny@yoyodyne.com.

2. If desired, provide optional information about time-out values and other configuration settings for the SMTP proxy using the *SMAP* button at the bottom of the Mail window. Refer to the online help for specific information about the available settings. Refer to Chapter 30, "Managing Content Scanning," for information on configuring content scanning.

Enabling SMTP Proxy Services

To enable the SMTP proxy service:

1. In the Mail window, click *SMAP*.
2. In the *SMAP* window, click Enable *SMAP* and *SMAPD*.
3. Click OK.
4. Before exiting the Gauntlet Firewall Manager, save and apply your changes.

The firewall enables the SMTP proxy the next time you reboot.

Configuring Other Settings

Configuring settings consists of these activities:

- “Advertising the Firewall as a Mail Exchanger” on page 218
- “Configuring Your Internal Mail Hub” on page 218
- “Verifying Your Setup” on page 218

Advertising the Firewall as a Mail Exchanger

Advertise the firewall as the mail exchange host for your domain. For more specific information, see *IRIX Admin: Networking and Mail*.

Configuring Your Internal Mail Hub

As long as you are using transparency to pass all packets for outside networks to the firewall (the default Gauntlet configuration), you do not need to configure your internal mail hub or mail agents. Because of the transparency, these systems forward all requests to the firewall.

If you are not using transparency, configure your internal mail hub to use the firewall as a mail forwarder, and direct clients to the internal mail hub. If you do not have an internal mail hub, configure the clients to use the firewall directly as a mail forwarder.

Verifying Your Setup

To verify your mail setup:

1. Verify your configuration by sending mail from an inside host to an outside host. Watch the logs on the firewall for error messages.
2. Run *mail* in verbose mode and send mail to a bouncing service, which will automatically generate a reply:

```
dimension-23: mail -v bouncer@bbnplanet.com
Subject: Test After Configuring Mail and the Gauntlet Firewall
This is a test.
```

.

The verbose mode ensures that you see the details of the delivery. The bouncer service sends you a return message shortly.

3. If you need to test header rewriting or other custom configurations, start *sendmail* in debug mode using the **-bt** option.

Using Mail

The firewall and the *smap* and *smapd* proxies for SMTP traffic are transparent to the user.

Understanding the POP3 Proxy

The POP3 proxy is an application-level gateway that provides configurable access control, authentication, and logging mechanisms. The POP3 proxy, which runs on the firewall, transfers mail between external workstations and internal mail servers, based on rules you supply:

- source IP address
- source hostname
- destination IP address
- destination hostname
- user name

Using these options, you can configure your firewall to allow specific hosts on outside networks to exchange mail with an internal mail server via POP3. An employee working with a laptop PC running Windows needs to read mail while traveling. They can use a mail user agent (such as Eudora Pro) on the laptop to collect their mail from the mail server inside the perimeter. The proxy uses the APOP command (part of the POP3 protocol) for strong authentication. The proxy logs all successful and unsuccessful mail connections, and the number of bytes transferred.

You can configure the POP3 proxy to allow inside workstations to exchange mail with POP3 servers outside the perimeter. However, according to most security policies (including the Gauntlet Firewall default), this is not considered a good idea. The POP3 protocol assumes that the SMTP proxy has already checked the formatting in the headers of incoming mail messages. In addition, allowing POP3 clients to communicate with outside mail servers adds another level of complexity. It bypasses the central control

center of the inside mail hub, which rewrites addresses and enforces other company policies. Your mail server should be behind the firewall on the inside network. All POP3 clients on the inside network can collect their mail from this mail server.

How the POP3 Proxy Works

The firewall runs the POP3 proxy (*pop3-gw*) as a daemon listening for requests on the standard POP3 port (TCP port 110). When the firewall receives requests for POP3 services on this port, the proxy checks its configuration information and determines whether the initiating host has permission to use POP3 services. If the host does not have permission, the proxy logs the connection and displays an error message.

If the host has permission, the POP3 proxy authenticates the user using APOP and logs the connection. The proxy then passes the message on to the POP3 server on the internal mail hub, and authenticates on behalf of the user using APOP. The proxy remains active until either side terminates the connection.

The default configuration allows outside hosts to connect to an internal mail server to collect mail. The firewall itself cannot run a POP3 server, because the POP3 proxy is running on the standard POP3 port.

Accessing POP3 Proxy Configuration

To access the POP3 proxy configuration:

1. From within the Gauntlet Firewall Manager, select Environment.
2. Click the Mail tab.
3. Click POP3.

The POP3 Mail and Proxy Configuration window displays.

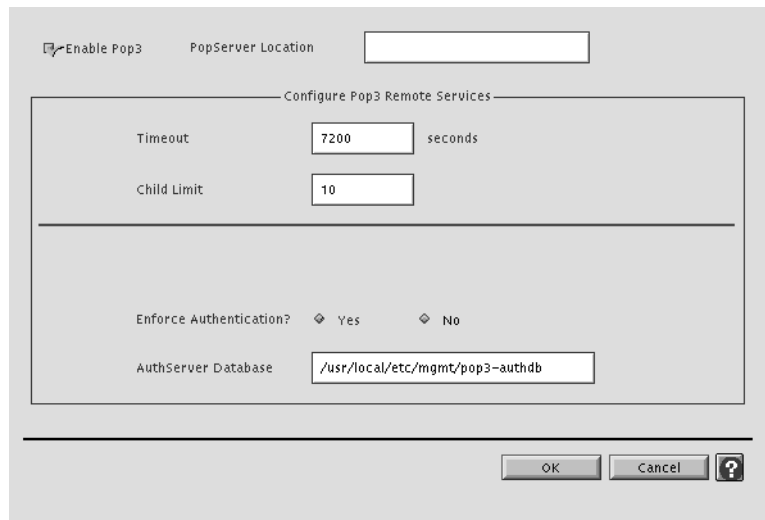


Figure 24-2 POP3 Mail and Proxy Configuration Window

Configuring the Firewall for POP3 Services

Configuring the Gauntlet Firewall involves planning, configuring the proxy to enforce your security policy, creating APOP accounts for users who will need to authenticate, and enabling the proxy.

Planning

When planning POP3 proxy settings, determine your policies for:

- source and destination addresses
- user access to POP3

Configuring POP3 Proxy Settings

To configure POP3 proxy settings, provide information about your POP3 mail settings in the POP3 Mail and Proxy Configuration window.

POP3 Server Location IP address of the system running your POP3 server.

Creating User Authentication Entries

If you are using APOP authentication, which is the recommended and default configuration, you must create a user authentication entry for each user who will access the POP3 proxy.

To create user authentication entries:

1. Create user authentication entries for each user. You can use the same user ID and have different passwords for different types of authentication. For example, you can use a strong authentication mechanism for other access, and APOP authentication for POP3 access.
2. Provide information about the APOP password. Check the Set Pop3 Password box to indicate that you want to create an APOP password. Enter the new APOP password, then reenter the same new password in the Verify field.
3. Make a note of the POP3 password as you need to enter this value on the POP3 server and provide it to the user.

Enabling POP3 Proxy Services

To enable the POP3 proxy service:

1. In the Mail window, click Pop3.
2. In the POP3 window, click Enable Pop3.
3. Click OK.
4. Add the POP3 configuration to the service groups you want to use the POP3 proxy.
5. Before exiting the Gauntlet Firewall Manager, save and apply your changes.

The firewall enables the POP3 proxy the next time you reboot.

Configuring Your Internal POP3 Mail Server

You must configure your internal POP3 mail server so that it accepts requests for service from the firewall.

To configure your internal POP3 mail server:

1. Configure your POP3 mail server to accept POP3 requests from the firewall. If you need to specify an IP address, remember to use the internal IP address for the firewall.
2. Ensure that the POP3 mail server is using the POP3 port (110).
3. Configure your POP3 mail server to support APOP.

Configure the APOP password for each user. Use the same APOP password that you specified when creating user authentication entries on the firewall.

Using POP3 to Exchange Mail

Because the POP3 proxy requires authentication, users need to follow different procedures to use POP3 services.

To retrieve electronic mail using POP3 with authentication:

1. Configure the mail user agent and set the name of the POP3 server to the firewall.
2. Retrieve mail, causing the user agent to connect to the firewall.
3. Authenticate to the proxy by supplying your APOP password.
4. Continue as though the firewall were not there.

Note: The order of these steps may differ for different user agents.

The example below shows a user working on an outside network who needs to retrieve mail from the mail server on the inside network.

First, the user configures the mail reader to get mail via POP3 from the firewall. The following figure shows the configuration window for Eudora Pro for Windows, a popular mail application.

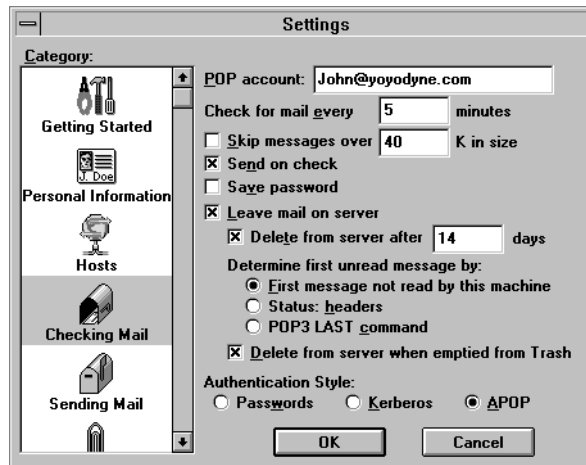


Figure 24-3 Eudora Pro Configuration for APOP

John, working on his laptop (cavalier.yoyodyne.com) at home, configures his mail reader to connect to the firewall (fire-out.yoyodyne.com) to get his mail.

Next, John retrieves his mail. As part of the connection, the proxy requests authentication information from the user agent, which prompts the user. After authenticating, the proxy transfers the request to the internal POP3 mail server (mail.yoyodyne.com), authenticates using the value stored on the firewall, and retrieves his mail.

Using the POP3 Proxy with Multiple POP3 Servers

You have multiple POP3 servers within your organization, all behind the same firewall. If you want to direct different users to each POP3 server, use the POP3 proxy in this configuration. It is not necessary to modify the configuration on the firewall

To access different POP3 servers, enter your user name when prompted to authenticate. Use this format:

```
user%pop3server@firewall
```

- *user* specifies the user name in the firewall's authentication system.
- *pop3server* specifies the name of the POP3 server.
- *firewall* specifies the name of the firewall.

Managing Packet Screening

Some sites make use of network monitoring tools to manage their networks. These tools often use protocols such as UDP and ICMP. You cannot use the plug or authenticated circuit proxies to allow these protocols through the Gauntlet firewall because these proxies support only TCP-based traffic. You may also want to allow a TCP-based application through the Gauntlet Firewall. However, the application makes use of multiple ports for communicating and is unsuitable for use with either the plug or authenticated circuit proxies.

The Gauntlet Firewall includes a packet screening facility. Use the packet screening facility to allow various types of traffic, such as UDP and ICMP packets, through the Gauntlet Firewall.

Caution: Allowing packets through the Gauntlet Firewall is a risk. Traffic that goes through the Gauntlet Firewall via the packet screen does not also pass through the proxies. Therefore, the Gauntlet Firewall cannot provide logging or authentication of this traffic as it does on traffic that goes through the proxies. Use of packet screening is optional. Configure packet screening rules only after performing a careful risk analysis. Only advanced firewall systems administrators who can judge the risks and consequences of rules they create should configure packet screening rules. A mistake in configuring packet screening can severely compromise the security of your protected network!

The following sections explain how the packet screening facility works, how packet screening rules work, and how to configure packet screening rules:

- “Understanding Packet Screening” on page 226
- “How Packet Screening Works” on page 227
- “How Packet Screening Rules Work” on page 227
- “Accessing Packet Screening Configuration” on page 232
- “Adding Packet Screening Rules” on page 233
- “Modifying Packet Screening Rules” on page 235

- “Deleting Packet Screening Rules” on page 235
- “Changing Order of Precedence” on page 236
- “Verifying Your Configuration” on page 236

Understanding Packet Screening

The packet screening facility, implemented as part of the IRIX kernel, reads rules you create (and default Gauntlet Firewall rules) and processes packets accordingly. Use the packet screening facility to discard packets, pass packets on to the proxies, or bypass the proxies entirely. You can configure the Gauntlet firewall to screen packets based on these options:

- source IP address
- destination IP address
- interface
- protocol
- source port
- destination port

Using these options, you can configure the Gauntlet Firewall to allow a UDP-based service through on a particular port. For example, a host on one side of your Gauntlet Firewall could use *ntpd*, which uses a UDP-based protocol, to set the time on a host on your service network.

The Gauntlet Firewall default packet screening rules implement IP-spoofing checks. These rules deny any packets that appear on the outside interface of the firewall with source IP addresses that match IP addresses of hosts on the inside, trusted network.

The Gauntlet Firewall logs packets denied by the packet screening facility. These logging capabilities are more extensive than those provided without the Gauntlet firewall in place. However, the packet screening facility does *not* normally log packets permitted through the screen. You do not receive as much logging information when using packet screening as when using an application proxy.

How Packet Screening Works

The packet screening facility is part of the IRIX kernel. When the boot process for the firewall begins, the packet screening rules are empty. This causes the firewall to refuse all traffic. As part of the boot process, the packet screening facility loads its screening rules from configuration files. Each time the firewall receives a packet, the kernel compares header information in the packet to appropriate screening rules and determines what to do with the packet.

If header information does not match any of the rules, the firewall drops the packet. This is in keeping with the “that which is not permitted is denied” philosophy of the Gauntlet Firewall.

How Packet Screening Rules Work

The packet screening facility understands two different sets of rules: local and forward. Together, these rules cover any packets the Gauntlet Firewall receives.

Local rules apply to any packet destined for the Gauntlet Firewall itself. For example, if the Gauntlet Firewall receives a *ping* request from an untrusted host for the outside interface of the firewall (*ping* fire-out.yoyodyne.com), the packet screening facility uses local rules to determine what to do with *ping* packets.

Forward rules apply to any packets destined for hosts other than the Gauntlet Firewall. For example, if the Gauntlet Firewall receives a TELNET request from a trusted host to connect with an untrusted host (*telnet* dialin.bigu.edu), the packet screening facility uses forward rules to determine what to do with TELNET packets.

Packet Screening Action Rules

Each packet screening rule specifies one of three actions: deny, permit, or absorb.

Deny Rules

A deny rule instructs the Gauntlet Firewall to drop the packet and log information. The kernel does not notify the packet sender that it dropped the packet. The default Gauntlet Firewall forward and local screening rules include rules to detect IP spoofing. These rules deny packets that appear on the outside interfaces of the Gauntlet Firewall pretending to

come from inside by presenting source IP addresses that match the IP addresses of hosts on the inside, trusted network.

Permit Rules

A permit rule instructs the Gauntlet Firewall to process the packet. If the packet is destined for the Gauntlet Firewall itself, the firewall accepts the packet. The kernel passes handling of the packet to the appropriate program, such as one of the proxies. If the packet is destined for another host, the firewall routes the packet to the destination host, bypassing the proxies.

Local screening rules include a rule that allows the Gauntlet Firewall to accept packets destined for the IP address of each interface. For example, the firewall is listed in DNS as the mail exchanger for your domain. Local screening rules on the firewall include rules that permit the Gauntlet Firewall to accept these packets.

If your security policy permits, you could create a forward screening rule that permits the firewall to forward ICMP packets. This screening rule allows you to use the *ping* and *traceroute* networking tools to monitor hosts on your internal network on either side of the firewall.

Absorb Rules

An absorb rule instructs the Gauntlet Firewall to process the packet as if the packet were destined for the firewall itself. Absorb rules are generally included as part of the forward rules. The local rules affect only those packets that are destined for the Gauntlet Firewall itself. Including an absorb rule in the local rules is redundant.

The Gauntlet Firewall includes default absorb rules that implement transparency from the inside to the outside. These absorb rules tell the packet screening facility to absorb the packet and process the packet locally, even though the destination IP address is some other system.

For example, a user on a trusted host wants to access her account at Big University directly (*telnet dialin.bigu.edu*). The default route for the trusted network goes through the Gauntlet Firewall. If there is no absorb rule in place, the packet screen uses the permit and deny rules in the forward rules to determine what to do with the packet. The kernel never passes the packet to the proxies.

And, if the forward rules contain an absorb rule, the packet screen accepts the packet for delivery locally. The packet screening facility processes the packet according to the local rules, and passes the packet onto the appropriate proxy service on the firewall itself. The TELNET proxy uses its own rules to determine if the trusted host can TELNET to the untrusted host.

Packet Screening Field Rules

The packet screening rules allow you to permit or deny network traffic, based on several values in packet headers:

- source IP addresses
- destination IP addresses
- network interface
- protocol
- source ports
- destination ports

Source IP Address and Destination IP Address Rules

Rules indicate the source IP address and mask and the destination IP addresses and mask. The address mask indicates how much of the IP address is significant. For example, an IP address of 10.12.1.0 with a netmask of 255.255.255.0 matches all IP addresses beginning with 10.12.1. The IP address 0.0.0.0:0.0.0.0 matches every packet's IP address.

The packet screening utility applies the address mask in the screening rule to the IP address in the packet in a bitwise AND operation. If the result matches the IP address in the screening rule, then the packet screening facility considers the rule a match.

While the packet screening facility can detect attempts to spoof trusted IP addresses, the packet screening facility cannot detect one untrusted network host masquerading as another. Be careful when configuring rules that rely on source IP addresses on an untrusted network.

Interface Rules

Rules indicate the interface on which the packet arrives at the Gauntlet Firewall. These match the interface names on your firewall, for example ec0, ec1, and ec2. When using the packet screening options in the administration tools, you can use symbolic names for the interfaces: inside, outside, third (for your service net), and any (for any network).

The packet screening facility applies the rule as packets arrive at the firewall. This ensures that the packets adhere to the defined network configuration, including the name of the network interface that received the packet.

Protocol Rules

Rules indicate the type of protocol the IP packet uses. Valid values are in RFC 1700. Common values include 6 (TCP) and 17 (UDP). Use the asterisk wildcard (*) to match any protocol.

You can also specify subtypes of the ICMP protocol. Valid subtypes are:

- ECHO
- ECHOREPLY
- IREQ
- IREQREPLY
- MASKREPLY
- MASKREQ
- PARAMPROB
- REDIRECT
- ROUTERADVERT
- ROUTERSOLICIT
- SOURCEQUENCH
- TIMXCEED
- TSTAMP
- TSTAMPREPLY
- UNREACH

For example, Yoyodyne wants to allow only ECHO and ECHOREPLY packets. When creating the packet filtering rule, they use the following syntax as the protocol:

```
i.cmp : ECHO | ECHOREPLY
```

On another firewall, Yoyodyne wants to allow all ICMP packets *except* ECHO and ECHOREPLY. When creating the packet filtering rule, they use the following syntax as the protocol:

```
i.cmp : !ECHO & !ECHOREPLY
```

Source Port and Destination Port Rules

Rules indicate the source port and destination ports. These rules match ports on which the particular service communicates. Valid values are any number between 1 and 65,534. Use the asterisk wildcard (*) to match any port.

If the protocol is other than TCP or UDP, the packet screening facility considers the port to be the wildcard *, matching any port.

Order of Precedence of Packet Screening Rules

As the packet screening facility reads through the screening rules, the sequence of the rules is significant. The packet screening facility reads rules from top to bottom, stopping when it finds a match. When the interface, source IP address, source port, destination IP address, destination port, and protocol match a rule, the packet screen returns the appropriate permission or denial.

The packet screening rules are loaded into the kernel in this sequence:

1. User-defined rules specified manually in the netperm-table using the **authenIP** keyword.
2. IP spoofing rules and other user-defined rules specified manually in the netperm table using the **netconfig** keyword.
3. User-defined rules from the packet screening configuration option in the administration tools.
4. VPN rules that implement virtual private networks, trusted networks, and passthrough networks.
5. User-defined rules specified manually in the netperm table.
6. Transparency rules.

Accessing Packet Screening Configuration

To access packet screening configuration:

1. From within the Gauntlet Firewall Manager, select Environment.
2. Click the Packet Screening tab.

The Packet Screening window displays.

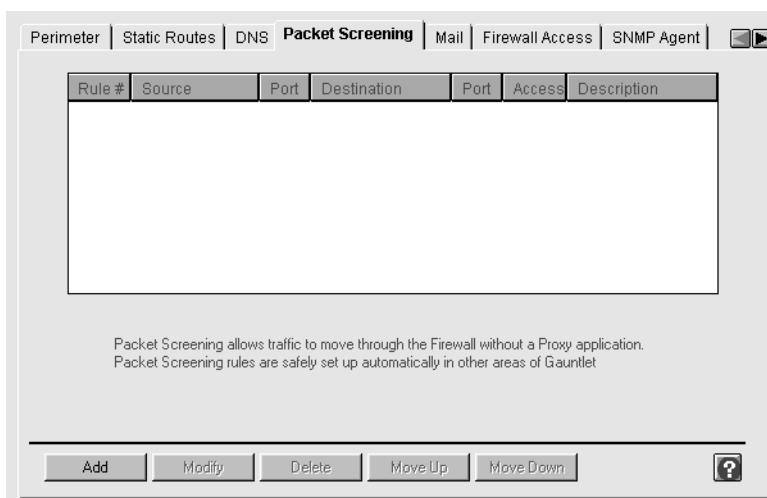


Figure 25-1 Packet Screening Window

Adding Packet Screening Rules

Creating packet screening rules involves planning, creating actual rules, and loading screening rules into the kernel.

Planning Packet Screening Rules

Warning: Allowing packets through the Gauntlet Firewall is a risk. Traffic that goes through the Gauntlet Firewall via the packet screen does not also pass through the proxies. Therefore, the Gauntlet Firewall cannot provide logging or authentication of this traffic as it does on traffic that goes through the proxies. Use of packet screening is optional. Configure packet screening rules only after performing a careful risk analysis. Only advanced firewall systems administrators who can judge the risks and consequences of rules they create should configure packet screening rules. A mistake in configuring packet screening can severely compromise the security of your protected network!

When planning packet screening rules:

1. Determine the type of traffic you want to pass through the Gauntlet Firewall.
2. Perform a careful risk analysis and ensure that this traffic is an acceptable risk.
3. If you plan to create forward rules that allow access to your inside network, make sure that you are using registered, routable IP addresses for your inside networks. You cannot use one of the reserved (RFC 1918) addresses for your networks.

Creating Packet Screening Rules

The packet screening rules editor in the Gauntlet Firewall Manager is appropriate for most configurations and is recommended for creating screening rules. You can also create packet screening rules manually. Use this manual method only if you have an unusual configuration, such as four interfaces. Refer to the *Gauntlet Netperm Reference Guide* for more information on manually adding packet screening rules.

To create packet screening rules:

1. In the Packet Screening window, click *Add*.
The Add Packet Screening Rule window displays.

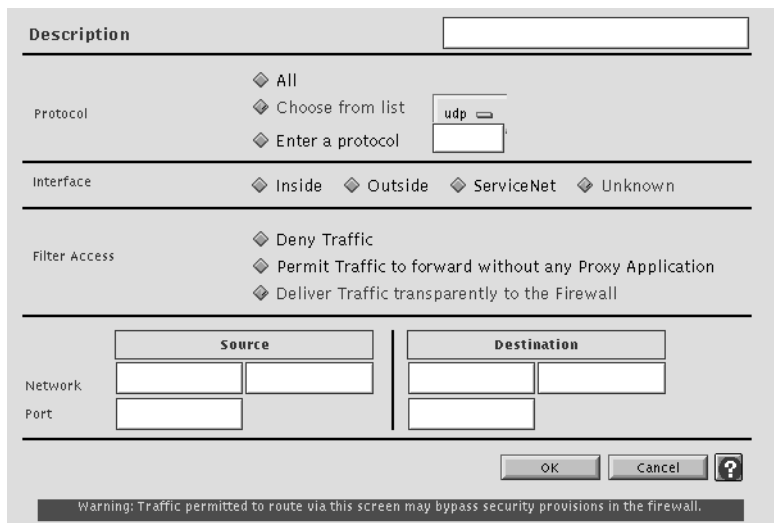


Figure 25-2 Add Packet Screening Rule Window

2. Provide information about your packet screening rule.

- Description** Description for this packet screening rule.
- Protocol** Protocol for which this rule applies. If you want to specify a subtype of the ICMP protocol, use the syntax detailed in “Protocol Rules” on page 230.
- Interface** Interface for which this rule applies.
- Filter Access** Action the packet screening rule takes. If you want the firewall to deny this type of traffic, select “Deny Traffic.” If you want the firewall to absorb this type of traffic, select “Deliver Traffic transparently to the Firewall.” If you want the firewall to forward this traffic to the destination host, select “Permit Traffic” to forward without any Proxy Application.
- Source IP Address** IP address and mask of the source host for this packet screening rule. Use 0.0.0.0:0.0.0.0 to indicate any host.
- Source Port** Port number for source traffic.
- Destination IP Address** IP address and mask of the destination host for this packet screening rule. Use 0.0.0.0:0.0.0.0 to indicate any host.

Destination Port number for destination traffic.

3. Click *OK*.
4. Define the opposite rule on the opposite interface, if necessary.
5. Order your new user restriction rule, as described below, so that the firewall uses your new rule in the right order.

Loading Packet Screening Rules

The Gauntlet Firewall kernel contains the packet screening rules.

To load packet screening rules, apply your changes before exiting the Gauntlet Firewall Manager.

The next time you reboot your firewall, the firewall begins to use the new packet screening rules.

Modifying Packet Screening Rules

To modify packet screening rules:

1. Select the rule you want to modify.
2. Change the properties about the rule you want to modify.
3. Repeat the preceding steps for related rules.

Deleting Packet Screening Rules

To delete packet screening rules:

1. Select the rule you want to delete.
2. Click *Delete*.
3. Repeat the same steps for related rules.

Changing Order of Precedence

Remember that the order in which you place your packet screening rules is very important. The firewall reads them from top to bottom and applies the first one that matches. You generally want to place the most restrictive rules first.

To change the order of precedence:

1. Select the packet screening rule you want to move.
2. Click *Move Up* or *Move Down* as many times as necessary to move the rule to the desired position.

Verifying Your Configuration

Try to send traffic that uses your new packet screening rule through the Gauntlet Firewall. Watch the log for errors. Use the *ipfs* utility trace option to display detailed information about each packet processed:

```
# ipfs -t on
```

When you are done testing, confirm that the trace option is off. The trace option creates large logs:

```
# ipfs -t off
```

PART FIVE

Managing Additional Firewall Services

Chapter 26

Managing Logging and Reporting

Chapter 27

Verifying Integrity

Chapter 28

Managing Virtual Private Networks

Chapter 29

Managing Web and Gopher Servers

Chapter 30

Managing Content Scanning

Chapter 31

Managing URL Filtering

Chapter 32

Managing the Network Management Agent

Chapter 33

Login Shell

Managing Logging and Reporting

Logging (the creation of logs by the firewall) is an important part of a properly configured firewall. Gauntlet Firewall system administrators use information in logs to analyze usage statistics, monitor activities, check for problems, and investigate potential attacks. The logging features of the Gauntlet Firewall provide system administrators with information about activities to and through the firewall. The logging features present information in several formats. Configure the Gauntlet Firewall's logging and reporting features to match your organization's security policies.

This chapter describes the concepts behind logging and reporting systems, configuring these systems, and understanding the logs and reports. The chapter consists of the following sections:

- "Understanding Logging and Reporting" on page 239
- "Creating Logs" on page 240
- "Configuring Logs" on page 241
- "Creating Reports" on page 242
- "Configuring Reports" on page 244
- "Reading Logs and Reports" on page 248

Understanding Logging and Reporting

The Gauntlet Firewall follows the philosophy that it is easy to compress, consolidate, summarize, and delete log information; it is impossible to retroactively gather log information on an event that has already occurred. Disk space is a lot cheaper than spending many hours debugging a problem that a program would have written to the logs. For these reasons, the components of the Gauntlet Firewall log a wide variety of activities and attributes.

These are the components of the Gauntlet Firewall:

- firewall kernel
- proxies
- authentication management system
- DNS
- *sendmail*

These are the attributes logged:

- source IP address
- destination IP address
- source port
- destination port
- user name
- session date and time
- number of bytes transferred
- individual commands (for some activities)
- successful access attempts
- unsuccessful access attempts
- errors in configuration

Creating Logs

The proxies, kernel, and authentication management system automatically write information to the logs. These programs call the standard IRIX system log command (*syslog*) to write information to the standard IRIX log file in */var/adm/SYSLOG*. You do not need to do anything special to create the logs. Even if you choose not to do anything with the information in the logs, the programs still write the information. You never know when you might need it.

The system log file also contains information from other programs, such as *bind*, *cron* and other IRIX utilities that use the *syslog* command.

As with any other information that the *syslog* function writes, the firewall log information is ASCII text. People and shell scripts can easily parse the information.

The Gauntlet Firewall uses other standard IRIX tools to manage logs. Every night, the *cron* daemon runs a shell script that rotates, compresses, and removes log files. The Gauntlet Firewall stores the last two days of logs in ASCII format. By default, the firewall also stores the most recent 14 days in compressed format.

Configuring Logs

The default logging options included with the Gauntlet Firewall meet the needs of most security policies. You do not need to set or modify any options if you wish to use the default configuration, which logs all of the information described above, and retains the logs for 14 days. You can, however, customize the contents and retention of the log.

Configuring Proxy Logging

A few of the proxies (FTP, HTTP, lp, and the Info Server) can log specific commands. For example, the FTP proxy can create a log entry for each command (STOR, RETR, CWD, LIST) it receives.

To modify commands that the proxies log:

1. From within the Gauntlet Firewall Manager, select Services.
2. Select the tab for the proxy you wish to configure.
3. Turn on the logging options for that proxy.

Configuring Log Retention Time

The Gauntlet Firewall allows you to configure the number of days of logs you want to keep on the firewall. Be sure that you have sufficient disk space on your firewall for the logs. Even compressed log files can be large files.

To set log retention time:

1. From within the Gauntlet Firewall Manager, select Reports.
2. Click the Configure tab.

The Configure window displays.

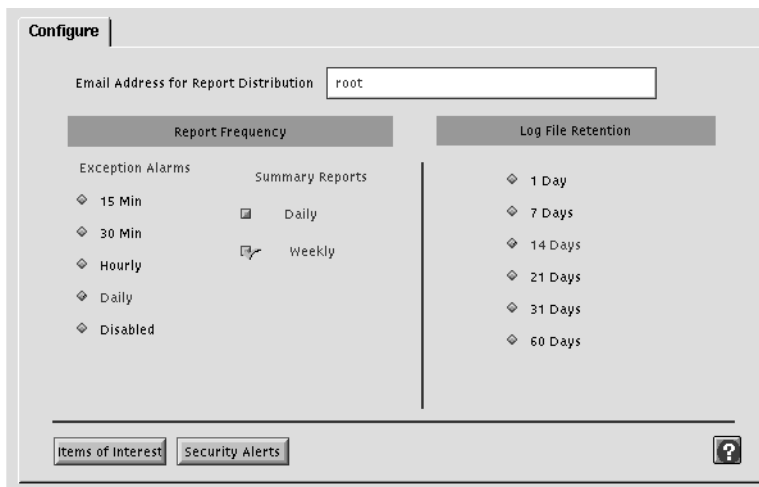


Figure 26-1 Configure Window

3. Under Log File Retention, select the number of days you want the firewall to keep the log files.

Creating Reports

The Gauntlet Firewall contains several reporting mechanisms that sort through the log files and summarize information. The Gauntlet Firewall automatically generates reports. The *cron* daemon runs a set of shell scripts that parse the information in */var/adm/SYSLOG*. You do nothing to create reports; the Gauntlet Firewall does it automatically.

The Gauntlet Firewall includes two types of reports:

- Service Summary reports
- Exception reports

Creating Service Summary Reports

The Service Summary reports include usage and user information on a per service basis. For example, the default report for the TELNET gateway indicates the top 100 clients by connections, the top 100 clients by amount of traffic, and the top 100 denied clients.

To create the Service Summary reports, each night the *cron* daemon on the Gauntlet Firewall runs the scripts to summarize the logs for each service.

When you turn on this option (by default, it is off), the Gauntlet Firewall e-mails the Service Summary reports to root each night. You can modify the e-mail recipient and disable the mail notification. The Gauntlet Firewall does not store the daily Service Summary Report.

Each week, the *cron* daemon on the Gauntlet Firewall runs scripts to summarize services for the past week. By default, each week the Gauntlet Firewall e-mails the Service Summary reports to root. You can modify the mail recipient and disable e-mail notification, but the Gauntlet Firewall still creates the Service Summary Report. The Gauntlet Firewall does not store the weekly Service Summary Report.

Creating Exception Reports

Exception reports include noteworthy items. The Gauntlet Firewall defines a list of items that are not noteworthy and ignores those entries in the reports. The Gauntlet Firewall treats all other events as possible security events. Thus, any item that you do not specifically tell the Gauntlet Firewall to ignore, it reports. Exception reports include information that could indicate a possible attack or other problems.

For example, the Gauntlet Firewall default is to ignore successful authentications when parsing the log file. Successful authentication attempts are a normal part of Gauntlet Firewall activity. However, unsuccessful authentication attempts could be a sign of a potential attack. Therefore, the Exception Report includes all unsuccessful authentication attempts.

To create Exception reports, the *cron* daemon periodically (the default is once a day) runs a reporting script. The script scans log files for security alerts it can ignore, as defined in a different configuration file.

The script summarizes all noteworthy items since the last time the script created a report. By default, the Gauntlet Firewall e-mails Exception reports to root. You can modify the

e-mail recipient, disable the e-mail notification, and set the frequency interval. The script does not store the Exception Report.

Creating Other Reports

Because the Gauntlet Firewall uses standard IRIX logging mechanisms, the logs it creates are in an easily accessible format. You can write your own scripts or programs to create reports to provide additional information.

Configuring Reports

The Gauntlet Firewall default reporting options meet or exceed the needs of most organization's security policies. To use the reporting default configuration, you do not need to set or modify options. The Gauntlet Firewall e-mails weekly Service Summary reports to root. The firewall uses a set of default activities as items that the reporting scripts ignore.

You can customize the report recipient, and enable or disable daily or weekly Service Summary reports. You can customize the events that the Gauntlet Firewall ignores in the Exception reports, enable and disable Exception reports, and customize the Exception reports reporting interval.

Accessing Report Configuration

To access the report configuration:

1. From within the Gauntlet Firewall Manager, select Reports.
2. Click the Configure tab.

The Configure window displays.

Configuring Report Recipients

You can configure the report recipients for Exception and Service Summary reports.

Note: You must send these reports to the same people. You cannot send the Service Summary reports to one person and the exception reports to another.

If you do not specify report recipients, the firewall e-mails the reports to root on the firewall.

To configure report recipients, type the name of the recipient in the e-mail address for Report Distribution field on the Configure window,

The report recipient can be one person, several people, one e-mail alias, or several e-mail aliases.

Configuring Report Frequency

If you have made changes to your firewall's configuration, you may wish to receive reports more often. This can help you to detect a misconfiguration more quickly. If you suspect that someone is trying to attack your firewall, you may also wish to increase the report frequency.

To configure the frequency of Exception reports, click on one of the intervals under Exception Alarms. After each interval, the firewall scans the log, creates a report, and e-mails the report.

To configure the frequency of Service Summary reports:

- If you want the firewall to create and e-mail Service Summary reports every day, under Summary Reports, select Daily.
- If you want the firewall to create and e-mail Service Summary reports every week, under Summary Reports, select Weekly.

Configuring Events to Ignore in Exception Reports

You can configure the patterns and security alerts that the event reporting scripts ignore when parsing the logs. This allows you to configure the Gauntlet Firewall to ignore patterns and routine security alerts.

The firewall still records all events in the log, even if events do not appear on the Exception Report.

You must use regular expressions to configure events to ignore. If you are not familiar with regular expressions, consult your operating system documentation or the programming reference section of a technical bookstore.

As an example, the Gauntlet Firewall creates information messages about the number of bytes transferred by the *lp* proxy. You do not find this information exceptional and do not want to display this information in the Exception reports. The regular expression

```
lp-gw.*bytes.*xferred
```

causes the reporting script to ignore messages in the logs about the number of bytes the *lp* proxy transfers.

To modify items of interest:

1. Click *Items of Interest*.

This starts an editor that displays the current list of items of interest.

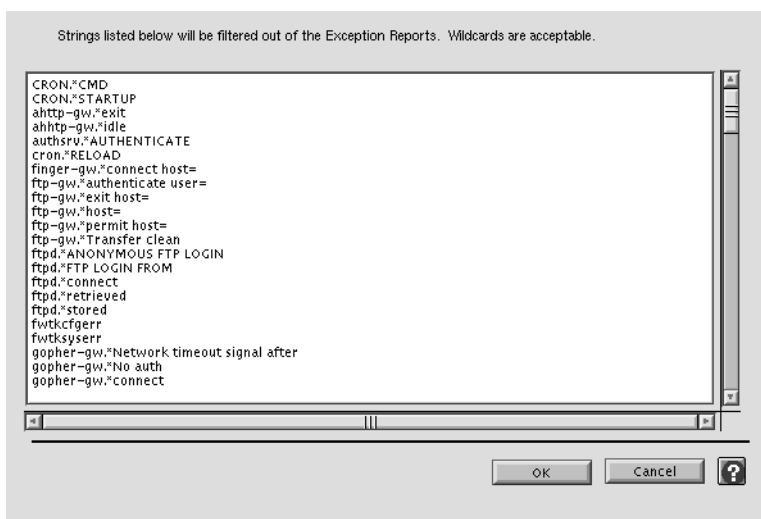


Figure 26-2 Items of Interest Window

2. Enter regular expressions for items you do not wish to see in your reports.
3. Click *OK*.

Security Alerts

A system on your trusted network often sends UDP requests on port 53 that cause security alerts on the Gauntlet Firewall. You know this is because of inadequate settings for time-outs in the DNS code that your system is running. You do not find this information exceptional and do not want to see it in your Exception reports. The regular expression

```
kernel: securityalert: udp from 10.0.1.126:53 to
```

causes the reporting script to ignore messages from the kernel indicating UDP errors on port 53 from the problematic system.

To modify security alerts:

1. Click *Security Alerts*.

This starts an editor that displays the current list of security alerts.

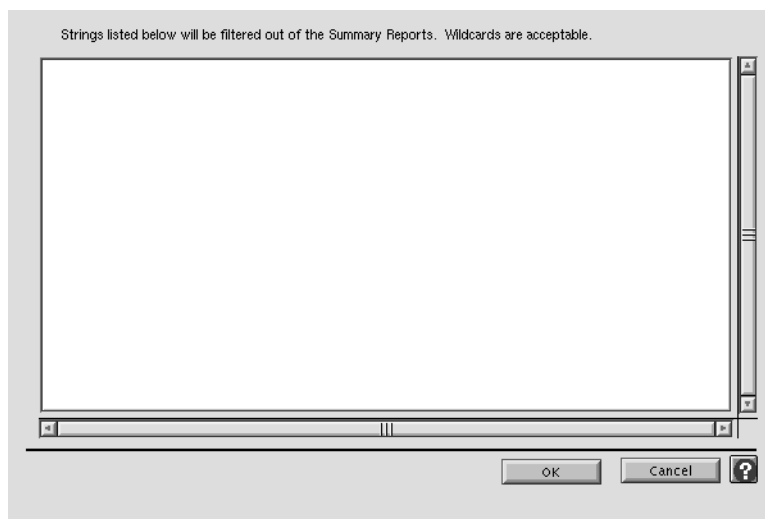


Figure 26-3 Security Alerts Window

2. Enter regular expressions for security alerts you do not wish to see in your exception reports.
3. Click *OK*.

Reading Logs and Reports

The Gauntlet Firewall writes logs and reports as ASCII text. People and shell scripts can easily use the logged information.

This section presents a brief overview of what the logs and reports look like and what each entry indicates.

Reading Logs

The log file (*/var/adm/SYSLOG*) contains a chronological list of events written by the kernel, proxies, authentication management system, and other processes. The example below shows events the Gauntlet Firewall logged in a two-minute period between 10:47:00 and 10:48:59.

```
Oct 30 10:47:22 fire-in http-gw[12079]: permit host=unknown/10.0.1.17 use of gateway
Oct 30 10:47:22 fire-in http-gw[12079]: log host=unknown/10.0.1.17 protocol=HTTP cmd=dir dest=www.sgi.com
path=/
Oct 30 10:47:23 fire-in http-gw[12079]: content-type= text/html
Oct 30 10:47:23 fire-in http-gw[12079]: exit host=unknown/10.0.1.17 cmds=1 in=2392 out=0 user=unauth
duration=6
Oct 30 10:47:23 fire-in http-gw[12080]: permit host=unknown/10.0.1.17 use of gateway
Oct 30 10:47:23 fire-in http-gw[12080]: log host=unknown/10.0.1.17 protocol=HTTP cmd=get dest=www.sgi.com
path=/art/actual/title.gif
Oct 30 10:47:25 fire-in http-gw[12080]: content-type= image/gif
Oct 30 10:47:27 fire-in http-gw[12080]: exit host=unknown/10.0.1.17 cmds=1 in=5581 out=0 user=unauth
duration=4
Oct 30 10:47:28 fire-in http-gw[12081]: permit host=unknown/10.0.1.17 use of gateway
Oct 30 10:47:28 fire-in http-gw[12081]: log host=unknown/10.0.1.17 protocol=HTTP cmd=get dest=www.sgi.com
path=/art/buttons/2.netsec.gif
Oct 30 10:47:28 fire-in http-gw[12081]: content-type= image/gif
Oct 30 10:47:28 fire-in http-gw[12081]: exit host=unknown/10.0.1.17 cmds=1 in=135 out=0 user=unauth
duration=0
Oct 30 10:48:24 fire-in smap[12082]: connect host=cosmo.clientsite.com/192.94.214.96
Oct 30 10:48:24 fire-in smap[12082]: host=cosmo.clientsite.com/192.94.214.96 bytes=1005
from=<bob@clientsite.com> to=<@fire-out.trusted.com:clancy@yoyodyne.com >
Oct 30 10:48:24 fire-in smap[12082]: exiting host=cosmo.clientsite.com/192.94.214.96 bytes=1005
Oct 30 10:48:39 fire-in sendmail[12084]: KAA12084: from=<bob@clientsite.com>, size=921, class=0,
pri=30921, nrcpts=1, msgid=<9510301544.AA04030@clientsite.com>, relay=uucp@localhost
Oct 30 10:48:39 fire-in smapd[12083]: delivered file=sma012082
Oct 30 10:48:40 fire-in sendmail[12086]: KAA12084: to=<@firewall.yoyodyne.com:clancy@yoyodyne.com>,
ctladdr=<bob@clientsite.COM> (6/0), delay=00:00:01, mailer=smtp, relay=mail.yoyodyne.com. [10.0.1.126],
stat=Sent (Ok)
```

Reading Service Summary Reports

Service Summary reports contain a concise overview of events by service (proxy). The example shows weekly information for TELNET activity through the Gauntlet Firewall.

Telnet/Rlogin Proxy Usage

 Top 100 telnet gateway clients (total: 308)

Connects	Host/Address	Input	Output	Total
-----	-----	-----	-----	-----
287	dimension.yoyodyne.com/	267484	11412	278896
6	eight.yoyodyne.com/10.0	495575	2249	497824
6	jersey.yoyodyne.com/10.	291915	3608	295523
3	lizardo.yoyodyne.com/10	4204	318	4522
2	john.yoyodyne.com/10.0.	472366	4719	477085
2	planet10.yoyodyne.com/1	123	64	187
1	blaze.clientsite.com/20	169588	1473	171061
1	unknown/204.254.155.2	0	0	0

Top 100 telnet gateway clients in terms of traffic

Connects	Host/Address	Input	Output	Total
-----	-----	-----	-----	-----
287	dimension.yoyodyne.com/	267484	11412	278896
2	john.yoyodyne.com/10.0.	472366	4719	477085
6	jersey.yoyodyne.com/10.	291915	3608	295523
6	eight.yoyodyne.com/10.0	495575	2249	497824
1	blaze.clientsite.com/20	169588	1473	171061
3	lizardo.yoyodyne.com/10	4204	318	4522
2	planet10.yoyodyne.com/1	123	64	187
1	unknown/204.254.155.2	0	0	0

Reading Exception Reports

Exception reports contain a chronological summary of security alerts and potential items of interest. Exception reports contain several sections: Security Alerts, System Warnings, and Possible Items of Interest.

Exception Reports: The Security Alerts Section

The security alerts section includes information about potential security violations. Common items in this section include:

- Requests for unserved ports. These requests come from other systems looking for a particular service on your firewall. For example, if someone tries to access print services on your firewall, and you are not running the lp proxy, the firewall creates a log message indicating that it received a request for service on port 515.
- DNS spoofing problems. These problems may be malicious attempts by a host on the outside network to act as a host on your inside network. More commonly, the firewall receives DNS requests from misconfigured DNS servers.

Exception Reports: The System Warnings Section

The system warnings section includes information about possible configuration errors. Common items in this section include:

- Incorrect lines in the netperm table. Check the spelling and capitalization. Make sure that you have also used the correct syntax.
- System resource problems.

Exception Reports: Possible Items of Interest Section

The possible items of interest section includes information from other components of the firewall, which may indicate a problem or concern. Common items in this section include:

- Notices about disk space
- Configuration problems in the sendmail program
- Users becoming root

Exception Reports: Example

The example shows information captured and logged over a one-day interval by the Gauntlet Firewall about security alerts, system warnings, and possible areas of interest.

Security Alerts

```
-----
Aug  5 00:01:41 fire-in kernel: securityalert: udp from 10.0.1.68:1140
to 129.241.131.10 on unserved port 29659
Aug  5 09:01:16 fire-in kernel: securityalert: tcp from 10.0.1.128:3616
to 10.0.1.159 on unserved port 14464
Aug  5 09:01:29 fire-in kernel: securityalert: tcp from 10.0.1.128:3621
to 10.0.1.159 on unserved port 14464
Aug  5 12:32:01 fire-in kernel: securityalert: tcp from 10.0.1.68:2089
to 199.201.68.101 on unserved port 8888
```

System Warnings

```
-----
Aug  5 08:41:49 fire-in authsrv[17675]: fwtkyserr: srvhear: read
error: Connection reset by peer
Aug  5 08:41:49 fire-in authsrv[17677]: fwtkyserr: srvhear: read
error: Connection reset by peer
```

Possible Items of Interest

```
-----
Aug  5 16:09:39 fire-in authsrv[18432]: administrator LIST
Aug  5 16:10:38 fire-in su: BAD SU crawhide to root on /dev/tty0
Aug  5 16:10:44 fire-in su: crawhide to root on /dev/tty0
Aug  1 17:02:31 fire-in gui[13316]: permit
host=fire-in.trusted.com/10.0.1.159 use of server(5)
Aug  1 17:02:32 fire-in gui[13316]: protocol=http
file=/usr/local/etc/guidb/D/Dgui/Hstartup0html
Aug  1 17:02:32 fire-in gui[13316]: exit in=22 out=3320
Aug  1 17:02:32 fire-in ahttp-gw[15822]: exit id=283 in=3416 out=240
Aug  1 17:02:32 fire-in gui[13316]: permit
host=fire-in.trusted.com/10.0.1.159 use of server(5)
Aug  1 17:02:32 fire-in gui[13316]: protocol=http
file=/usr/local/etc/guidb/D/Dgui/Dimages/Hbackground00gif
Aug  1 17:02:32 fire-in gui[13316]: exit in=32 out=3090
Aug  1 17:02:32 fire-in gui[13316]: permit
host=fire-in.trusted.com/10.0.1.159 use of server(5)
Aug  1 17:02:32 fire-in gui[13316]: protocol=http
file=/usr/local/etc/guidb/D/Dgui/Dimages/Hiconbar-startup0gif
Aug  1 17:02:32 fire-in gui[13316]: exit in=36 out=2999
Aug  1 17:02:32 fire-in gui[13316]: permit
host=fire-in.trusted.com/10.0.1.159 use of server(5)
```

```
Aug  1 17:02:32 fire-in gui[13316]: protocol=http
file=/usr/local/etc/guidb/D/Dgui/Dimages/Hbanner-startup0gif
Aug  1 17:02:32 fire-in gui[13316]: exit in=35 out=4425
Aug  1 17:02:32 fire-in ahttp-gw[15822]: exit id=285 in=3095 out=245
Aug  1 17:02:32 fire-in ahttp-gw[15822]: exit id=284 in=3186 out=241
```

Verifying Integrity

Even if you have followed all the appropriate precautions to make your firewall secure, you may still want an extra level of assurance that no person or process has modified your system. To provide this assurance, the Gauntlet Firewall includes facilities to assist you in verifying the integrity of your system.

The following sections describe the concepts behind system integrity and some common administrative tasks:

- “Understanding System Integrity” on page 253
- “How System Integrity Works” on page 254
- “Configuring Integrity Checks” on page 254
- “Creating an Integrity Database” on page 256
- “Updating the Integrity Database” on page 257
- “Verifying System Integrity” on page 257

Understanding System Integrity

The Gauntlet integrity database is a collection of cryptographic checksums or message digests for most files on your filesystem. The database contains a checksum for each file, and includes information about the user ID, group ID, and mode.

To verify the integrity of the system, you can create a new database of checksums, and compare the values with the existing database. Changes in the checksums for a file indicate that someone or something has modified the file in some manner.

The database does not contain information about files that can change often, such as the mail spool, the log files, and system aliases. You expect these files to change, so the checksums would always be different. You can modify the list of files that the integrity checker ignores. Unless you specifically tell the Gauntlet Firewall to ignore a particular item, it provides information on that item.

How System Integrity Works

The Gauntlet Firewall allows you to create integrity databases. When you create an integrity database, the firewall runs a scan program (*scan*) that walks the directory tree and creates MD5 message digest checksums for each file. It ignores the files listed in a *scan* configuration file. The *scan* program writes the checksums to the integrity database.

Updating the database reinitializes the integrity database to reflect the current state of your firewall.

Configuring Integrity Checks

Configuring integrity checks consists of accessing the Integrity Check window and selecting files to ignore.

Accessing Integrity Checking

To access integrity checking:

1. From within the Gauntlet Firewall Manager, select Maintenance.
2. Click the Integrity Check tab.

The Integrity Check window displays.

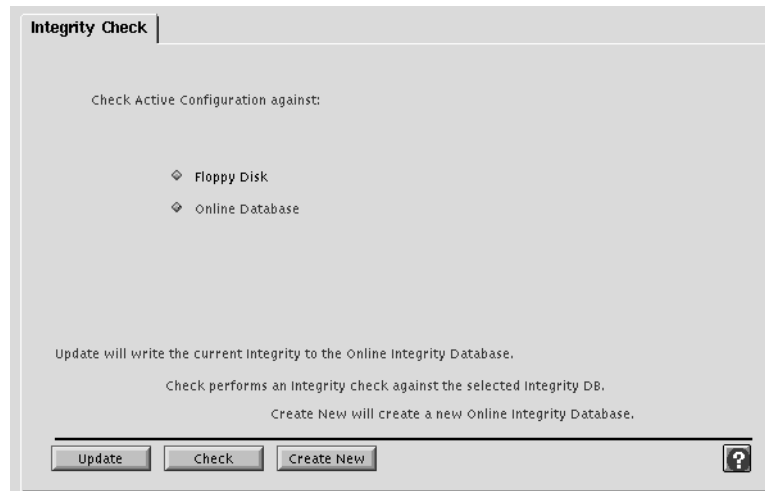


Figure 27-1 Integrity Check Window

Configuring Files to Ignore

You can modify the list of files and directories that the *scan* program ignores when creating and checking databases. This allows you to ignore directories and files that you know are volatile.

To configure the files to ignore:

1. Log in to the firewall and become root.
2. Use your favorite text editor to edit `/usr/local/etc/checksums/scan.conf`.
3. Modify the list of directories and files the integrity checker ignores.

The next time you create or update an integrity database, the firewall does not create database entries for the files you added.

Creating an Integrity Database

You can create an integrity database anytime you wish. You should create one after you first configure your firewall. Generally, you create a new database when you have made major modifications to your firewall. Of course, you can do this more or less often as your security policy dictates.

You can create an integrity database on the firewall. An online database is easy to access, and does not require someone to physically place media into the firewall. However, an online database can be accessed by others who have access to the firewall.

Note: Creating an integrity database scans your entire directory tree. On systems with large disks, this may take as long as 10 or 20 minutes.

Planning an Integrity Database

When planning an integrity database:

1. Make sure you are satisfied with the current state of your Gauntlet Firewall.
2. Make sure you have sufficient space for your database. The database file is usually between 2 MB and 5 MB in size, but it can be larger.

Creating the Database

To create an integrity database:

1. On the Integrity Check window, click *Online Database*.
2. Click *Create New*.

The firewall creates the integrity database.

Protecting the Integrity Database

You use the integrity database to verify that nothing has modified your system. Therefore you must protect the database itself from tampering. Consider the following options to protect your integrity database (*/usr/local/etc/checksums/gauntlet.sum*):

- Copy the integrity database to removable media that you keep off line for safekeeping. Store this copy of the database according to your security policies.
- If you leave the database online, protect the database file by removing write permission for groups and others.
- Store the copy of the initial integrity database that you created during your installation with your original distribution media.

Updating the Integrity Database

Updating the integrity database creates a small database file that contains only the updates.

Updating the Database

Before you start, make sure you have already created an integrity database.

To update the integrity database:

1. On the Integrity Check window, click *Online Database*.
2. Click *Update*.

The firewall updates the integrity database.

Verifying System Integrity

When you wish to check the integrity of the files on your system, you can verify the current filesystem checksums against the known, good filesystem checksum information stored in the integrity database. This process runs the scan program (*scan*), which walks the directory tree and creates MD5 checksums for each file. It ignores the files listed in a *scan* configuration file. The *scan* program compares the checksums for each file with the information in the integrity database, and writes any differences to a temporary file (*/tmp/scan.processID*). The *scan* program leaves the changes list in the */tmp* directory, so you can view them later.

The firewall can compare checksums with information in a database stored online. It can also compare with a database on a floppy diskette.

You can verify your system integrity anytime you wish. Administrators often perform these types of tasks monthly. You may also wish to check system integrity any time you notice unusual behavior on your firewall. You can, of course, define what you consider unusual according to your security policies.

Check Integrity

To check integrity against an online database:

1. On the Integrity Check window, click *Online Database*.
2. Click *Check Integrity*.

The firewall creates a new integrity database and compares it with the old database.

Viewing and Understanding the Results

Review the changes noted in the changes file and make sure they are acceptable changes. For example, you may have added users to your authentication database since you last created your integrity database. The authentication database information is different, and the integrity check notes this. This is an acceptable difference.

The *scan* program does not tell you what has changed about the file. It only tells you that the file has changed.

To view the results of an integrity check:

1. Log in to the firewall and become root.
2. Use your favorite text editor to view the file (*/tmp/scan.processID*) that contains the list of changes.

Managing Virtual Private Networks

Packets on the Internet flow through a variety of wires and fibers owned and managed by a variety of organizations. The opportunities for someone or something to monitor these packets are large.

The Gauntlet Firewall can be used to create a Virtual Private Network (VPN). VPNs use encryption to allow secure communication between various points within the network.

The following sections explain how you can use your Gauntlet Internet Firewall to exchange encrypted traffic with other Gauntlet Firewalls:

- “Understanding VPNs” on page 260
- “How Virtual Private Networks Work” on page 264
- “Accessing Encryption Key Configuration” on page 265
- “Working With Encryption Keys” on page 265
- “Accessing the VPN Configuration” on page 268
- “Working With the VPN Configuration” on page 268

Note: This feature is available only in the Unites States domestic version of the Gauntlet product.

Understanding VPNs

When using a single firewall, the defense perimeter includes the network of systems that sit behind the firewall, inside the perimeter. Communication with any other systems or networks outside the perimeter is over some untrusted network, such as the Internet. A VPN extends the defense perimeter to include other networks and systems.

For example, Yoyodyne has offices in Maryland and California, each protected by a Gauntlet Internet Firewall. When these offices communicate, they use the Internet. Yoyodyne can create a VPN and extend the defense perimeter from its corporate headquarters in Maryland to include the network of systems behind the defense perimeter in its California office, as shown in Figure 28-1.

A VPN is considered private because all traffic that passes through the firewall to another part of the VPN is encrypted. Any program watching the packets flow by would simply see a stream of encrypted data. Without the key used to encrypt the data, the information isn't useful to snoopers. Because the remote host or network shares a key with the firewall, it can decrypt and process the encrypted packets that it receives. In Figure 28-1, all traffic between the firewall in the Maryland office and the firewall in the California office is encrypted.

A VPN is considered a virtual network because you are extending the network from the systems that are physically within the defense perimeter to include other systems or networks that are not. A VPN may or may not trust the other network.




Figure 28-1 Example VPN

Privacy With Trust (Trusted Link)

A VPN with trust expands the concept of trust (as in trusted networks) to include not only the systems within your defense perimeter but also all of the systems within the remote defense perimeter. For all intents and purposes, all of these systems are part of the same network within the same defense perimeter. Any activities that you allow within your network can be used with systems on the remote network.

For example, Yoyodyne allows users in the Maryland office to use the network time protocol (NTP) within the network to set the clocks on their systems. If Yoyodyne sets up a VPN with the California office using a trusted link, they can now use *ntp* with systems in the California office.

You can create trusted links for host-to-host, network-to-network, or host-to-network communications. This allows you to trust individual hosts or entire networks.

A VPN also allows any IP services you desire to pass between the two firewalls. The services simply need to be IP-based. You can allow applications that use the user datagram protocol (UDP) or the transmission control protocol (TCP).

In addition to sharing a defense perimeter against the rest of the world, sites that create a VPN must share the security perimeter in other ways. These sites should share the same policies, procedures, and administrative control. If the security policy for the Maryland office does not allow TELNET from remote locations, then the security policy for the California office should match this. If the policies differ, someone can simply come in through the California office and then connect directly to a system in the Maryland office, which is part of the same VPN.

Privacy Without Trust (Private Link)

A VPN without trust does not expand the concept of trust to include the systems within the remote defense perimeter. In this case, the traffic between the two networks is encrypted, providing the privacy. Once it decrypts the traffic, the remote firewall still considers the request as being from an untrusted network. The request is the same as any other that comes from an untrusted network, but with the additional benefit of encryption.

For example, Yoyodyne sets up an untrusted VPN between the Maryland and California offices. Traffic between the two offices is still encrypted. When the firewall for the California office receives and decrypts a TELNET request from a system at the Maryland office, it treats the request as it would any other untrusted network. They cannot send UDP packets between the two networks or trust NTP from the other site as they could using a VPN with privacy with trust.

You can create private links for host-to-host, network-to-network, or host-to-network communications. The most common use of privacy without trust creates a private link between two networks.

Sites that create a VPN without trust must of course share the encryption key that gives them the privacy. However, they can now use different policies and procedures and have different administrative control.

Encryption Through Multiple Firewalls (Passthrough Link)

A VPN can use encryption through a series of firewalls. In this case, the traffic between the outer firewalls is encrypted, but the firewalls in between simply pass the encrypted data through. They do not decrypt the data nor do they have the encryption key.

For example, Yoyodyne sets up a VPN (with or without trust) between the firewall for the accounting department in Maryland and the firewall for the accounting department in California. On the firewall for the entire Maryland office (which includes the accounting department), Yoyodyne creates a passthrough link. This link simply passes the encrypted traffic from the accounting firewall in Maryland to the accounting firewall in California. The administrators in the California office must create a similar passthrough link on their firewall to pass encrypted traffic to the accounting firewall in the California office.

You can create passthrough links for host-to-host, network-to-network, or host-to-network communications. The most common use of a passthrough link specifies a host-to-host link for two firewalls.

How Virtual Private Networks Work

The firewall handles VPNs by examining all outbound traffic and encrypting any traffic between hosts that are marked as encrypted peers. The exact sequence of events varies depending on whether there is privacy with trust or just privacy.

When the firewall is about to send a packet, it checks to see if the source and destination are listed in a table of encrypted pairs. If the source and destination match an entry in the table, the firewall hands the packet to the *swIPe* driver for encryption.

Encrypting the Data

The *swIPe* driver uses the Data Encryption Standard (DES) to encrypt the data using the key provided for this VPN during firewall-to-firewall configuration. The new packet contains encrypted data and a header that indicates this is a special encrypted protocol. The firewall then sends the encrypted packets across the Internet (or other untrusted network) to the firewall for the remote network.

When the remote firewall receives the packet on its outside interface, the IP input layer recognizes this as an encrypted packet because of the special protocol. This information indicates that the firewall should send any packets with this special protocol to the *swIPe* driver.

If the source and destination addresses in the packet indicate that it is part of a passthrough link, the *swIPe* driver forwards the packet without modifying it.

Decrypting the Data

The *swIPe* driver decrypts the data using the same key used to encrypt the data. The *swIPe* driver passes the now decrypted data back to the IP input layer. This now handles the packet as it would handle any other packet that it receives on the outside interface.

Routing the Packet

If the VPN between the two networks uses privacy with trust, the routing layer forwards the packet to the appropriate host on the inside network. If the VPN between the two networks uses just privacy with no trust, the routing layer hands the packet to the appropriate service or proxy. The proxies treat this packet as they would any other packet from any other untrusted network.

Accessing Encryption Key Configuration

The first step to creating a VPN is to create a key.

To access VPN configuration:

1. From within the Gauntlet Firewall Manager, select VPNs.
2. Click the Swipe Keys tab.

The Swipe Keys window displays.

Working With Encryption Keys

Before creating your VPN, you must create encryption keys that the firewall can use to encrypt and decrypt traffic.

Planning Encryption Keys

Before you start adding encryption keys, determine whether you wish to use the same key for every VPN. You must use the same key at each end of a VPN. However, if you have one VPN to your California office and one VPN to your client in New York, you may not wish to use the same key for both networks.

Creating Encryption Keys

To create an encryption key:

1. In the Swipe Keys window, click *Add*.
The Add Swipe Key window displays.

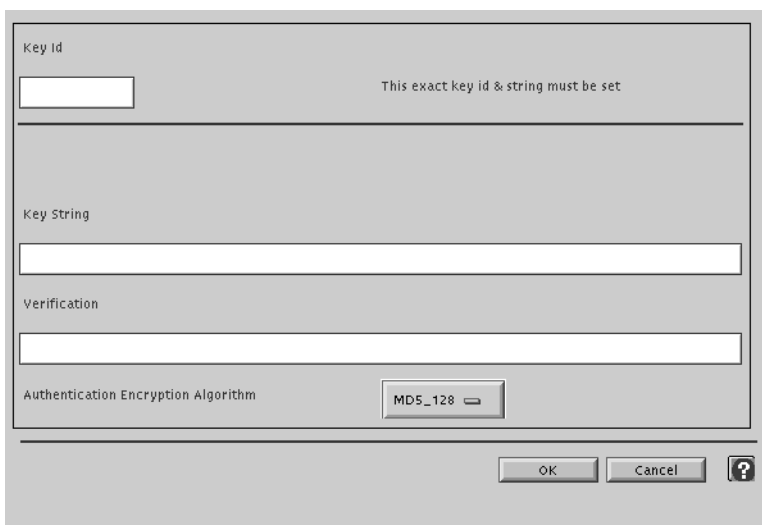


Figure 28-2 Add Swipe Key Window

2. Provide information about your encryption keys.

KeyID	ID number for this key. When you create the trusted or private link, you specify the key you wish to use by Key ID. If you wish to use this key to create a link with a Gauntlet Firewall running Version 4.0 or greater, enter a number between 0 and 999. If you wish to use this key to create a link with a Gauntlet Firewall running a version preceding Version 4.0, enter 0.
-------	---

Key String	<p>Alphanumeric string. The swIPe driver uses an MD5 hashing algorithm to create a pseudo-random number. The swIPe driver uses this value as the DES key.</p> <p>Ensure that the key string is not simply a word found in a dictionary. For best security, use a long string that includes multiple non-alphabetic characters. For example, Yoyodyne uses the phrase</p> <pre>Try&and*guess(this)one.</pre>
Verification	Enter the same value you entered for the Key String.
Authentication Encryption Algorithm	Specifies the authentication algorithm that the swIPe driver uses to verify the contents and source of incoming encrypted packets. Select MD5_128. The larger the hash size, the less likely it is that someone can spoof the verification.

3. Click *OK*.

Modifying Encryption Keys

To modify an encryption key:

1. Select the key you wish to modify.
2. Click *Modify*.
3. Change the settings for this encryption key.
4. Click *OK*.
5. Inform the administrator at the other end of the VPN of your changes.

Deleting Encryption Keys

Note: You cannot delete an encryption key that is currently defined as part of a VPN.

To delete an encryption key:

1. Click the encryption key you wish to delete.
2. Click *Delete*.

Accessing the VPN Configuration

To access VPN configuration:

1. From within the Gauntlet Firewall Manager, select VPNs.
2. Click the VPNs tab.

The VPNs window displays.

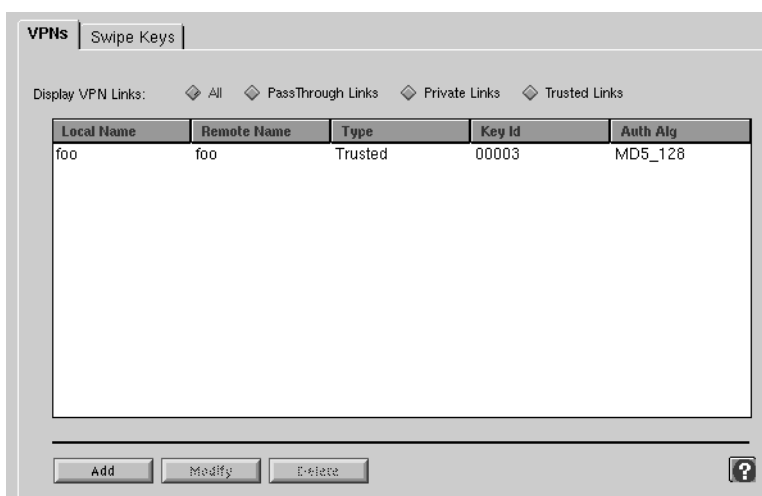


Figure 28-3 VPNs Window

Working With the VPN Configuration

This section discusses how to work with VPNs, from the planning stage to the actual implementation and deletion.

Planning VPNs

When planning the VPNs for your network:

1. Make sure your firewall is working as you would like before you create a VPN.
2. Determine whether you wish to use trusted links or private links.

3. Determine whether you need to create any passthrough links on firewalls that lie between your encrypted links.
4. Coordinate your efforts with the administrator of the remote network. Discuss your security policies and procedures. Prepare to synchronize the firewalls as you configure them.

Creating a Trusted or Private VPN

To create a trusted or private VPN:

1. On the VPNs window, click *Add*.

The Add VPN Link Configuration window displays.

Link Type: Private Trusted PassThrough Mode: Static

Local Network

Network Name

IP Address

Mask

Interface

Remote Network

Network Name

IP Address

Mask

Gateway

Authentication Algorithm SHA1

Key Id 00003

Sequence On Off

OK Cancel ?

Figure 28-4 Add VPN Link Configuration Window

2. Provide information about your VPN.

Link Type	Select the type of link you are creating.
Local Network Name	Name for your local network or host. This name is for your information only, to help you track the various addresses and subnet masks.
Local Network Address and Mask	IP address for the local network or host. Specify the IP address and subnet mask for the host or network. If you are creating a trusted link between the networks inside two firewalls, enter the IP address and mask for your network. If you are creating a private link between two networks that are behind firewalls, enter the IP address of the outside interface of the firewall.
Local Interface	Name of the interface to which the local network is connected.
Remote Network Name	Name for the remote network or host. This name is for your information only, to help you track the various addresses and subnet masks.
Remote Network Address and Mask	IP address for the remote network or host. Specify the IP address and subnet mask for the host or network.
Remote Gateway Address	IP address for the outside interface of the remote firewall. After it encrypts the outgoing packets, the firewalls sends them to this address.
KeyID	Select the Key ID for the encryption key you wish to use for this link.

3. Click *OK*.

Creating a Passthrough Link

To create a passthrough link:

1. In the VPNs window, click *Add*.

The Add VPN Link Configuration window displays.

2. Provide information about your passthrough link.

Local Network Address and Mask IP address for the local network or host that is using the encrypted link. Specify the IP address and subnet mask for the host or network.

This is often the IP address of the outside address of the firewall that is using the encrypted link.

Local Network Name IP address for the local network or host. Specify the IP address and subnet mask for the host or network. If you are creating a trusted link between the networks inside two firewalls, enter the IP address and mask for your network. If you are creating a private link between two networks that are behind firewalls, enter the IP address of the outside interface of the firewall.

Remote Network Address and Mask IP address for the remote network or host that is using the encrypted link. Specify the IP address and subnet mask for the host or network.

This is often the IP address of the outside address of the firewall that is using the encrypted link.

Remote Network Name Name for the remote network or host that is at the other end of the encrypted link. This name is for your information only, to help you track the various addresses and subnet masks.

3. Click *OK*.

Modifying VPNs

To modify a VPN:

1. On the VPNs window, select the VPN you wish to modify.

2. Click *Modify*.

The Modify VPN Link Configuration window displays.

3. Change the settings for this VPN.

4. Click *OK*.

5. Inform the administrator at the other end of the VPN of your changes.

Deleting VPNs

To delete a VPN:

1. Select the VPN you wish to delete.
2. Click *Delete*.

Starting Your VPN

To bring up your VPN:

1. Coordinate your settings with the administrator of the remote VPN.
2. When exiting the Gauntlet Firewall Manager, click *Save*, *Apply*, and *Reboot* to make your configuration take effect.

When your firewall reboots, your VPN is enabled.

Testing Your VPN

After you have brought up your VPN, you may wish to confirm your configuration.

To verify your configuration:

1. Log in to the firewall and become root.
2. Use the trace option of the encryption policy configuration utility (*/usr/local/etc/ipe*) to display kernel diagnostic messages to the console:

```
# cd /usr/local/etc
# ipe -t on
```

3. Watch the console for errors. Verify that there is diagnostic information for each of your links.

The string *Swipe_input* means the firewall encrypted a packet, while *Swipe_output* means the firewall decrypted a packet.

4. Turn off the trace function when you are done to avoid the overhead of the diagnostic feature.

```
# ipe -t off
```

To verify the configuration of a trusted link:

1. Log in to a host other than the firewall within the local network.
2. Use *ping* to reach a host within the remote network.

For example, the Yoyodyne administrator in Maryland pings the mail hub on the California office network:

```
% ping mail.ca.yoyodyne.com
```

Specify the remote host by IP address to make sure that it is not a DNS problem:

```
% ping 10.0.6.12
```

To verify the configuration of a private link:

1. Log in to a host other than the firewall within the local network.
2. Connect to a host within the remote network.

For example, the Yoyodyne administrator uses TELNET to connect to the mail hub on the California network:

```
% telnet 10.0.6.12
```

The remote firewall should pass this request to its TELNET proxy, as it would any request from an untrusted network. By default, the proxy should display the TELNET proxy prompt.

Stopping Your VPN

You can disable your VPN by removing the configuration for the VPN. The next time you reboot your firewall, the firewall does not start the VPN. If you need to disable the VPN and leave your firewall running, use the encryption policy configuration utility (*/usr/local/etc/ipe*) to disable the VPN:

```
# ipe -c
```

Managing Web and Gopher Servers

Sometimes it is not feasible to run a separate WWW or Gopher server outside your firewall. Because of hardware or other constraints, you cannot devote a separate system to be your WWW server. Or, you do not expect enough traffic to justify another system, but still want to offer WWW services to your customers. Instead, you want to run the WWW server securely on the firewall itself. Most WWW servers are large programs, making it harder to ensure that they do not have any security holes.

The Info Server included with the Gauntlet Firewall securely services requests for HTTP and Gopher services.

The following sections explain how the Info Server works, how to configure the server for the various protocols, and how to use the server:

- “Understanding the Info Server” on page 275
- “How the Info Server Works” on page 276
- “Accessing Info Server Configuration” on page 279
- “Configuring the Firewall to Run an Info Server” on page 280
- “Using the Info Server” on page 284

Understanding the Info Server

The Gauntlet Info Server is a minimal Web server. The server, which runs on the firewall, works with a set of management tools to service HTTP and Gopher requests. You can configure the server to allow connections based on:

- source IP address
- source hostname

You would use the Gauntlet Info Server in place of another HTTP server (such as the CERN or Netscape HTTP servers) or Gopher server (such as the University of Minnesota Gopher Server).

The Info Server implements a minimalist design, in which the server handles only the file requests. A variety of management tools (on a per service basis) actually provide the data. These smaller programs are easier to analyze and verify that there are no holes. Simpler code is easier to verify.

How the Info Server Works

The Info Server (*info-gw*) runs on the firewall as a daemon listening for TCP-based requests on port 8000. When the firewall receives a request, it forks a child copy of the Info Server, leaving the parent Info Server to continue listening for requests.

The child Info Server process looks at the request and places it in one of several categories (such as Gopher or HTTP).

It checks the appropriate configuration information and determines whether the requesting host has permission to use the desired service. If the host does not have permission, the Info Server logs the connection and displays an error message.

If the host has permission to use the service, the Info Server uses its internal database (by default in */usr/local/etc/infodb*) to find the requested file or to go to the requested directory. The client thinks it is talking to a regular HTTP or Gopher server, even though it is not.

How the Database Works

When the Info Server processes a request, it does not use standard directory commands to traverse the file hierarchy on the firewall. Instead, the Info Server uses a database manager, which translates the HTTP or Gopher request into the internal database structure. The database manager then tells the Info Server the actual name of the file, which the server displays or returns to the client. The database uses */usr/local/etc/infodb* as the root directory for the database.

The database structure restricts the number of characters that can exist in a filename and translates others. It uses particular letters to designate particular types of files and directories. The database uses the first letter of file and directories to indicate the type of file or directory type.

Info Server Directories

The database structure recognizes only those directories with names that start with the letter D. When the Info Server receives an HTTP request for a file in the *images* directory, the database manager translates the request and looks in the *Dimages* directory.

The database structure also translates other characters in directory names. It translates the dot (.) character in file names to the zero (0) character. When the Info Server receives a request to go to the directory */../etc*, the database manager translates the request and looks for the directory *D/D00/D00/Detc*. Because the root directory of the database is actually */usr/local/etc/infodb*, the Info Server is actually looking for */usr/local/etc/infodb/D/D00/D00/Detc*.

Note: The Info Server always looks for files within its own directory tree. It does not and *cannot* move back out of its directory tree to other areas of the systems, as some HTTP or Gopher servers do.

Info Server Data Files

The database structure recognizes only those data files that start with the letter A. When the Info Server receives a request for the file *readme*, the database manager translates the request and look for the file *Areadme*. HTTP and Gopher requests both return these files.

The database structure looks for HTTP header files (for HTTP version 1.0) in files that start with H. These contain information about the document, such as content length and encoding, and are sent with each download.

The database structure also limits the characters that can exist in filenames. It translates the dot (.) character in filenames to the zero (0) character. When the Info Server receives a request for the file *latest.gz*, the database manager translates the request and looks for the file *Alatest0gz*.

In many cases, the files that start with A and H are actually symbolic links to the real text or binary file. For example, the file *Alatest0gz* would actually be a symbolic link to *latest.gz*. For text files, the A file is generally a copy of the actual data file with every line terminated with a carriage return/line-feed pair. You do not need to create files specifically for use with the Info Server. You merely need to create symbolic links or copies of the files that the database understands.

Info Server Queries and Executable Programs

The database structure recognizes only those query programs and executables with names that start with the letter Q. When the Info Server receives an HTTP request that contains a query, such as `animals?dogs`, the database manager translates the request and tries to run the program *Qanimals*. The database manager passes all the information after the query marker (`?` for HTTP requests and a Tab for Gopher requests) to the query program.

The Q files are usually symbolic links to the executable program. For example, the program *Qimagemap* could actually be a symbolic link to *imagemap*.

Info Server Gopher Menu Files

When the Info Server receives a request to display a Gopher menu, it instead returns a specific file that contains the list of files that you wish to display for that directory. For example, when the Info Server receives a Gopher request for the menu in a directory, the database manager translates this request and returns a file with a name that begins with G in the current directory. The client displays a menu of files that looks like the list of files you might see with any other Gopher server.

The G files are actual files that contain Gopher menus. You create the files, listing only those files that you wish to appear in the menu. For example, the *Gmenu* file could contain only the list of files that you want anyone to view, even though you have other files in the directory.

Accessing Info Server Configuration

To access the Info Server configuration:

1. From within the Gauntlet Firewall Manager, select Services.
2. Click the Info tab.

The Info window displays.

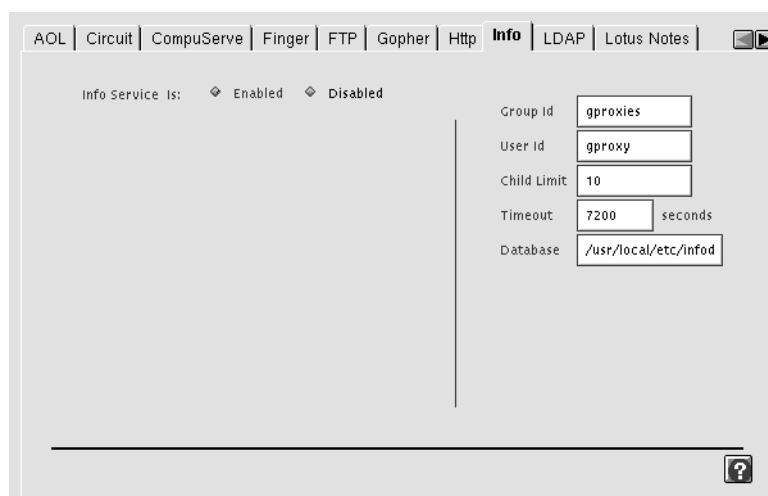


Figure 29-1 Info Window

Configuring the Firewall to Run an Info Server

Configuring the Gauntlet Firewall to run an Info Server involves planning, creating files, placing files on the firewall, adding files to the database, and enabling the proxy.

Planning an Info Server

When planning an Info Server:

1. Determine which services (HTTP, or Gopher) you wish to offer.
2. Determine who will put the files onto the firewall. Remember that if you want your WWW or Gopher administrator to have access, you need to give them an account on the firewall, which is not recommended. Instead, make arrangements with your WWW or Gopher administrator to periodically transfer files for them.

Creating Files for an Info Server

Create your text and executable files as you would for use with any HTTP or Gopher server. You do *not* need to modify references to directories or to executables within your documents.

Placing Info Server Files on the Firewall

To set up your files for use with the Info Server on the firewall:

1. Log into the firewall and become root.
2. Create your directory structure under `/usr/local/etc/infodb/D`.

Prefix each directory with the letter D when you create the directory. For example, if you want to keep all of your pictures in the images directory, use these commands:

```
# cd /usr/local/etc/infodb/D
# mkdir Dimages
```

3. Copy all your files (HTML, text files, executables, and pictures) to the appropriate directory.

Adding Files to the Info Server Database

This process creates the A and H files for HTML files, the Q files for queries, and so forth. The process differs slightly for text and binary files.

Adding Text Files to the Info Server Database

Adding text files to the database creates the necessary A and H files in the database. Use the *addtext* program (*/usr/local/etc/infodb/tools/addtext*).

To add text files to the database:

1. Create the A and H files with these commands:

```
addtext file ctfiletype
```

file is the name of the text file. *ctfiletype* is one of the default header file types used to create an HTTP version 1.0 header file, it can be either of these:

- *cthtml*—HTML text header (default)
- *cttext*—Text header

Consult */usr/local/etc/infodb/tools* for a list of currently available sample headers. Use these files as templates to create your own header files, if necessary.

2. Repeat this process for every file you wish to have accessible via the Info Server.

Adding Binary Files to the Info Server Database

Adding binary files to the database creates the necessary A and H files for images. Use the *addfile* program (*/usr/local/etc/infodb/tools/addfile*).

To add binary files to the database:

1. Create the A and H files with these commands:

```
addfile file cftype
```

file is the name of the binary file. *cftype* is one of the default header file types used to create an HTTP version 1.0 header file, and can be one of the following:

- ctavi—AVI movie header
- ctgif—GIF image header
- cthtml—HTML text header
- ctjpg—JPEG image header
- ctps—PostScript header
- ctqt—QuickTime movie header
- cttext—Text header
- ctzip—ZIP header

Consult */usr/local/etc/infodb/tools* for a list of currently available sample headers. Use these files as templates to create your own header files, if necessary.

2. Repeat this process for each binary file you want to have accessible via the Info Server.

Adding Query Files to the Info Server Database

Adding query files to the database creates the necessary symbolic links for the query file.

To add query files:

1. Create the symbolic link:

```
# ln -s file Qfilename
```

file is the path and file of the actual query executable. *Qfilename* is the name of the executable prepended with a Q and any periods converted to the zero (0) character.

2. Repeat this process for each query file you wish to have accessible via the Info Server.

Creating Info Server Gopher Menu Files

Creating Gopher menu files actually creates the text file that the Info Server displays when it receives a request for a Gopher menu.

To create Gopher menus:

1. Execute the list command and redirect it to a file that starts with G. You may wish to restrict the files that the command displays, so that it looks like a normal Gopher menu. See the *makedirlist* script for examples of redirecting list files to text files for the Info Server.
2. Modify the resulting file and add the other standard Gopher menu fields.

Enabling the Info Server

To enable the Info Server:

1. On the Info window, click Enabled.
2. Add the Info Server configuration to the service groups you want to use the Info Server.
3. Before exiting the Gauntlet Firewall Manager, save and apply your changes.

The firewall enables the Info Server.

Verifying Your Info Server Setup

Access your Info Server as you would any other HTTP or Gopher server, specifying port 8000 in the URL. Watch the log messages.

Using the Info Server

Advertise your HTTP or Gopher Server to your customers or the world. Be sure to:

- Advertise the outside IP address of the firewall for customers outside your defense perimeter. For example, Yoyodyne tells customer to visit this web site:

`http://fire-out.yoyodyne.com:8000/welcome.html`

For added security, you should not advertise that the Info Server is running on the firewall. For example, Yoyodyne should create another record in its DNS database indicating that the name `www.yoyodyne.com` maps to the outside interface of the firewall. This would allow them to use this URL:

`http://www.yoyodyne.com:8000/welcome.html`

- Advertise the inside IP address of the firewall for your own users. For example, Yoyodyne tells its employees on the inside network to use this URL:

`http://10.0.1.100:8000/welcome.html`

- Specify which connections should use port 8000 for HTTP and Gopher requests.

Updating Info Server Files

To update files on the Info Server:

1. Make changes to your original file, not the file that the Info Server creates.
2. Add the file to the database. This replaces the previous version of the file with your new file.

Managing Content Scanning

Your security policy allows users to transfer files from other systems to their systems on the network. They can exchange mail with colleagues, view Web sites, and transfer files from other sites on the Internet. These activities are essential to your business. You want to make sure that the files that come into your network are not infected with viruses. The Gauntlet Firewall includes content scanning as a configurable option for transferring files (via FTP), viewing Web sites (via HTTP), and exchanging e-mail (via SMTP).

This chapter discusses the concepts behind content scanning and explains how it works, how to configure the proxies for content scanning, and how to configure your content scanning engine. The chapter consists of the following sections:

- “Understanding Content Scanning” on page 285
- “How Content Scanning Works” on page 287
- “Accessing Content Scanning Configuration” on page 288
- “Configuring the Firewall to Use Content Scanning” on page 289
- “Configuring the Content Scanner” on page 291

Understanding Content Scanning

Several of the Gauntlet proxies provide content scanning. The content scanning option runs as part of the proxies. You can configure content scanning on the following protocols:

- FTP
- HTTP
- SMTP

Because content scanning is a configurable option in several different proxies, you can configure the firewall to scan transferred files for some hosts and not for others. For example, you could enable content scanning for e-mail only, if you were concerned about the types of files people are sending to your employees.

FTP

File transfers via FTP can contain threatening contents. You can use the content scanning capability to identify and block these transfers. Scanning can be configured independently for file retrieval (GET) and transmission (PUT). When you enable content scanning, the firewall inspects the file before it delivers the file.

SMTP Mail

It is important to note that the firewall scans only mail that enters the firewall via SMTP. The firewall does not scan mail that passes through the firewall using the POP3 protocol. Because the POP3 proxy is intended for use from outside to inside only, this is generally not a problem. In this recommended configuration, all of your company's e-mail arrives via SMTP protocol and is scanned at the firewall.

If you configure the POP3 proxy to allow employees on the inside network to connect to POP3 servers on the outside network, you cannot use content scanning on their e-mail.

HTTP

Web pages can contain malicious content and applications. When you enable content scanning, the firewall scans Web pages before it displays them. It scans files transferred via the HTTP proxy, whether they were retrieved via HTTP or FTP URLs.

Note: Even though the Gopher proxy is an implementation of the HTTP proxy, the firewall does not scan Gopher traffic.

Content Scanning Engine

The firewall makes decisions on the data based on information from a content scanning engine running on another system. The types of content for which the firewall scans depends completely on the content scanning engine you are using. Common types of files that these engines can scan include the following:

- e-mail messages
- attachments to e-mail messages (for example, PostScript documents)
- executable files
- compressed or encoded files (for example, ZIP, UUENCODE, MIME)
- Java applets

Consult the documentation included with your content scanning engine to see what types of files the engine can scan.

How Content Scanning Works

Content scanning is another check that the firewall performs before passing packets from one side of the firewall to the other. When the firewall receives an SMTP, HTTP, or FTP packet, the appropriate proxy service performs the standard source and destination checks. If the proxy is configured to allow requests from the source to the destination, the proxy scans the files.

The proxies use the Content Vectoring Protocol (CVP) to communicate with the content-scanning engine. As it scans the file, the proxy uses the criteria you have set on the content-scanning engine to determine if the file is infected. If the file is not infected, the proxy passes the packet to the destination host just as it would have if content scanning were disabled. If the file is infected, the proxy checks its configuration settings and either discards the infected file or asks the scanning engine to attempt to repair the infected file.

When the proxy discards the infected file, it logs the error and drops the packet. The firewall does provide notification when it blocks a file. When blocking FTP traffic, the firewall provides status messages to the requesting client that notify the user that the firewall blocked the transfer. When blocking SMTP traffic, the firewall logs the attempt and sends e-mail to the recipient configured on the firewall to receive bad e-mail (by

default, root). When blocking HTTP traffic, the firewall provides information to the browser to display a page indicating that the firewall blocked the requested page.

If the proxy is configured to repair infected files, the firewall asks the content engine to provide a repaired file. If the content engine successfully repairs the file, it returns the repaired file. The firewall passes the file to the destination host. If the file is not infected and the proxy is not configured to repair infected files, the firewall passes the file to the destination host.

With the exception of status messages that the firewall displays when it blocks content, content scanning is transparent to the user. Users do not need to change their file transfer, mail, or browsing habits. However, users may see some small delays in both FTP and HTTP traffic, as the content scanning engine must also look at each packet.

Accessing Content Scanning Configuration

You can access content scanning configuration from each of the proxy services that support content scanning.

To access content scanning configuration for mail:

1. From within the Gauntlet Firewall Manager, select Environment.
2. Click the Mail tab.

The Mail window displays.

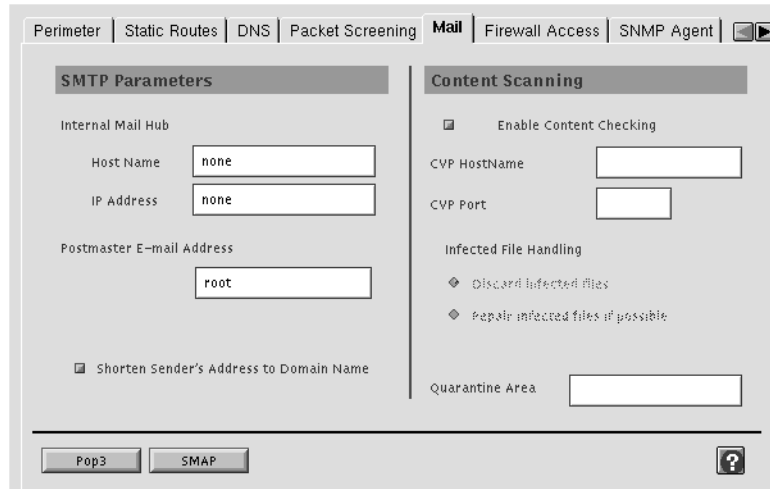


Figure 30-1 Mail Window

To access content scanning configuration for the FTP proxy service:

1. From within the Gauntlet Firewall Manager, select Services.
2. Click the FTP tab.
3. Select the configuration set for which you want to enable content scanning.
4. Click *Modify*.

To access content scanning configuration for the HTTP proxy service:

1. From within the Gauntlet Firewall Manager, select Services.
2. Click the HTTP tab.
3. Select the configuration set for which you want to enable content scanning.
4. Click *Modify*.

Configuring the Firewall to Use Content Scanning

Configuring the firewall consists of planning, configuring, and enabling content scanning.

Planning the Firewall Configuration for Content Scanning

When planning configuration for content scanning:

1. Determine whether you wish to configure content scanning for FTP, HTTP, or e-mail.
2. If you plan to hold infected e-mail files in a quarantine area, make sure you have sufficient disk space.

Configuring and Enabling Content Scanning

To configure and enable content scanning:

1. On the content scanning configuration page for the service, click *Enable Content Scanning*.

2. Provide information about content scanning.

CVP Hostname IP address or hostname of the system on which the content scanning engine is running.

CVP Port Port number on which the content scanning engine is running.

Infected File Handling Specifies what the firewall should do with files that the content scanner reports to be infected.

If you are using Norton AntiVirus for Firewalls, click *Discard Infected Files*. When Norton AntiVirus finds a virus type it cannot repair, it replaces the contents of the file with a warning message and marks the file as clean. The firewall believes that the file was repaired and delivers it to the user, and the logs on the firewall state that the transfer was repaired. In reality, the content scanning engine did not repair the file, it replaced the file.

Quarantine Area Directory where the firewall should place mail messages that the content scanner reports to be infected.

3. Enable the proxy service normally. Do not forget to add the service to a service group or create service group rules.
4. Before exiting the Gauntlet Firewall Manager, save and apply your changes. The firewall enables content scanning the next time you reboot your firewall.

Configuring the Content Scanner

The Gauntlet Firewall uses the Content Vectoring Protocol (CVP) to communicate with content scanning engines. Currently, Symantec Norton AntiVirus for Firewalls is supported.

To configure your firewall:

1. If at all possible, install the content scanning engine on a system that is on the same network segment as the firewall, either the trusted or service net. Do *not* place the scanning engine on the outside network.
2. When configuring information about the host that will provide content to be scanned (that is, the firewall), enter the IP address of the interface that is closest to the content scanning engine. For example, if the content scanner is on the inside network, enter the IP address of the inside interface of the firewall.
3. For each of the protocols that you are scanning, set the scan policy to:
 - scan all items
 - accept all items

This ensures that the content-scanning engine looks at all traffic it receives.

Managing URL Filtering

While there is a wealth of valuable information available on the Internet, not all of it is appropriate at all times and for all purposes. With the Gauntlet Firewall, you have the option of explicitly blocking or allowing specific URLs, and you can also use the Cyber Patrol filtering software from The Learning Company to block access to objectionable material.

The following sections explain the concepts behind URL filtering and Cyber Patrol, describe how they work, and tell you how to configure them.

- “Understanding URL Filtering” on page 293
- “Configuring URL Filtering” on page 294
- “Understanding Cyber Patrol” on page 297
- “Configuring Cyber Patrol” on page 299
- “Enabling Cyber Patrol Services” on page 304

Understanding URL Filtering

URL filtering works in two ways:

- By denying or allowing users inside the firewall access to specified URLs on the Internet
- By filtering out of URL headers specified xurl-encoded characters (usually used to filter information that you do not want the untrusted world to know about)

Both kinds of URL filtering are configured using the URL Filtering window of the HTTP proxy.

Denying and Allowing Access to Specified URLs

Using the URL Filter Configuration window, you can specify that certain URLs be denied to your users. If a user attempts to access a denied URL, they will receive a message indicating that access to this URL is denied. You can use this feature to deny access to any URL whether or not you are using Cyber Patrol.

If you are using Cyber Patrol, these denied sites are in addition to the sites denied by Cyber Patrol. You can also use the URL Filter Configuration window to allow access to URLs that would otherwise be blocked by Cyber Patrol.

Note: Any URL you permit or deny access to using URL filtering takes precedence over Cyber Patrol settings.

Filtering URL Headers

For security purposes, you may want to hide from the untrusted world certain xurl-encoded characters that are normally contained in outgoing URL headers.

Consult the HTML RFC or other HTML specification documents for more information about xurl-encoded characters.

Configuring URL Filtering

Configuring URL filtering for your Gauntlet Firewall involves planning, configuring the settings, and enabling the new settings.

Planning URL Filtering

When planning for URL filtering:

1. Determine which URLs you for which you wish to deny your users access, and which sites blocked by Cyber Patrol you wish to allow your users to access (if you are using Cyber Patrol).
2. Decide what xurl-encoded characters you wish to filter out of outgoing URL headers.

Configuring URL Filtering Settings

To configure URL filtering settings:

1. From within the Gauntlet Firewall Manager, select Services.
2. Click the HTTP tab.
The HTTP window displays.
3. Click the *Add* button.
The Add HTTP Services window displays.
4. Click the *URL Filtering* button.
The URL Filtering window displays.

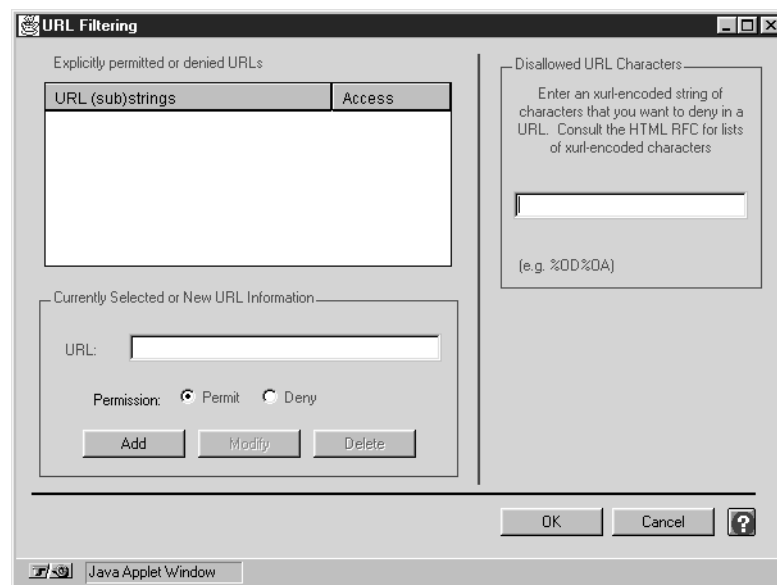


Figure 31-1 URL Filtering Window

To enter a URL to be denied or permitted:

1. Enter an entire URL or a substring of a URL in the URL field. For example:
 - .jpg matches all URLs with .jpg somewhere in the URL.
 - yoyodyne.com matches all URLs with yoyodyne.com in the URL.
2. Select Permit or Deny from the Permission options.
3. Click *Add*.

The specified URL displays on the list of permitted or denied URLs.

To change the settings of a denied or permitted URL:

1. Select the URL whose settings you wish to change from the list of URLs.
2. Make the desired changes.
3. Click *Modify*.

To delete a permitted or denied URL:

1. Select the URL you want to delete from the list of URLs.
2. Click *Delete*.

The URL disappears from the list of URLs.

To enter disallowed xurl-encoded characters:

1. Click in the Disallowed URL Characters field.
2. Enter an xurl-encoded string of characters.

Enabling URL Filtering Services

To enable URL filtering settings:

1. On the URL Filtering window, click *OK*.
2. Add the HTTP proxy to the appropriate service groups.
3. Before exiting the Gauntlet Firewall Manager, save and apply your changes.

The firewall enables URL Filtering.

Understanding Cyber Patrol

Cyber Patrol is filtering software from The Learning Company that blocks access to sites on the Internet. The CyberNOT database, a listing of Web sites, is divided into content categories (for example, Violence/Profanity, Intolerance, and so on).

If a category is active, your users will be denied access to all Web sites (URLs) that are part of that category—unless a specific URL that is part of the category is specifically designated as permitted using the URL Filter Configuration window.

You can make a different set of categories active for work time hours and leisure time hours. This allows you to make categories such as Sports & Entertainment or Search Engines available to your users during leisure time hours.

You can download and automatically install the most current version of the CyberNOT database at any time during the 30-day trial period or after you have obtained a valid Cyber Patrol license.

Note: To purchase a Cyber Patrol license, contact your local Network Associates sales representative or dealer, or call 1-888-FIREWALL.

The Cyber Patrol categories are:

- **Violence/Profanity**—Pictures or text exposing extreme cruelty, or physical or emotional acts against any animal or person that are primarily intended to hurt or inflict pain. It includes obscene words, phrases, and profanity, defined as text that uses, but is not limited to, George Carlin's seven censored words more often than once every 50 messages (newsgroups) or once a page (Web sites).
- **Partial Nudity & Art**—Pictures exposing the female breast or full exposure of either male or female buttocks except when exposing genitalia. Excludes all swimsuits, including thongs.
- **Full Nudity**—Pictures exposing any or all portions of the human genitalia. Excluded from the Partial Nudity and Full Nudity categories are sites containing nudity or partial nudity of a wholesome nature; for example, Web sites containing publications such as National Geographic or Smithsonian Magazine; or sites hosted by museums such as the Guggenheim, the Louvre, or the Museum of Modern Art.
- **Sexual Acts/Text**—Pictures or text exposing anyone or anything involved in explicit sexual acts and/or lewd and lascivious behavior, including masturbation, copulation, pedophilia, and intimacy involving nude or partially nude people in

heterosexual, bisexual, lesbian, or homosexual encounters. Also includes phone sex ads, dating services, adult personals, CD-ROMs, and videos.

- **Gross Depictions/Text**—Pictures or descriptive text of anyone or anything that are crudely vulgar or grossly deficient in civility or behavior or that show scatological impropriety. Includes such depictions as maiming, bloody figures, or indecent depiction of bodily functions.
- **Intolerance**—Pictures or text advocating prejudice or discrimination against any race, color, national origin, religion, disability or handicap, gender, or sexual orientation. Any picture or text that elevates one group over another. Also includes intolerant jokes or slurs.
- **Satanic or Cult**—Pictures or text advocating devil worship, an affinity for evil, or wickedness, or the advocacy to join a cult. A cult is defined as a closed society headed by a single individual where loyalty is demanded and leaving is punishable.
- **Drugs & Drug Culture**—Pictures or text advocating the illegal use of drugs for entertainment. Includes substances used for other than their primary purpose to alter the individual's state of mind, such as glue sniffing. Does not include currently illegal drugs legally prescribed for medicinal purposes (for example, drugs used to treat glaucoma or cancer).
- **Militant/Extremist**—Pictures or text advocating extremely aggressive and combative behaviors or advocacy of unlawful political measures. Topics include groups that advocate violence as a means to achieve their goals. Includes "how to" information on weapons making, ammunition making, or the making or use of pyrotechnics materials. Also includes the use of weapons for unlawful reasons.
- **Sex Education**—Pictures or text advocating the proper use of contraceptives. This topic would include condom use, the correct way to wear a condom, and how to put a condom in place. Also included are sites relating to discussion about the use of the Pill, IUDs, and other types of contraceptives. Also includes discussion sites on how to talk to your partner about diseases, pregnancy, and respecting boundaries. Excluded from this category are commercial sites wishing to sell sexual paraphernalia.
- **Gambling/Questionable/Illegal**—Pictures or text advocating materials or activities of a dubious nature that may be illegal in any or all jurisdictions, such as illegal business schemes, chain letters, copyright infringement, computer hacking, phreaking (using someone's phone lines without permission), and software piracy. Also includes text advocating gambling relating to lotteries, casinos, betting, numbers games, online sports, or financial betting, including non-monetary dares.

- **Sports & Entertainment**—Internet resources considered appropriate for a leisure setting, but which may be deemed irrelevant, unproductive, or inappropriate to a work environment. Includes content in the following categories: automotive, entertainment (television and radio stations and schedules, and music information such as bands, record companies, and concert schedules), finance (stock reports, investment companies, college financial aid), interactive games, software companies, hobbies (collecting, gardening), newspapers and magazines, real estate and apartment guides, shopping, sports, theater and movies, and vacation planning.
- **Alcohol, Beer, Wine & Tobacco**—Pictures or text advocating the sale, consumption, or production of alcoholic beverages and tobacco products.
- **Search Engines**—Commercial Web sites used primarily as Internet search engines (Yahoo, Lycos, Alta Vista, and Excite, for example). Like the Sports/Leisure category, these sites may be deemed irrelevant, unproductive, or inappropriate to a work environment.

Note: Any online content with more than three instances in 100 messages or any easily accessible pages with graphics or text that fall within the definition of a category puts the source into the category. The categories do not pertain to sites containing opinion or educational material, such as the historical use of marijuana or the circumstances surrounding 1940s anti-Semitic Germany.

Configuring Cyber Patrol

Configuring Cyber Patrol for your Gauntlet Firewall involves planning, configuring the settings, and enabling the new settings.

Planning

When planning your Cyber Patrol setup:

1. Determine whether you wish to use Cyber Patrol filtering.
2. Decide which CyberNOT categories you wish to be active and inactive.
3. Consider how you would like to define work time hours and leisure time hours.
4. Decide whether you want to purchase a Cyber Patrol license.
5. Determine whether or not you have the latest CyberNOT database (they are updated weekly).

Configuring Cyber Patrol Settings

To configure Cyber Patrol settings:

1. From within the Gauntlet Firewall Manager, select Services.
2. Click the HTTP tab.
The HTTP window displays.
3. Click *Add*.

The Add HTTP Services window displays.

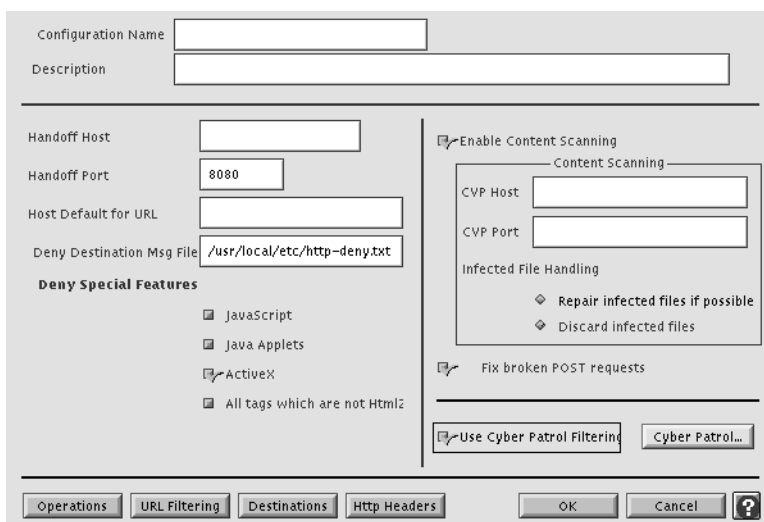


Figure 31-2 Add HTTP Services Window

4. Make sure there is a check mark in the Use Cyber Patrol Filtering box.
5. Click the *Cyber Patrol* button.
The Cyber Patrol Configuration window displays.

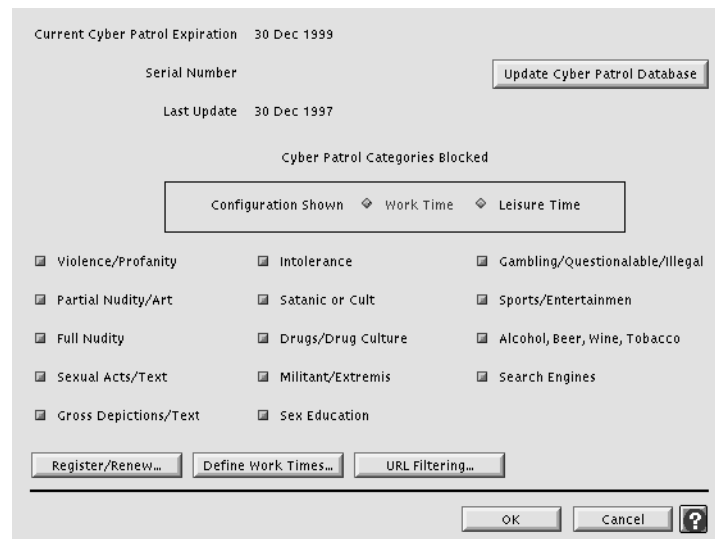


Figure 31-3 Cyber Patrol Configuration Window

To make CyberNOT categories active or inactive for Work or Leisure Time hours:

1. From the Configuration Shown options, click *Work Time* or *Leisure Time*.
2. Put a check mark next to the names of those categories you wish to be blocked for your users; remove the check mark from those categories you wish to be available to your users.

Note: You do not need to register your 30-day trial subscription; just start using it.

To register or renew your Cyber Patrol license:

1. On the Cyber Patrol Configuration window, click *Register/Renew*.
The Cyber Patrol Renewal/Registration window displays.

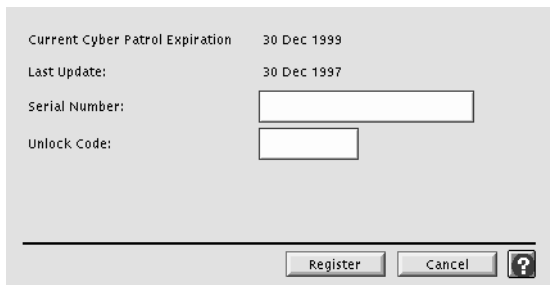


Figure 31-4 Cyber Patrol Renewal/Registration Window

2. To register your Cyber Patrol license, enter your serial number in the Serial Number field and click *Register*.
The registration information at the top of the window displays the updated information.
3. To renew your Cyber Patrol license, enter your unlock code in the Unlock Code field and click *Renew*.

Note: If your original serial number isn't in the Serial Number field, you must re-enter it.

A new CyberNOT database is automatically downloaded and installed, and the Cyber Patrol Configuration window displays.

To define work and leisure time hours:

1. In the Cyber Patrol Configuration window, click *Define Work Times*.

The Work Time Range Configuration window displays.

All times listed below are considered work time. Any time outside this set of times is considered leisure time. Enter new or modify existing work time range entries below.

Days Of The Week	Start Time	Stop Time

Currently selected or new work time range information

Start Time

End Time

Enter Time in form HHMM PM

Sunday Thursday
 Monday Friday
 Tuesday Saturday
 Wednesday

Figure 31-5 Work Time Range Configuration Window

2. Establish work time hours by entering start and end times, selecting the appropriate days, and clicking *Add*.

The time frame you established displays in the list. You do not need to configure leisure time hours; all hours outside of work time hours are considered leisure time hours.

3. Click *OK*.

The Cyber Patrol Configuration window redisplay.

To update the CyberNOT database with the current version:

1. In the Cyber Patrol Configuration window, click *Update Cyber Patrol Database*.

The current CyberNOT database is automatically downloaded and installed.

Enabling Cyber Patrol Services

To enable Cyber Patrol settings:

1. In the Cyber Patrol Configuration window, click *OK*.
2. Add the HTTP proxy to the appropriate service groups.
3. Before exiting the Gauntlet Firewall Manager, save and apply your changes.

The firewall enables Cyber Patrol services.

Managing the Network Management Agent

Network management and monitoring are crucial in today's increasingly complex and heterogeneous network environment. Swift detection and response to a failing mission-critical networked resource can prevent considerable financial losses. Most of today's network management platforms use the industry-standard Simple Network Management Protocol (SNMP) to communicate with resources being managed. The Gauntlet Firewall includes an SNMP agent that provides a limited amount of information to SNMP managers. The SNMP agent can be used to obtain information about the firewall such as its DNS name and the IP addresses of the firewall's interfaces.

Note: The Gauntlet Firewall also includes an SNMP proxy. Refer to Chapter 13, "Managing Network Management Services," on page 121 for more information on the SNMP proxy.

This chapter discusses the concepts behind the SNMP agent, explains how it works, how to configure and start the agent, and lists steps for configuring a selection of network management products to use the SNMP agent with Gauntlet Firewalls. The chapter consists of the following sections:

- "Understanding the SNMP Agent" on page 306
- "Accessing SNMP Agent Configuration" on page 307
- "How the SNMP Agent Works" on page 306
- "Configuring the Firewall as an SNMP Agent" on page 307
- "Configuring SNMP Network Managers" on page 309
- "Understanding SNMP Agent Replies" on page 309

Understanding the SNMP Agent

The SNMP agent is a daemon running on the firewall that accepts queries from SNMP managers. These queries must be SNMP version 1, defined by RFC 1157. You can configure the agent to allow queries from SMNP managers based on:

- source IP address
- source hostname

How the SNMP Agent Works

The firewall runs an instance of the SNMP agent (*snmpd*) as a daemon listening on a configurable port (by default, UDP port 1610). When the agent receives a query, it checks its configuration information and determines whether the initiating host has permission to access the agent. If the initiating host does not have permission to query the agent, the agent logs the attempt and drops the SNMP packet. If the host has permission, the agent responds using standard SNMP responses.

If the manager provides incorrect community information, the agent also logs the request and discards the packet. If the manager requests an unsupported MIB object, the agent provides the appropriate SNMP response.

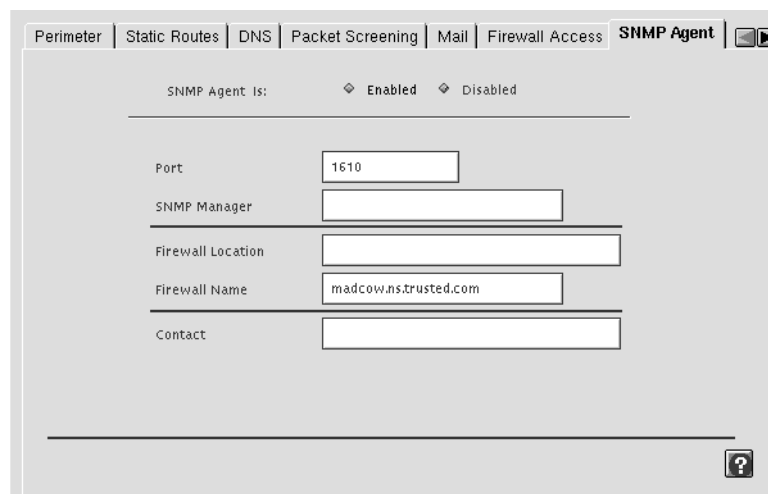
The agent does not listen on the standard SNMP port (UDP port 161) because the Gauntlet Firewall also includes an SNMP proxy that uses the SNMP port. If you are not using the SNMP proxy, you can run the SNMP agent on the SNMP port.

Accessing SNMP Agent Configuration

To access the SNMP agent configuration:

1. From within the Gauntlet Firewall Manager, select *Environment*.
2. Click the SNMP Agent tab.

The SNMP Agent window displays.



The screenshot shows the 'SNMP Agent' configuration window. At the top, there are several tabs: 'Perimeter', 'Static Routes', 'DNS', 'Packet Screening', 'Mail', 'Firewall Access', and 'SNMP Agent'. The 'SNMP Agent' tab is active. Below the tabs, the status is 'SNMP Agent Is: Enabled', with radio buttons for 'Enabled' and 'Disabled'. The 'Enabled' radio button is selected. Below the status, there are several input fields: 'Port' with the value '1610', 'SNMP Manager' (empty), 'Firewall Location' (empty), 'Firewall Name' with the value 'madcowns.trusted.com', and 'Contact' (empty). A help icon (question mark) is located in the bottom right corner of the window.

Figure 32-1 SNMP Agent Window

Configuring the Firewall as an SNMP Agent

Configuring the Gauntlet Firewall involves planning, configuring proxy settings, and enabling the agent.

Planning SNMP Agent Settings

When planning the SNMP agent settings:

1. Determine which SNMP manager should receive the cold start manager sent by the agent. Ensure that you have the IP address of that host.
2. Decide which SNMP managers can use the SNMP agent and create the appropriate network groups.

Configuring SNMP Agent Settings

To configure SNMP agent settings, provide information about the agent on the SNMP Agent window:

SNMP Manager IP address of the SNMP manager that can communicate with the agent on the Gauntlet Firewall.

Firewall Name Name of the firewall that the SNMP agent reports to the SNMP manager. By default, this is the host name of the firewall.

Note: Some network management software (such as HP OpenView and CA Unicenter TNG) use the *ping* program to verify that an agent is reachable before they send the actual SNMP request. Such management programs may also try to ping all interfaces on a system. By default, the firewall responds to pings of the directly connected interface, but does not respond to ping traffic for other interfaces. If your network management software uses the *ping* program to ping all interfaces, you may want to add packet screening rules to your firewall. These rules must allow ICMP between the network manager and the other interfaces of the firewall.

Enabling the SNMP Agent

To enable the SNMP agent:

1. In the SNMP Agent window, click *Enabled*.
2. Add the SNMP agent to the service group that controls access to the firewall for the network containing the SNMP network manager.
3. Before exiting the Gauntlet Firewall Manager, save and apply your changes.

The next time you reboot your firewall, the firewall enables the SNMP agent.

Configuring SNMP Network Managers

To configure your network manager, configure the network manager so that it looks for the agent on the firewall on UDP port 1610 instead of 161.

Understanding SNMP Agent Replies

The Gauntlet Firewall uses standard SNMP responses. Refer to your SNMP documentation for more detailed explanations. This section lists the types of responses the Gauntlet Firewall agent provides.

SNMP Agent and Management Information Base

The SNMP agent currently supports a small subset of the management information base (MIB) defined in RFC 1213 (MIB-II). The agent provides the following MIB values:

- all the mib-2.system values
- all the mib-2.snmp values
- mib-2.interfaces.ifNumber
- mib-2.ip.ipAddrTable values

All of the values that the agent supports are read-only values. An SNMP manager cannot set these values using SNMP. Some of the mib-2.system values are configurable through the Gauntlet Firewall Manager. The agent provides all of the supported values regardless of the community the agent is in.

SNMP Agent Community

The agent can be a member of the following communities:

- public
- private
- regional
- proxy
- core

SNMP Agent Object ID

The value provided for the mib-2.system.sysObjectID object is .1.3.6.1.4.1.602.1.1.1.5.1.

Trap

The Gauntlet Firewall sends only one trap, the cold start trap.

Login Shell

You need to log into the firewall occasionally to manage it. You may log in directly at the console or remotely via TELNET or Rlogin. Occasionally, you may also need to send files to or from the firewall using FTP. Whenever you access the firewall remotely, you are sending your password (and your root password) in the clear across your internal network to the firewall. While you would like to believe that this is secure, you want to be prudent.

One way of protecting your password is to log into the firewall using some form of strong authentication that uses one-time passwords or time-based responses. The login shell program included with the Gauntlet firewall allows you to use the same strong authentication scheme for logging into the firewall itself as you do for activity between opposite sides of your security perimeter.

This chapter discusses the concepts behind the login shell program and explains how it works, how to configure the program, and how to use it. The chapter consists of the following sections:

- “Understanding the Login Shell Program” on page 311
- “How the Login Shell Program Works” on page 312
- “Configuring the Firewall to Use the Login Shell Program” on page 312
- “Using the Login Shell Program” on page 315

Understanding the Login Shell Program

The login shell program is a wrapper program that authenticates the user (using strong authentication) before passing control to the real login shell. It provides authentication and logging.

How the Login Shell Program Works

A user logs into the firewall via the console, TELNET, or *rlogin*. This calls the standard login program to process the login. The login program asks for a user name. The login program reads the */etc/passwd* file and determines that this user does not require a password (because the password field is empty). It then passes the information to the program specified in the shell field, the login shell program (*/usr/etc/login-sh*).

The login shell program prompts the user for the appropriate response for their authentication method (S/Key password, and so on). The login shell program checks its configuration information (in the *netperm* table). It authenticates the user using the authentication server specified in the *netperm* table. The login shell program then reads the shell file configuration file (by default, */usr/local/etc/login-shellfile*). It passes the login information to the executable specified for the user in the shell configuration file, which is normally the user's shell. The user is now logged into the firewall and ready to work.

The standard FTP daemon does not use *login*, so it will not invoke the login shell program for authentication. This is not usually a problem, because running the standard FTP daemon on the firewall is strongly discouraged.

Configuring the Firewall to Use the Login Shell Program

Configuring the Gauntlet Firewall involves planning, enabling remote login, creating user accounts, configuring the proxy to enforce your policy, and securing other applications.

Planning Remote Login

When planning for remote login, determine which users you want to allow to access the firewall remotely.

Enabling Remote Login

You must configure the firewall to allow remote login from other hosts.

Use the Gauntlet administration tools to allow the firewall to run the TELNET or *rlogin* daemons.

Adding Support for the Login Shell

You must add support for the login shell so that the operating system recognizes the login shell as a valid shell.

To add support for the login shell, edit */etc/shells* and add a line indicating the location and path of the login shell:

```
/usr/etc/login-sh
```

Creating User Accounts

To create user accounts:

1. Create or modify the user account on the firewall.

```
scooter::518:10:Scooter Lindley:/home/scooter:/usr/etc/login-sh
```

- Leave the password empty, because the login shell uses your strong authentication information.
- If you include a password, you are prompted to authenticate twice: once for the information you enter here, and once for your strong authentication information.

2. Specify *login-sh* as the shell.
3. Create the user's home directory, if necessary:

```
mkdir /home/scooter
```

4. Add the user to the appropriate group by editing */etc/groups*.

Configuring the Proxy Rules

If you are running the Gauntlet Firewall default configuration, you do not need to modify configuration rules for the login shell. If you have chosen a different authentication server or a different location for your shell file information, you must modify */usr/local/etc/netperm-table* to reflect your configuration. Refer to the *Gauntlet Netperm Table Reference Guide* for more information about modifying the *netperm* table.

Configuring the Shell

You must provide information for each user indicating their final (real) shell. After the login shell authenticates the user, it starts the user's final shell.

To configure the shell, edit the *shellfile* file (*/usr/local/etc/login-shellfile*) and add information about the final shell for that user:

username executable parameters

username Same user name that you specified when you created the user's account on the firewall.

executable Name of the executable that the login program executes after authenticating the user. This is typically the user's shell.

parameters Parameters to the executable program. The first parameter is typically a hyphen (-) and the shell name (csh, ksh, and so on).

For example:

```
scooter /usr/bin/tcsh -tcsh
```

Creating User Authentication Records

To create the record in the authentication database:

- Use the authentication management system to create authentication user entries for all users who will use the login shell on the firewall.
- Use the same user name that you specified when you created the user on the firewall.

Securing Other Applications

To secure other applications:

1. Disable programs (such as *chsh*) that allow users to change their shells. Either remove the executable or change the file permissions to 700:

```
chmod 700 chsh
```

Create accounts on the firewall for only those people who need to administer the firewall. They will all have access to the root password. Changing file permissions will not prevent them from changing their shell. If you are creating accounts for other users on the firewall (which is not recommended), changing file permissions will prevent them from changing their shell.

2. Verify that the *su* command is not aliased to *su -m* in your account (*.cshrc*, *.login*, and so on) on the firewall. The **-m** option attempts to retain the current environment. This causes your login shell (in this case, *login-sh*) to be executed by user root. Because there is no entry for root in the *login-shellfile*, *su -m* does not work.

Verifying Your Setup

Verify your installation by TELNETing to the firewall and connecting to the firewall itself. Connect to the firewall directly:

```
tn-gw-> c localhost
```

After you enter your user name, you are prompted for your strong authentication information.

Using the Login Shell Program

Using the login shell program can mean accessing the firewall from trusted or untrusted networks.

Accessing the Firewall From Trusted Networks

Log into the firewall (via the console, TELNET, or *rlogin*) as you did before. After you enter your user name, you are prompted for the response or password specified for your authentication scheme. Become root (via *su*) to do work as needed.

Accessing the Firewall From Untrusted Networks

Connect to the firewall as you did before, providing your strong authentication information to connect to the proxy. If you log into the firewall itself, you will need to authenticate again.

Logging into the firewall as root from an untrusted network is strongly discouraged. When you log in as root, you are sending your password across the untrusted network, making it vulnerable to detection.

Changing Password for User Account

When you are using the login shell, the password is actually the strong authentication password, not the standard IRIX password.

Do not use the *passwd* or *chpass* programs on your IRIX system. To change your password, you must follow the instructions for changing your strong authentication information.

If you use the *passwd* or *chpass* programs, you will create an IRIX password. You will then need to provide both your IRIX password and your strong authentication information when you log into the firewall.

PART SIX

Appendices, Glossary, and Index

Appendix A

Installing and Upgrading to Gauntlet 4.1

Appendix B

Initializing Strong Authentication Tokens

Glossary

Index

Installing and Upgrading to Gauntlet 4.1

This appendix provides implementation-specific information on installing and upgrading to Gauntlet 4.1. This information supplements the Personal System Administration Guide and the Software Installation Administrator's Guide, which are both available online.

The appendix explains installation and upgrade in the following sections:

- “Gauntlet Execution Environment Subsystems” on page 320
- “Prerequisites for Installing Gauntlet” on page 321
- “Files In This Release” on page 324
- “Activating Your Gauntlet License” on page 325
- “Gauntlet Configuration” on page 324
- “Obtaining and Installing a Software License” on page 326
- “Upgrading to Gauntlet 4.1” on page 328

Note: If you are upgrading to Gauntlet 4.1 from a previous version of Gauntlet, you should back up your firewall and follow the same instructions as someone installing Gauntlet for the first time. Make sure you install the license shipped with Gauntlet 4.1. Your previous license will not work.

Gauntlet Execution Environment Subsystems

Gauntlet Execution Environment includes these subsystems:

Table A-1 Gauntlet Execution Environment Subsystems

Subsystem	Description
gauntlet_eoe.books.Gauntlet_AG	Gauntlet for IRIX Administrator's Guide.
gauntlet_eoe.books.Gauntlet_NG	Gauntlet for IRIX Netperm Table Reference Guide.
gauntlet_eoe.man.gauntlet	On-line manual pages for some Gauntlet components.
gauntlet_eoe.man.relnotes	Gauntlet release notes.
gauntlet_eoe.sw.ace	Software support for Security Dynamics ACE authentication server.
gauntlet_eoe.sw.apop	Software support for APOP (mail) user authentication.
gauntlet_eoe.sw.ccard	Software support for CRYPTOCARD user authentication.
gauntlet_eoe.sw.gauntlet	Base Gauntlet software, required.
gauntlet_eoe.sw.gui	Gauntlet Firewall Manager software, required.
gauntlet_eoe.sw.mdauth	Software support for MD5 user authentication.
gauntlet_eoe.sw.safeword	Software support for Secure Computing SafeWord authentication server.
gauntlet_eoe.sw.skey	Software support for S/Key user authentication.
gauntlet_eoe.sw.vasco	Software support for VASCO Access Key II user authentication.
gauntlet_eoe.sw.radius	Software support for RADIUS authentication.
gauntlet_eoe.sw.digipass	Software support for Digipass authentication.

The Gauntlet distribution media also includes and some patches that are necessary for Gauntlet operation. The U.S. Domestic version also includes the gauntlet_encrypt subsystems. The patches and the gauntlet_encrypt subsystems have their own release notes.

Gauntlet Disk Space Requirements

This section lists the subsystems of the Gauntlet Execution Environment and their sizes.

Note: The listed subsystem sizes are approximate. See the IRIS Software Installation Guide for information on finding exact sizes.

Table A-2 Gauntlet Subsystem Sizes

Subsystem Name	Subsystem Size (kilobytes)
gauntlet_eoe.man.gauntlet	256
gauntlet_eoe.sw.ace	300
gauntlet_eoe.sw.apop	292
gauntlet_eoe.sw.ccard	452
gauntlet_eoe.sw.gauntlet	6828
gauntlet_eoe.sw.gui	4064
gauntlet_eoe.sw.mdauth	324
gauntlet_eoe.sw.safeword	168
gauntlet_eoe.sw.skey	500
gauntlet_eoe.sw.vasco	440

Prerequisites for Installing Gauntlet

Users of Gauntlet for IRIX have to have a software license and meet certain software prerequisites listed in this section.

Software License

Before you start to install Gauntlet 4.1, obtain a software license for the product. See “Activating Your Gauntlet License” on page 325 for more information.

IRIX Software Prerequisites

Your Silicon Graphics system must be running IRIX release 6.2, 6.3, 6.4, or 6.5 and various software subsystems that are part of the IRIX system must be installed on your system before you can run Gauntlet software.

To check your IRIX release level:

1. From the Desktop toolchest, choose Desktop> Unix Shell.
2. Type: `uname -r`
This returns the operating system version.
3. To check whether the `ipgate` and `named` subsystems are installed type:
`versions eoe.sw.ipgate eoe.sw.named`

If these subsystems are installed, you will see these lines:

```
I eoe          <date>  IRIX Execution Environment, 6.2
I eoe.sw       <date>  IRIX Execution Environment Software
I eoe.sw.ipgate <date>  IP Network Gateway Support
I eoe.sw.named <date>  Berkeley Internet Name Domain Server
```

(Notice the "I"(installed) at the beginning of each line.)

If these subsystems are not installed, follow the IRIX installation instructions to install them on your system before continuing.

Installing the Software

All Gauntlet subsystems can be installed while your system is running. These installation instructions give an overview of how to install the software from a CD-ROM drive that is connected directly to your system. To install the software from a remote CD-ROM drive, or for more detailed information, see the online InSight administrator manuals.

To install the Gauntlet Execution Environment, follow these steps:

1. Log in as root, as follows.
 - In a shell window, type `login root`
 - Provide the root password if required.

2. Insert the Gauntlet 4.1 CD-ROM into your CD-ROM drive.
3. In the shell window, type:

```
inst -f /CDROM/dist
```

4. At the inst> prompt, type:

```
inst> install default
inst> install patch*
inst> go
```

You may see messages from inst(1M) that indicate that it cannot proceed because of conflicts. Some patch files included on the Gauntlet 4.1 CD-ROM may include updates to IRIX software subsystems which you do not have installed. You can safely resolve such conflicts by choosing to not install that part of that patch using the conflicts command.

Below is an example of the kind of conflict which can safely be resolved. Note that the conflict in this example is that the base subsystem is not installed:

```
Patch patchSG0000639.eoel_sw.svr4net does not have base subsystem
eoel_sw.svr4net version 1021572036 to 1029999900 installed
1a. Do not install patchSG0000639.eoel_sw.svr4net (1029999906)
```

To resolve the above conflict, you would type at the inst> command line:

```
inst> conflicts 1a
```

See the *Software Installation Administrator's Guide* for help with other conflicts. After all conflicts are resolved, type **go** at the >inst prompt to try again.

The software is installed when you see this message:

```
Installation and/or removal succeeded. You can insert another tape
or CD-ROM now. Type "quit" if you are ready to leave the
installation tool.
```

5. At the inst> prompt, type:

```
inst> quit
```

You may see some exit messages.

6. Reboot your system.

The new software will be used.

7. After you've installed the software, you have to configure it. See "Configuring the Firewall" on page 19 for more information.

Files In This Release

To find out exactly which files were installed, use the `showfiles(1M)` command. To see all the files that were installed as part of the Gauntlet Execution Environment, type:

```
showfiles gauntlet_eoe
```

If you would like to see files of specific subsystems, for example, what files were installed for S/Key support, call `showfiles(1M)` with the subsystem, as follows:

```
showfiles gauntlet_eoe.sw.skey
```

See Table A-1 for a list of all subsystems.

To find out which files would be installed before installing the Gauntlet Execution Environment subsystems, follow these steps:

1. In a shell window, type

```
# inst -f /CDROM/dist
```

2. At the `inst>` prompt, type:

```
inst> admin At the admin> prompt, type:
```

```
admin > files gauntlet_eoe
```

3. To exit

- At the `admin>` prompt, type **return**.
- At the `inst>` prompt, type **quit**

Gauntlet Configuration

You must configure Gauntlet before it works correctly. Gauntlet provides a text-based interface to help you configure your firewall for management via the Web. To get started, run the `/usr/local/etc/gauntlet-admin` command. When you finish setting up your firewall, reboot your machine. You can continue configurations using the Web-based Gauntlet Firewall Manager located at `http://your_host:21000/auth/gui.html`.

“Upgrading to Gauntlet 4.1” on page 328 explains how to use the upgrade tool to automatically change most configuration information.

Activating Your Gauntlet License

Your copy of Gauntlet for IRIX requires a software license in order to operate. This chapter defines some important software licensing terms, describes the basic procedure for obtaining, installing, and testing a software license for Gauntlet for IRIX, and lists addresses and phone numbers for contacting Silicon Graphics License Administration.

For more information about software licenses, see the *FLEXlm User's Guide*, which provides detailed information on using and administering software licenses. It is included in the system software documentation; the online version is in the subsystem `license_eoe.books.FlexLM_UG`.

A Few Definitions

A software license is a collection of information that, after installation, allows you to use a licensed software product on one or more systems. Software license information includes license type, license expiration date, a license password, system hostname, and host ID number (`lmhostid`), and additional information concerning the license and licensed software.

You can find the host name using the command `/sbin/uname -n`, and the host ID number using the command `/usr/sbin/lmhostid`. The license must be installed on the system that has the host name included in the software license information. If the host ID in the license is "ANY", the software license can be installed on any system.

There are two types of software licenses, node-locked and concurrent:

- *node-locked*—A node-locked license is installed on a particular system (node) and allows the licensed software to run on that system.
- *concurrent*—A concurrent license allows the licensed software to run on one or more systems in the network simultaneously. The number of systems (nodes) allowed is included in the software license information. The system on which a concurrent license is installed must be configured as a license server. (See Chapter 1 of the *FLEXlm User's Guide* for more information about license servers.)

There are two durations of software licenses, temporary and permanent:

- *temporary*—A temporary license allows you to use the licensed software for a short period, typically a week to several months. The expiration date tells you the last date on which you can use the licensed software. Temporary licenses are often shipped with software so that you can use the software before a permanent license is issued.
- *permanent*—A permanent license allows you to use this release of the licensed software for a long time. Permanent licenses are issued only for software that has been purchased.

Obtaining and Installing a Software License

To obtain and install a software license, follow these steps:

1. Check whether you have received software license information.

Software license information is distributed in several ways: on labels attached to a Software License Registration card, on sheets of paper included with the product, or by mail, FAX, electronic mail, or via the World Wide Web.

2. Determine if you need to install a software license.

You may or may not need to install a software license for Gauntlet for IRIX:

- If you are updating your system to this release of Gauntlet for IRIX, you need to install a new license at this time.
- If you have received both a temporary license and a permanent software license, install the permanent license; do not install the temporary license.
- If you have received a permanent license, you should install it because it enables you to use the software that you have purchased.
- If concurrent licenses are used at your site and you plan to use an already-installed license, you can install and use the licensed software on your system without installing a license.
- If you have received a temporary software license but do not need to use the software immediately, you may choose to wait to install a license until you obtain the permanent license.

3. Request a software license if you don't have a software license at all, or if you have a temporary license and need a permanent license.

To obtain a software license, fill out the Software License Registration card that was included with the software (or the replica in the FLEXlm User's Guide). Send the information on the card by electronic mail (preferred), FAX, or mail to Silicon Graphics License Administration or your local service provider. After your request is received by Silicon Graphics or your local service provider, you should receive a software license within two business days.

4. Identify the system on which you will install the software license.

Because software license information usually must be installed on a particular system, follow these guidelines:

- Use the `/sbin/uname -n` command to identify the system on which the license is intended to be used.
- If the sysinfo is "ANY," you can install the license on any system you choose.
- If the host name included with the software license information doesn't match the host name of the system on which you want to install the license, contact Silicon Graphics License Administration.

5. Install the software license (temporary or permanent).

Check the license type listed in the software license information to find out whether the license is a node-locked license or a concurrent license. The installation procedure depends on the license type:

- If you are installing a node-locked license, use the LicenseManager(1M) tool. You can bring up the tool by choosing "License Manager" from the System toolchest.
- If the license is a concurrent license, you may need to configure the system on which you plan to install the license as a license server. (See Chapter 1 of the FLEXlm User's Guide for more information about license servers.)

6. Verify that the software license has been successfully installed.

If the software license is not working, running the command `/etc/init.d/gauntlet start` will result in a number of warnings about missing licenses.

Note: If you installed a temporary license and you are entitled to a permanent license, replace the temporary license with a permanent license as soon as possible to ensure uninterrupted use of Gauntlet for IRIX.

Upgrading to Gauntlet 4.1

If you are upgrading to Gauntlet 4.1 from a previous version of Gauntlet, you can for the most part, follow the same installation instructions as a new user (see “Installing the Software” on page 322.)

However, the file formats used by Gauntlet 4.1 are in some cases quite different from those used by Gauntlet 3.2. This means that your old configuration files will not work with your new version of Gauntlet.

Gauntlet 4.1 for IRIX provides an upgrade program, */usr/local/etc/gauntlet-upgrade*, which translates your old configuration files into the new formats wherever possible. This helps you avoid reentering all your configuration data,

Upgrade Instructions

To use the upgrade program:

1. Log into your Gauntlet firewall as root.
2. Enter the command:

```
/usr/local/etc/gauntlet-upgrade
```
3. The *gauntlet-upgrade* program will inform you of its progress as it translates your files, concluding with the word “Done”.
4. You are now ready to set up your firewall using the text-based *gauntlet-admin* interface, as described in previous chapters.

How the Upgrade Program Works

The *gauntlet-upgrade* programs performs the following functions:

- Translates data stored by the Gauntlet 3.2 administrative interface to formats readable by the Gauntlet 4.1 Java-based Firewall Manager.
Variables previously stored in */usr/gauntlet/cgi-data/*.g* are now stored in */usr/local/etc/mgmt/gauntlet.conf*.
- Copies files which are in new locations under Gauntlet 4.1, including the authorization database, info server database and integrity verification files.

- Translates Gauntlet 3.2 trusted and untrusted network configurations to Gauntlet 4.1 format.
- Translates Gauntlet 3.2 explicit routing setup to Gauntlet 4.1 static routes.
- Converts swIPe and PC Extender configuration files.
- Converts plug gateway configuration files.

Upgrade Considerations

Keep in mind the following issues when using the *gauntlet-upgrade* program:

- Run *gauntlet-upgrade* immediately after you have installed Gauntlet 4.1, but before you have configured your firewall using *gauntlet-admin*. Changes made by *gauntlet-admin* would be lost during the upgrade process.
- The *gauntlet-upgrade* program is not designed to be run more than once.
- *gauntlet-upgrade* attempts to upgrade all aspects of your firewall which are configurable through the Gauntlet 3.2 administrative interface.
- *gauntlet-upgrade* replicates your Gauntlet 3.2 configuration whenever possible. However, in some cases an automatic upgrade may not be feasible. When using the Gauntlet 4.1 Firewall Manager GUI for the first time, carefully check each aspect of your configuration and make updates where needed.
- Manual changes to the Netperm table are not automatically upgraded; you have to make these changes through the Gauntlet 4.1 Firewall Manager or by manually editing */usr/local/etc/netperm-table*.
- Manual changes to ipfilterd rules are not automatically upgraded; you have to copy your changes to */usr/local/etc/mgmt/template.ipfilterd.conf*.

Initializing Strong Authentication Tokens

The Gauntlet Firewall supports several different strong authentication systems. The steps to create users is slightly different for each system.

The following sections explain how to configure the authentication systems supported by the Gauntlet Firewall:

- “Access Key II Authentication” on page 331
- “Digipass Authentication” on page 339
- “SafeWord Authentication Server” on page 342
- “SecurID System Authentication” on page 346
- “S/Key System” on page 351
- “RADIUS Authentication” on page 356
- “Reusable Passwords” on page 360

This appendix also explains how to configure your Gauntlet Firewall to use these systems. Refer to Chapter 6, “Users and User Groups,” on page 57 for an explanation of how the Gauntlet Firewall user authentication system works.

Access Key II Authentication

Your Access Key system includes a set of hand-held authenticators that you can use for authenticating to the Gauntlet Firewall.

How Access Key II Authentication Works

The Access Key II system, from VASCO Data Security, uses a random challenge password. When the firewall prompts for authentication, Access Key II provides a challenge. The user enters a PIN (if one is required) and the challenge into the Access Key II. The Access Key II responds with a password. The user enters this value at the Gauntlet prompt, and the Gauntlet authentication server verifies this value.

Configuring the Access Key II

To configure the Access Key II:

1. Create a key for the Access Key II according to the documentation included with the key.
2. This creates a file (*keyfile.log*) that contains the key. Place this file into a location accessible from the firewall.

Adding an Access Key II User

To add users, log into the firewall console or use TELNET to log in remotely.

To add an Access Key II user:

1. Log into the firewall and become root.
2. Copy the Access Key II keyfile (*keyfile.log*) to a temporary directory (such as */tmp/vasco*) on the firewall.
3. Load the key information into the user authentication management system using the key initialization tool (*/usr/etc/vasco_init*):

```
# cd /usr/etc
# ./vasco_init /tmp/vasco/keyfile.log
```

This tool creates a user in the authentication management system and loads the key for this user. It creates the user name by prepending the letter *i* to the serial number for that Access Key II. The user is initially disabled.

The key initialization tool reads only the first record in the file. If you need to create multiple Access Key II users, consider writing a script to create individual key files and run the key initialization tool.

4. Make a note of the user name that the initialization program displays so you can change it to something easier for the user to remember.

```
Record loaded for user: i2-0005899-4
```

5. Use the Authentication Manager tool to change the name of the user to something easier to remember:

```
# /usr/etc/authmgr
authmgr-> rename i2-0005899-4 jnolan 'John Nolan'
```

6. Enable the new user:

```
authmgr-> enable jnolan
```

7. Make the information active by exiting the authentication manager.
8. Provide the user with their Access Key II and user name.

Using Access Key II With the Gauntlet Firewall

This section first lists the steps for authenticating with Access Key II, then provides an example.

Authenticating With Access Key II

To authenticate to the Gauntlet Firewall using Access Key II:

1. Access a proxy that requires authentication:

```
telnet fire-out.yoyodyne.com
Trying 204.255.154.100...
Connected to fire-out.yoyodyne.com
Escape character is '^|'
```

2. Enter your user name:

```
Username: jnolan
```

3. Note the challenge the proxy displays:

```
Challenge 9683-5263:
```

4. On the Access Key II, press **p**. Note that the Access Key II displays 0000.

5. On the Access Key II, enter your PIN to enable the Access Key II.

6. On the Access Key II, enter the challenge that the proxy displays (without the dash). Note that the proxy displays a response:

```
9683-5263
```

7. At the response prompt, on your keyboard enter the response that the Access Key II displays (with or without the dash) and press **Enter**.

```
Challenge: 9683-5263 eh5ce3
```

Access Key II Example

This example shows a sample TELNET session from a system outside the firewall to a system inside the firewall.

```
blaze.clientsite.com-28: telnet fire-out.yoyodyne.com
Trying 204.255.154.100...
Connected to fire-out.yoyodyne.com
Escape character is '^]'.
Username: jnolan
Challenge: 9683-5263 eh5ce3
Login Accepted
fire-out.yoyodyne.com telnet proxy (Version 4.0a) ready:
tn-gw> c dimension
Trying 10.0.1.120 port 23...
Connected to dimension.yoyodyne.com
IRIX 6.5 (dimension) (ttyp5)
login: jnolan
Password: guess#me$now (password does not display)
Welcome to dimension.yoyodyne.com
3:57PM up 16 days, 5:35, 4 users, load averages: 0.03, 0.01, 0.00
dimension-26:
```

CRYPTOCARD RB-I

Your CRYPTOCARD RB-I system includes a set of hand-held authenticators (called CRYPTOCARDS) that you can use for authenticating to the Gauntlet Firewall.

How It Works

The CRYPTOCARD system uses a shared-secret key. You enter the key into the CRYPTOCARD when you initialize it, and into the Gauntlet authorization database when you configure that user. When the proxies want to authenticate, they ask the authentication server. The authentication server uses the shared secret to generate a challenge, which the proxy displays. You enter the challenge into your CRYPTOCARD, which uses the shared secret to create a response. You enter the response at the prompt. The proxy passes it back to the authentication server, which compares what you entered with what it expected and allows or denies access.

Configuring the CRYPTOCARD

Configuring the CRYPTOCARD involves generating a shared secret and initializing the CRYPTOCARD.

Generating the Shared Secret

To generate the shared secret:

1. Log in to the firewall and become root.
2. Run the key program that generates random keys for CRYPTOCARDS (*/usr/local/etc/ccardkey*):

```
# cd /usr/local/etc  
# ./ccardkey
```
3. Enter a random seed string and press **Enter**, as in this example:
Enter a line of random text as a seed: **alid oe02 I -01 [2qppdk 9**
Use any random set of nonsense words or characters.
4. Make note of the set of eight three-character groups (the shared secret) and the checksum that the program displays (you need these to configure the CRYPTOCARD and the authentication database), as in this example:
Enter into CRYPTOCARD: **044 346 000 315 035 171 045 011**
Checksum: **412-7559**

Initializing the CRYPTOCARD

To initialize the CRYPTOCARD:

1. Turn on the brand new, unprogrammed unit.
2. Within half a second enter **225371** and press **Enter**. The display shows **locked**.
3. Press **Enter**. The display shows **Options?**
4. Enter **111** and press the right arrow key to set the Pin Entry Feedback to on, Decimal Display to on, and Telephone Display to on (consult your CRYPTOCARD documentation for other options).
5. Enter **003** and press the right arrow key to set the UserID to none, Tries to unlimited, and Minimum PIN Length to 3 (consult your CRYPTOCARD documentation for other options).

6. Enter **001** and press the right arrow key to set the TimeOut Length to 30 seconds, Language choice to English, and Number of Keys to one (consult your CRYPTOCARD documentation for other options).
7. Press **Enter**. The display shows `KEY1?`
8. Enter the first group of three characters from the shared secret and press the right arrow key.
9. Enter the second through eighth set of characters, pressing the right arrow key after each set.
The display goes blank after the last set.
10. Press **Enter**. The display shows a number (the checksum).
11. Compare the checksum the CRYPTOCARD created with the one the key initialization program on the firewall created.
 - If the checksums match, press **Enter**.
 - If the checksums *do not* match, press **Clear** (note that display shows `KEY1?`) and then reenter the shared secret. The display shows `NEW PIN?`
12. Enter a three- to eight-digit PIN and press **Enter**. The display shows `Verify`.
13. Enter the three- to eight-digit PIN again and press **Enter**. The display shows `Card OK`.
Make note of the PIN you entered so that the user can change the PIN once you give them the unit.

Adding a CRYPTOCARD User

Use the Gauntlet Firewall Manager to add users.

To add a CRYPTOCARD user:

1. From the Users tab, click Add.
The Add User window displays.
2. Provide information about the user.

UserID	User name. Remember that this user ID does <i>not</i> need to match any other user IDs for this user.
Name	Descriptive information about this user.

Status	Select Enabled to activate the account.
Group	If you want to make this user a member of a group, select the name of the group.
Authentication Method	Select CRYPTOCARD as the authentication method.
Password	Eight three-character groups (shared secret), including spaces, you used for the CRYPTOCARD:Passwd: 044 346 000 315 035 171 045 011The display uses XXXs to hide the password
Verification	The same eight three-character groups, including spaces, that you entered as the password.

3. Click OK to make your changes take effect.
4. Provide the user with their user name.

Using the CRYPTOCARD with the Gauntlet Firewall

This section discusses using the CRYPTOCARD in the following sections:

- “Using the CRYPTOCARD for the First Time” on page 337
- “Authenticating With CRYPTOCARD” on page 338
- “Changing Your CRYPTOCARD PIN” on page 338
- “CRYPTOCARD Example” on page 339

Using the CRYPTOCARD for the First Time

To use the CRYPTOCARD for the first time:

1. Turn on the CRYPTOCARD. Note that the display shows PIN?
2. Enter your PIN and press **Enter**. Note that the display shows New Pin?
3. Enter a new three to eight digit PIN and press **Enter**. Note that the display shows Verify.
4. Enter the new three to eight digit PIN again and press **Enter**. Note that the display shows Ready.

Authenticating With CRYPTOCARD

To authenticate to the Gauntlet firewall using a CRYPTOCARD:

1. Access a proxy that requires authentication.

```
telnet fire-out.yoyodyne.com
Trying 204.255.154.100...
Connected to fire-out.yoyodyne.com
Escape character is '^]'
```
2. Enter your user name.

```
Username: jbigboot
```
3. Note the challenge the proxy displays:

```
Challenge 817-8618:
```
4. Turn the CRYPTOCARD on.
5. On the CRYPTOCARD, enter your PIN and press **Enter**. Note that the display shows Ready.
6. On the CRYPTOCARD, press **Enter** then enter the challenge the proxy displays (without the dash) and press **Enter**. Note that the proxy displays a response:

```
195-3454
```
7. At the response prompt, enter the response the CRYPTOCARD displays (with or without the dash) and press **Enter**.

```
Challenge 817-8618: 195-3454
```

Changing Your CRYPTOCARD PIN

Consult your CRYPTOCARD documentation for information on changing the PIN on a CRYPTOCARD. You do not need to make any changes to the Gauntlet Firewall authentication database when you change the PIN on an CRYPTOCARD.

CRYPTOCARD Example

This example shows a sample TELNET session from a system outside the firewall to a system inside the firewall.

```
blaze.clientsite.com-28: telnet fire-out.yoyodyne.com
Trying 204.255.154.100...
Connected to fire-out.yoyodyne.com
Escape character is '^]'.
Username: jbigboot
Challenge: 817-8618: 195-3454
Login Accepted
fire-out.yoyodyne.com telnet proxy (Version 4.0a) ready:
tn-gw> c dimension
Trying 10.0.1.120 port 23...
Connected to dimension.yoyodyne.com
IRIX 6.5 (dimension) (ttyp5)
login: jbigboot
Password: guess#me$now (password does not display)
Welcome to dimension.yoyodyne.com
3:57PM up 16 days, 5:35, 4 users, load averages: 0.03, 0.01, 0.00
dimension-26:
```

Digipass Authentication

Your Digipass system includes a set of hand-held authenticators (called Digipasses) that you can use for authenticating to the Gauntlet Firewall.

How Digipass Authentication Works

The Digipass system uses a time-based password. The Digipass card generates a passcode. You enter a secret password into the Gauntlet user authentication system when you configure that user. When the firewall prompts for authentication, the user selects the appropriate authentication application and enters a PIN on the card. The card displays a passcode, which the user enters at the Gauntlet prompt. The Gauntlet authentication server verifies this value to allow or deny access.

Configuring the Digipass

To configure the Digipass:

1. Create a key for the Digipass according to the documentation included with the key. This creates a file (*cinit_a.dgp*) that contains the token secret for that key. It also creates a 14-digit decryption key.
2. Place this file in a location accessible from the firewall.

Adding a Digipass User

To add users, log into the firewall console or use TELNET to log in remotely.

To add a Digipass user:

1. Log into the firewall and become root.
2. Copy the Digipass key file (*cinit_a.dgp*) to the same directory on the firewall in which the initialization tool is located (*/usr/etc*).
3. View the key file and make note of the 14-digit encryption key.
4. Load the key information into the user authentication management system using the key initialization tool (*/usr/etc/digi_init*) and the 14-digit decryption key:

```
# cd /usr/etc
# ./digi_init 12345678901234
```

This tool creates a user in the authentication management system and loads the key for this user. It creates the user name by prepending the letter *i* to the serial number for Digipass. The user is initially disabled.

Note: The key initialization tool reads only the first record in the file. If you need to create multiple Digipass users, consider writing a script to create individual key files and run the key initialization tool.

5. This tool also creates a PIN file (*ipin.txt*), which lists the user names and PINs for each Digipass. View the PIN file and note the user names in this file so that you can change it to something easier for the user to remember.

Token	PIN
1000000	1234
1000001	6789

6. Use the Authentication Manager tool to change the name of the user to something easier to remember:

```
# /usr/etc/authmgr
authmgr-> rename i1000000 jgant 'John Gant'
```

7. Enable the new user:

```
authmgr-> enable jgant
```

8. Make the information active by exiting the authentication manager.
9. Provide the user with their Digipass and user name, and the PIN listed in the PIN file for that Digipass.

Using Digipass With the Gauntlet Firewall

This section first lists the steps for authenticating with Digipass, then provides an example.

Authenticating With Digipass

To authenticate to the Gauntlet Firewall using Digipass:

1. Access a proxy that requires authentication:

```
telnet fire-out.yoyodyne.com
Trying 204.255.154.100...
Connected to fire-out.yoyodyne.com
Escape character is '^]'
```

2. Enter your user name:

```
Username: jgant
```

3. Turn the Digipass on.

4. On the Digipass, press **1**. The display shows `PIN?`

5. On the Digipass, enter your PIN and press **=**.

Your PIN displays as asterisks (*). The display shows “...” as it performs the calculation.

6. At the response prompt, enter the response that the Digipass displays and press **Enter**.

```
Code: 0190302588
```

Digipass Example

This example shows a sample TELNET session from a system outside the firewall to a system inside the firewall:

```
blaze.clientsite.com-28: telnet fire-out.yoyodyne.com
Trying 204.255.154.100...
Connected to fire-out.yoyodyne.com
Escape character is '^]'.
Username: jgant
Code: 0190302588
Login Accepted
fire-out.yoyodyne.com telnet proxy (Version 4.0a) ready:
tn-gw> c dimension
Trying 10.0.1.120 port 23...
Connected to dimension.yoyodyne.com
IRIX 6.5 (dimension) (ttyp5)
login: jgant
Password: guess#me$now (password does not display)
Welcome to dimension.yoyodyne.com
3:57PM up 16 days, 5:35, 4 users, load averages: 0.03, 0.01, 0.00
dimension-26:
```

SafeWord Authentication Server

The SafeWord Authentication Server is compatible with a group of hand-held authenticators that you can use for authenticating to the Gauntlet Firewall.

Note: SafeWord authentication is supported only by Gauntlet Firewalls running on Solaris systems.

SafeWord is compatible with:

- ActivCard
- CryptoCard
- DES Gold card
- DES Silver card
- Digipass (outside the U. S. only)
- SafeWord AccessCard

- SafeWord MultiSync
- SecureNet Key
- Softoken
- WatchWord

How SafeWord Authentication Works

This system, from Enigma Logic, provides an interface to the SafeWord Authentication Server for Gauntlet authentication. The Gauntlet authentication server uses the authentication information registered for a user with the SafeWord Authentication Server.

Configuring the SafeWord Authentication Server

To configure the SafeWord Authentication Server, create user accounts for your users on the SafeWord Authentication Server.

Configuring the Firewall for SafeWord Authentication

You must modify one file on the firewall so it knows where the SafeWord Authentication Server is (typically on a system other than the firewall).

To configure your firewall for use with a SafeWord Authentication Server:

1. Log into the firewall and become root.
2. Edit the SafeWord configuration file (*/usr/local/etc/mgmt/swec.cfg*). Set the SafeWord Authen. Server Name to the name of the system on which the SafeWord Authentication Server is running:

```
02 SafeWord Authen. Server Name:dimension 0 0 7482
```

Adding a SafeWord User

You can create two types of users in the Gauntlet Firewall for use with SafeWord Authentication Server: individual users or a default user. Individual users are unique user names for each user in your SafeWord Authentication Server system.

Creating a default user allows you to authenticate users without manually creating entries for every user in the Gauntlet authentication database. When a user logs in and the authentication server does not find the information in the Gauntlet authentication database, the authentication server sends the user information to the SafeWord Authentication Server. The authentication server also creates a record for that user in the Gauntlet authentication database.

For example, suppose you create an account for jparrot on your SafeWord Authentication Server. You create an account for default on the Gauntlet Firewall. When Jamie Parrot authenticates, she still uses the user name jparrot, and the authentication server sends the information to the SafeWord Authentication Server.

To add a user using SafeWord:

1. From the Users tab, click *Add*.

The Add User window displays.

2. Provide information about the user.

UserID	Enter the user name. Remember that this user ID does <i>not</i> need to match any other user IDs for this user.
Name	Enter descriptive information about this user.
Status	Click enable to activate the account.
Group	If you want to make this user a member of a group, select the name of the group.
Authentication Method	Select SafeWord as the authentication method.
Password	Leave this field blank. The Gauntlet authentication system uses the value registered with the SafeWord Authentication Server.
Verification	Leave this field blank.

3. Click *OK* to make your changes take effect.
4. Provide the user with the selected token and user ID.

Using SafeWord Authentication Server With the Gauntlet Firewall

This section first lists the steps for authenticating with SafeWord, then provides an example.

Authenticating With SafeWord Authentication Server

To authenticate to the Gauntlet Firewall using SafeWord Authentication Server:

1. Access a proxy that requires authentication.

```
telnet fire-out.yoyodyne.com
Trying 204.255.154.100...
Connected to fire-out.yoyodyne.com
Escape character is '^]'
- - Safeword@ Security Check V4.30.000 - -
```

2. Enter your user name:

```
ID: jparrot
```

3. Note the challenge the proxy displays for example:

```
Challenge: 17
```

4. On the selected token, enter the challenge the proxy displays.

5. At the response prompt, enter the response the token displays and press **Enter**.

```
Enter Password: 12HAAF
```

SafeWord Authentication Server Example

The following example shows a sample TELNET session from a system outside the firewall to a system inside the firewall:

```
blaze.clientsite.com-28: telnet fire-out.yoyodyne.com
Trying 204.255.154.100...
Connected to fire-out.yoyodyne.com
Escape character is '^]'.
- - Safeword@ Security Check V4.30.000 - -
ID: jparrot
Challenge: 17
Enter Password: 12HAAF
Login Accepted
fire-out.yoyodyne.com telnet proxy (Version 4.0a) ready:
tn-gw> c dimension
Trying 10.0.1.120 port 23...
Connected to dimension.yoyodyne.com
IRIX 6.5 (dimension) (ttyp5)
login: jparrot
Password: guess#me$now (password does not display)
Welcome to dimension.yoyodyne.com
3:57PM up 16 days, 5:35, 4 users, load averages: 0.03, 0.01, 0.00
dimension-26:
```

SecurID System Authentication

The Gauntlet Firewall includes support for the SecurID system, which you can use for authenticating with your firewall.

How SecurID System Authentication Works

The SecurID system is a synchronous authentication system. The Gauntlet Firewall includes client code that is capable of communicating with a separate SecurID ACE/Server to authenticate users.

The SecurID tokens generate and display unpredictable codes that change at a regular time interval (typically every 60 seconds). When a user attempts to log in using the SecurID system, the ACE/Server can independently verify the code that the user enters and either allow or deny login.

Configuring the ACE/Server

To configure the ACE/Server:

1. Make sure that your ACE/Server is using DES encryption and not SDI encryption. Use the *sdinfo* or *sdsetup* programs on your ACE/Server to determine which type of encryption your ACE/Server is using.
2. Create user accounts for your users using the ACE/Server.
3. Register the firewall as a client system on your ACE/Server. Any users that will authenticate from the firewall must be registered on the ACE/Server showing the firewall as one of their clients.
4. Be sure to use the IP address or hostname for the inside address of the firewall if your ACE/Server is running on a system on your inside network.
5. Look at the */etc/services* file on the system running the ACE/Server and make note of the service name and port on which the ACE/Server is listening. For example:

```
securid          755/udp          # securid ACE services
securidprop      5510/tcp
```

6. Place the ACE/Server configuration file (*/var/ace/sdconf.rec*) onto a floppy diskette or into a location accessible from the firewall.

Configuring the Firewall

To configure the firewall:

1. Log in to the firewall and become root.
2. Look at the */etc/services* file on the firewall. Make sure that the service name and port number specified for the SecurID service on the firewall are the same as the ones in the */etc/services* file on the ACE/Server.
3. Copy the ACE/Server configuration file (*/var/ace/sdconf.rec*) to the same directory on the firewall (*/var/ace/sdconf.rec*).
4. Use the Gauntlet Firewall Manager to add information about the ACE/Server:
 - From within the Gauntlet Firewall Manager, select Environment.
 - Click the Authentication tab. The Authentication window displays.
 - Click the SecurID button on the left side of the window. The SecurID Server window displays.

- Enter the hostname or IP address of the Gauntlet Firewall as you registered it the ACE/Server. For example, Yoyodyne might enter:
`fire-in.yoyodyne.com`.
- Before exiting the Gauntlet Firewall Manager, save and apply your changes.

Adding SecurID Users

You can create two types of users in the Gauntlet Firewall for use with SecurID: individual users or a default user. Individual users are unique user names for each user in your SecurID system. For example, if you create an account for jyaya on your ACE/Server, you must also create an account for jyaya on the Gauntlet Firewall.

Creating a default user allows you to authenticate users without manually creating entries for every user in the Gauntlet authentication database. When a user logs in and the authentication server does not find the information in the Gauntlet authentication database, the authentication server sends the user information to the ACE/Server. The authentication server also creates a record for that user in the Gauntlet authentication database.

For example, suppose you create an account for jyaya on your ACE/Server. You create an account for default on the Gauntlet Firewall. When John Yaya authenticates, he still uses the user name jyaya, and the authentication server sends the information to the ACE/Server.

Adding Individual SecurID Users

Use the Gauntlet Firewall Manager to add users.

To add a SecurID user:

1. From the Users tab, click Add.
2. Provide information about the user.

UserID	User name. This must match the user name you have registered with the ACE/Server.
Name	User name.
Status	Click Enable to activate the account.
Group	If you want to make this user a member of a group, select the name of the group.

Authentication Method	Select SecurID as the authentication method.
Password	Leave this field blank. The Gauntlet authentication system uses the value registered with the ACE/Server.
Verification	Leave this field blank.

3. Click OK to make your changes take effect.
4. Provide the user with the user name.

SecurID Default Users

You can only have one default user. If you are also using the SafeWord Authentication Server, you can only use the default user for one of the two authentication systems.

To add a default SecurID user, follow the steps listed under “Adding Individual SecurID Users.” Enter default as the UserID.

Using SecurID with the Gauntlet Firewall

This section first lists the steps for authenticating with SecurID, then provides an example.

Authenticating With SecurID

To authenticate to the Gauntlet firewall using SecurID:

1. Access a proxy that requires authentication:

```
telnet fire-out.yoyodyne.com
Trying 204.255.154.100...
Connected to fire-out.yoyodyne.com
Escape character is '^]'.
```

2. Enter your user name and press **Enter**:

```
Username: jyaya
```

3. At the response prompt, enter the response appropriate for the type of SecurID token you have:

- Standard Token or Key Fob: Enter your PIN (if enabled) followed by the SecurID code displayed on your token, with no spaces in between. Then press **Enter**:
Enter PASSCODE: **1234481283**
 - PINPAD Token: With the PINPAD token, enter your PIN into the card itself, and press the diamond key. The PASSCODE is simply the SecurID code displayed on the token.
Enter PASSCODE: **429162**
 - Press the **Enter** key.
4. On occasion, the system prompts you to enter the next code that appears on your token to resynchronize your SecurID token with the ACE/Server:
- Enter the next cardcode:
- Wait until the code changes on your token, and then enter the new code (without your PIN) at the prompt and press **Enter**.
- Enter the next cardcode: **617325**

SecurID Example

This example shows a sample TELNET session from a system outside the firewall to a system inside the firewall:

```
blaze.clientsite.com-28: telnet fire-out.yoyodyne.com
Trying 204.255.154.100...
Connected to fire-out.yoyodyne.com
Escape character is '^]'.
Username: jyaya
Enter PASSCODE: 429162
Login Accepted
fire-out.yoyodyne.com telnet proxy (Version 4.0a) ready:
tn-gw> c dimension
Trying 10.0.1.120 port 23...
Connected to dimension.yoyodyne.com
IRIX 6.5 (dimension) (ttyp5)
login: jyaya
Password: guess#me$now (password does not display)
Welcome to dimension.yoyodyne.com
3:57PM up 16 days, 5:35, 4 users, load averages: 0.03, 0.01, 0.00
dimension-26:
```


S/Key System

The Gauntlet Firewall includes support for the S/Key system, which you can use for authenticating with your Gauntlet Firewall. You can use both S/Key and S/Key5 with the Gauntlet Firewall.

How the S/Key System Works

The S/Key system is a one-time password authentication system. You enter a secret password into the Gauntlet authentication database when you configure that user. When the proxies want to authenticate, they use the secret password to generate a series of nonsense words and a sequence number. The proxy displays the sequence number. You enter the secret password and the sequence number into a key generation program, which provides a series of nonsense words. You enter the series of nonsense words at the proxy prompt. The proxy passes it back to the authentication server, which compares what you entered with what it expected and allows or denies access.

Adding an S/Key User

Use the Gauntlet Firewall Manager to add users.

To add an S/Key user:

1. From the Users tab, click Add.
2. Provide information about the user:

UserID	Enter the user name. Remember that this user ID does <i>not</i> need to match any other user IDs for this user.
Name	Enter descriptive information about this user.
Status	Click Enable to activate the account.
Group	If you want to make this user a member of a group, select the name of the group.
Authentication Method	If you are using S/Key or S/Key4, select Skey as the authentication method. If you are using S/Key5, select Skey5 as the authentication method.

Password Enter a random string, for example:
 Passwd: **try!and@guess#this\$one**
 The display uses XXXs to hide the password

Verification Enter the same random string again.

3. Click *OK* to make your changes take effect.
4. Provide the user with the user name. If you are creating the user's one-time passwords, provide a set of passwords. If you want users to create their own one-time access passwords, provide them with their random string and instructions on how to generate one-time access passwords.

Generating One-Time Access Passwords

You can generate one-time access passwords yourself and provide them to users all together. You can also allow your users to generate their own one-time passwords themselves as they need them. Using tools on the firewall, you can generate lists of one-time passwords and give them to each user. This involves determining the key value and generating the one-time passwords.

Determining the Key Value

To determine the key value, search the key value file for the user's information:

```
fire-in# grep jparker /etc/skeykeys
```

This displays information about the user:

```
jparker 0664 fi19289                    a6eb4adfeec9bad9 Jul 17,1996 15:57:49
```

The first number is the sequence number and the last string is the key. You need these numbers to generate the one-time passwords.

Generating the One-Time Passwords

This section discusses using the firewall to generate keys for users. You are urged to install the key program on another trusted host inside the firewall and generating the keys on that system.

To generate the one-time passwords:

1. Log in to the firewall and become root.
2. Run the key program (*/usr/bin/key*) to generate a limited number of one-time passwords, specifying the number of passwords, the sequence number, and the key. Redirect it to a file so that you can provide the passwords to the user.

- S/Key

```
# cd /usr/bin
# ./key -n 5 664 fi19289 > /tmp/jparker.key
```

- S/Key5

```
# cd /usr/bin
# ./key -m 5 -n 5 664 fi19289 > /tmp/jparker.key
```

3. Enter the secret password for the user for whom you are creating one-time passwords:

```
Reminder - Do not use key while logged in via telnet or dial-in.
Enter secret password: try!and@guess#this$one
```

4. View the list of one-time passwords you have created:

```
660: RIME SLUM DRY MYRA GORE ELBA
661: LUCY DISK MOSS BACH TUSK BODE
662: JANE HURT SELF RING MILE HOB
663: GOWN BOLT YET BEAD LYON PIT
664: PAR HOOK FLUE BIAS TANK WEEK
```

5. Send this file to a system on your trusted network using FTP, print it and give it to the user. Be sure to delete all copies of this file (on the firewall and on the trusted host) as soon as you have the printout.

Allowing Users to Generate One-Time Passwords

There is an alternative to the administrator using the key program on the firewall to generate one-time passwords for each user. You can obtain the key program for a variety of platforms from the Bellcore FTP site (see <ftp://ftp.bellcore.com/pub/nmh/skey> for more information). You can install the key program on trusted hosts within your network, allowing users to generate their own passwords all at once or as they need them. There are also versions for the Microsoft Windows family and the Macintosh, allowing users to generate keys from their desktops or laptops.

Be sure to:

1. Provide each user with the secret password that you created when you created that user's entry in the Gauntlet authentication database.
2. Remind users not to use the key program from hosts on the untrusted networks.

Using S/Key With the Gauntlet Firewall

This section first lists the steps for authenticating with the S/Key system, then provides an example.

Authenticating With the S/Key System

To authenticate to the Gauntlet Firewall using S/Key:

1. Access a proxy that requires authentication:

```
blaze.clientsite.com-28: telnet fire-out.yoyodyne.com
Trying 204.255.154.100...
Connected to fire-out.yoyodyne.com
Escape character is '^]'.

```
2. Enter your user name:

```
Username: jparker

```
3. Note the challenge the proxy displays:

```
Key Challenge: s/key 663 fi19289

```
4. Locate the one-time password with the corresponding sequence number from your list:

```
663: GOWN BOLT YET BEAD LYON PIT

```

Or use your key program to generate the one-time password.




Figure B-1 One-Time Password Window

5. Respond with a one-time password:

```

Key Challenge: s/key 663 fi19289GOWN BOLT YET BEAD LYON PIT
Login Accepted
fire-out.yoyodyne.com telnet proxy (Version 4.0a) ready:
tn-gw>

```

S/Key Example

This example shows a sample TELNET session from a system outside the firewall to a system inside the firewall:

```

blaze.clientsite.com-28: telnet fire-out.yoyodyne.com
Trying 204.255.154.100...
Connected to fire-out.yoyodyne.com
Escape character is '^]'.
Username: jparker
Key Challenge: s/key 663 fi19289GOWN BOLT YET BEAD LYON PIT
Login Accepted
fire-out.yoyodyne.com telnet proxy (Version 4.0a) ready:
tn-gw> c dimension
Trying 10.0.1.120 port 23...
Connected to dimension.yoyodyne.com
IRIX 6.5 (dimension) (ttyp5)
login: jparker
Password: guess#me$now (password does not display)
Welcome to dimension.yoyodyne.com
3:57PM up 16 days, 5:35, 4 users, load averages: 0.03, 0.01, 0.00
dimension-26:

```

RADIUS Authentication

The Gauntlet Firewall includes support for the RADIUS authentication protocol, which you can use to authenticate to your firewall.

How RADIUS Authentication Works

RADIUS stands for Remote Authentication Dial-In User Service, an authentication protocol specified by the IETF. Many different authentication vendors, including SecurID, Safeword, VASCO, and CRYPTO, support RADIUS.

Note: The Gauntlet Firewall uses RADIUS in a way that differs slightly from the standard use. RADIUS is normally used as a dial-in authentication service. The Gauntlet Firewall uses it as an authentication method.

RADIUS authentication can be used with your Gauntlet Firewall in conjunction with a strong authentication method such as Safeword or CRYPTO, or by itself using a plain RADIUS password.

Note: It is important that if you use RADIUS authentication in conjunction with a strong authentication method, you should explain to your users how your strong authentication method works and how to use it.

To use RADIUS with your Gauntlet Firewall, the user connects to the firewall in order to access the trusted network. The firewall prompts the user for a user name and RADIUS password, and then, acting as the RADIUS client, encrypts the authentication information with the RADIUS shared secret and sends it to the RADIUS authentication server.

Note: If you are using a strong authentication method, your use of RADIUS may be somewhat different.

The RADIUS authentication server decrypts the authentication information with the shared secret and begins the authentication process. The RADIUS authentication server authenticates the user, and sends this information to the Gauntlet Firewall, which grants access permission to the user.

Note: You must have a RADIUS authentication server up and running, know its shared secret, and have the user configured on the server before you can use the Gauntlet Firewall to support a RADIUS user.

Configuring the RADIUS Authentication Server

To configure the RADIUS authentication server, create user accounts for your users on your RADIUS authentication server.

Enabling RADIUS Support

Use the Gauntlet Firewall Manager to enable RADIUS support.

To enable RADIUS support:

1. From within the Gauntlet Firewall Manager, select Environment.
2. Click the Authentication tab.

The Authentication window displays.

3. Click the *RADIUS* button on the left side of the window.

The RADIUS Authentication Servers window displays.

4. Make the appropriate entries:

Shared Secret for the Primary RADIUS Server Shared secret string for your primary RADIUS server.

Host IP address or host name of your primary RADIUS server.

Port Port to be used to access your primary RADIUS server. The default is 1645.

Shared Secret for the Secondary RADIUS Server Shared secret string for your secondary RADIUS server. Make an entry in this field only if you are using a secondary RADIUS server.

Host IP address or host name of your secondary RADIUS server. Make an entry in this field only if you are using a secondary RADIUS server.

Port Port to be used to access your secondary RADIUS server. Make an entry in this field only if you are using a secondary RADIUS server.

5. Before exiting the Gauntlet Firewall Manager, save and apply your changes.

Adding a RADIUS User

Use the Gauntlet Firewall Manager to add a RADIUS user.

To add a RADIUS user:

1. From within the Gauntlet Firewall Manager, select Firewall Rules.
2. Click the Users tab.
The Users window displays.
3. From the Users tab, click *Add*.
The Add/Modify window displays.

4. Provide information about the user.

UserID	User name. Remember that this user ID does not need to match any other user IDs for this user.
Name	Descriptive information about this user.
Status	Click <i>Enable</i> to activate the account.
Group	If you want to make this user a member of a group, select the name of the group.
Authentication Method	Select RADIUS as the authentication method.
Plain Password	Check this box to tell the RADIUS server to use a plain password for this user; leave blank to use a strong authentication method.

5. Click OK.
6. Provide each user with a user ID.

Using RADIUS With the Gauntlet Firewall

The examples shown in this section use plain RADIUS passwords. If you are using a strong authentication method, the text will vary depending on the method you are using.

Authenticating With RADIUS

To authenticate to the Gauntlet Firewall using a plain RADIUS password:

1. Access a proxy that requires authentication:

```
telnet fire-out.yoyodyne.com
Trying 204.255.154.100...
Connected to fire-out.yoyodyne.com
Escape character is '^']'
```

2. Enter your user name and press **Enter**.

```
Username: jnolan
```

3. Enter your RADIUS password and press **Enter**. The password itself does not display; an X appears when each character is typed:

```
RADIUS Password: XXXXXXXX
```

4. If the user is a valid user, the login is accepted and the Gauntlet Firewall permits access to the trusted network.

```
Login Accepted
```

RADIUS Authentication Example

This example shows a sample TELNET session from a system outside the firewall to a system inside the firewall using a plain RADIUS password:

```
blaze.clientsite.com-28: telnet fire-out.yoyodyne.com
Trying 204.255.154.100...
Connected to fire-out.yoyodyne.com
Escape character is '^]'.
Username: jnolan
RADIUS Password: #####
Login Accepted
fire-out.yoyodyne.com telnet proxy (Version 4.0a) ready:
tn-gw> c dimension
Trying 10.0.1.120 port 23...
Connected to dimension.yoyodyne.com
```

Reusable Passwords

The Gauntlet Firewall includes support for reusable passwords. The reusable password system does *not* include support for password aging, minimum characters, or other security features of some reusable password systems.

Caution: Do not use the reusable passwords option for authentication from untrusted networks. You should not use reusable passwords. Reusable passwords are vulnerable to password sniffers and are easy to crack. This feature is provided for convenience and audit capability only.

How Reusable Passwords Work

You enter a secret password into the Gauntlet user authentication system when you configure that user. When an application on the firewall needs to authenticate you, they ask the Gauntlet authentication system. You enter the secret password at the prompt. The application passes it back to the authentication system, which compares what you entered with what it expected and allows or denies access.

Adding a User with Reusable Passwords

Use the Gauntlet Firewall Manager to add users.

To add a user using reusable passwords:

1. From the Users tab, click *Add*.
2. Provide information about the user.

UserID	Enter the user name. Remember that this user ID does <i>not</i> need to match any other user IDs for this user.
Name	Enter descriptive information about this user.
Status	Click Enable to activate the account.
Group	If you want to make this user a member of a group, select the name of the group.
Authentication Method	Select Password as the authentication method.

Password Enter a random string as the password.

Verification Enter the same random string you entered as the password.

3. Click *OK* to make your changes take effect.
4. Provide each user with a user ID and password.

Using Passwords with Gauntlet

This section first lists the steps for authenticating with reusable passwords, then provides an example.

Authenticating With Reusable Passwords

To authenticate to the Gauntlet Firewall using reusable passwords:

1. Access a proxy that requires authentication:

```
blaze.clientsite.com-28: telnet fire-out.yoyodyne.com  
Trying 204.255.154.100...  
Connected to fire-out.yoyodyne.com  
Escape character is '^]'.  
^Z  
^C
```

2. Enter your user name:

```
Username: jgomez
```

3. Enter your password:

```
Password: try!and@guess#this$one  
Login Accepted  
fire-out.yoyodyne.com telnet proxy (Version 4.0a) ready:  
tn-gw>
```

Reusable Password Authentication Example

This example shows a sample TELNET session from a system outside the firewall to a system inside the firewall:

```
blaze.clientsite.com-28: telnet fire-out.yoyodyne.com
Trying 204.255.154.100...
Connected to fire-out.yoyodyne.com
Escape character is '^]'.
Username: kgomez
Password: try!and@guess#this$one (password does not display)
Login Accepted
fire-out.yoyodyne.com telnet proxy (Version 4.0a) ready:
tn-gw> c dimension
Trying 10.0.1.120 port 23...
Connected to dimension.yoyodyne.com
IRIX 6.5 (dimension) (ttyp5)
login: kgomez
Password: guess#me$now (password does not display)
Welcome to dimension.yoyodyne.com
3:57PM up 16 days, 5:35, 4 users, load averages: 0.03, 0.01, 0.00
dimension-26:
```

Glossary

address scanning

Searching network address space via DNS inverse queries.

address spoofing

A host purporting to be another, usually trusted, host.

administrator

The individual responsible for a system or network or systems. The firewall administrator is responsible for the firewall.

application gateway

A protocol-specific data forwarder.

authenticated Post Office Protocol (APOP)

A version of POP that uses non-reusable passwords for authentication.

address resolution protocol (ARP)

Allows a host to find the physical address of a target host on the same physical network, given only the target's IP address. The protocol is used to dynamically bind an IP Address to a physical hardware address. The use of ARP is restricted to a single physical network and is limited to networks that support hardware broadcast.

authentication

Method to guarantee that the sender of information is who the sender purports to be.

bastion host

A secure computer that forms part of a security firewall and runs applications that communicate with computers outside an organization.

berkeley internet name domain (BIND)

In UNIX, DNS is implemented by the BIND.

cache

A portion of memory or storage that contains a “quick reference” to recently used information.

chroot

The *chroot* mechanism allows a program to irreversibly change its view of the filesystem by changing the where the root of the filesystem is. When a program *chroots* to a particular portion of a given filesystem, that portion becomes the whole filesystem and, in effect, the rest of the filesystem ceases to exist, from the program’s point of view.

common gateway interface (CGI)

The piece of HTTP that specifies how user information is communicated to the server and from it to external programs.

circuit level gateway

A protocol gateway for a specific service type.

connection laundering

A (usually) FTP or Gopher request initiated by an untrusted client that appears to come from a trusted server.

denial of service

An attack that is aimed entirely at preventing use of your own equipment.

data encryption standard (DES)

The most widely used symmetric cryptosystem.

domain name system (DNS)

The online distributed database system used to map human-readable machine names into IP addresses. DNS servers throughout the Internet implement a hierarchical namespace that allows sites to assign machine names and addresses.

domain

A part of the DNS naming hierarchy. Domain names consist of a sequence of names (labels) separated by periods (dots).

digital signature standard (DSS)

A mechanism for identifying a message source.

dual-homed

A dual-homed host has two network interfaces, and therefore two addresses, and acts as a router between the subnetworks to which those interfaces are attached.

encryption

A mechanism for changing the appearance of data. Encryption allows the creation of secure connections over insecure channels. Encrypting network traffic provides both privacy and authentication

finger

A service that looks up information about a user who has an account on the machine being queried. It tells whether or not that user is currently logged into the machine and may include the person's real name, login, phone number, office location, and other information.

firewall

A configuration of routers and networks placed between an organization's internal internet and a connection to an external internet to provide security.

file transfer protocol (FTP)

The TCP/IP standard, high-level protocol for file transfer from one machine to another. FTP uses TCP.

ftp daemon (ftpd)

One of the programs that implement FTP.

gateway daemon (gated)

A program run on a host or router that collects routing information from within one autonomous system and advertises the information to another autonomous system.

gateway

Dedicated host that interconnects two different services or applications.

group

A collection of users with a common security concern.

hardened

A term indicating that an operating system or application has been modified to eliminate elements that make it vulnerable to attack or failure.

hypertext markup language (HTML)

The language used to implement network resource pages.

hypertext transfer protocol (HTTP)

The primary application protocol that underlies the World Wide Web.

inetd

The program that listens for requests for services specified in the */etc/inetd.conf* configuration file. When it hears such a request, it starts the proper server to process the request.

inside network

The network of machines protected by the firewall (inside the security perimeter).

IP address

A 32-bit integer address assigned to each host on the Internet.

internet security scanner (ISS)

When ISS is run from another system and directed at your system, it probes your system for software bugs and configuration errors commonly exploited by crackers.

mail exchanger

A host that accepts e-mail; some mail exchangers forward the mail to other hosts.

mail exploder

Part of an electronic mail system that accepts a piece of mail and a list of addresses as input and sends a copy of the message to each address on the list.

mail gateway

A host that connects to two or more dissimilar electronic mail systems and transfers mail messages among them.

multicast route daemon (mrouted)

A program used with a multicast kernel to establish multicast routing.

name resolution

The process of mapping a name into a corresponding address. The domain name system provides a mechanism for naming hosts in which programs use remote name servers to resolve a host name into an IP address.

netacl

A program that provides the capability of a TCP Wrapper.

netperm table

The network permissions table (*/usr/local/etc/netperm-table*) that contains Gauntlet Firewall configuration information used by the kernel, proxies, and other applications. The configuration information is in the form of policy or applications rules.

network file system (NFS)

A protocol developed by Sun Microsystems, Inc. that uses IP to allow a set of cooperating computers to access each other's file systems as if they were local.

outside network

The network of machines not protected by the firewall (outside the security perimeter). When a firewall protects a network connected to the Internet, the outside network is the rest of the Internet.

packet filter

A method to select or deselect traffic from given network addresses.

packet filtering gateway

A system that serves as a gateway for protocols using simple packet content rules.

packet internet groper (ping)

The name of a program used with TCP/IP internets to test reachability of destinations by sending them an ICMP echo request and waiting for a reply.

plug gateway

A general-purpose program implemented as a proxy which allows data to flow from an inside host to an outside host.

post office protocol (POP)

A client-server protocol for handling user electronic mail boxes. The user's mailbox is kept on the server, rather than on the user's personal machine.

port

A specific pathway for data and control information.

port scanning

To probe given ports to determine what type of data or control information is normally passed via the given ports.

point to point protocol (PPP)

A protocol for framing IP across a serial line. A more recent protocol than SLIP.

promiscuous mode

Most Ethernet and token ring interfaces can operate in this mode to view all packets on an Ethernet or token ring.

protocol

A formal description of message formats and the rules that must be followed to exchange those messages

proxy

Specialized applications or programs that run on a firewall host. These programs take users' requests for Internet services (such as FTP and TELNET) and forward them according to the site's security policy. Proxies are replacements for actual services and serve as application-level gateways to the services.

proxy arp

The technique in which a host or router answers ARP requests intended for another by supplying its own physical address. The purpose is to allow a site to use a single IP network address with multiple physical networks.

public key encryption

An encryption technique that generates encryption keys in pairs. One of the pair, used to decrypt, is kept secret, and the other, used to encrypt, is published.

remote login (rlogin)

The remote login protocol developed for UNIX by the University of California in Berkeley. It offers a service similar to TELNET.

route daemon (routed)

A program devised for UNIX that implements the RIP protocol. Pronounced "route dee."

router

A special purpose, dedicated machine that attaches to two or more networks and forwards packets from one to the other. An IP router forwards IP datagrams among the networks to which it is connected. An IP router uses the destination address on the datagram to choose the next hop to which it forwards a datagram.

S/Key

A one-time password mechanism that allows a system to authenticate a user reliably. The S/Key encodes each key into a series of short words, so they are easier for a user to read and type, rather than generating random characters.

screening router

The type of router used in a packet filtering firewall.

security perimeter

The mechanisms used to protect a network of machines.

serial line internet protocol (SLIP)

A framing protocol used to send IP across a serial line. SLIP is popular when sending IP over dialup phone lines but has been largely replaced by PPP.

smap

A small program intended solely to handle incoming SMTP connections

smapd

A program that is invoked regularly (typically once a minute) to process the files queued in the queue directory, normally by handing them to *sendmail* for delivery.

simple mail transfer protocol (SMTP)

The TCP/IP standard protocol for transferring electronic mail messages from one host to another. SMTP specifies how two hosts interact and the format of control messages they exchange to transfer mail.

simple network monitoring protocol (SNMP)

A standard protocol used to monitor hosts, routers, and the networks to which they attach.

socket

The abstraction provided by the UNIX operating system that allows an application program to access the TCP/IP protocols.

socks

The SOCKS package is an example of the type of proxy system that requires custom clients.

source route

A route that is determined by the source. In IP, a source route consists of a list of routers a datagram should visit; the route is specified as an IP option. Source routing is most often used for debugging but should be rejected by most hosts.

strong authentication system

A system for verifying users which uses one-time, non-reusable passwords.

subnet

The portion of an IP address can be locally modified by using host address bits as additional network address bits. These newly designated network bits define a network within the larger network.

subnet addressing

An extension of the IP addressing scheme that allows a site to use a single IP network address for multiple physical networks by dividing the destination address into a network portion and local portion.

transmission control protocol/internet protocol (TCP/IP)

A suite of data communications protocols.

TCP wrapper

The TCP wrapper package monitors incoming network traffic and controls network activity. It is a simple but effective piece of publicly available software set up to run whenever certain ports (corresponding to certain services) are connected.

transparency

A method for providing network access through a firewall without user interaction with the firewall. Access that is allowed at a site is done invisibly to the user.

trusted network

The network of machines protected by the firewall.

user datagram protocol (UDP)

The TCP/IP standard protocol that allows an application program on one host to send a datagram to an application program on another. UDP uses IP to deliver datagrams but UDP includes a protocol port number, allowing the sender to distinguish between application programs on a given remote host.

untrusted network

The network of machines not protected by the firewall, but from which the firewall accepts requests.

uniform resource locator (URL)

A string that gives the location of a information. The string begins with a protocol type (for example, FTP, HTTP) followed by the domain name of a server and the path name to a file on that server.

virtual private network

A physically disparate set of networks that share a common security perimeter through secured internetwork communication.

web browser

A software program that lets you access the World Wide Web. Netscape Navigator and Microsoft Internet Explorer are well-known Web browsers.

well-known port

Any of a set of protocol port numbers assigned for specific uses by transport level protocols, (for example, SMTP and UDP). Each server listens at a well-known port so clients can locate it.

World Wide Web (WWW, the Web)

The large-scale information service that allows a user to browse through information. The Web offers a hypermedia system that can store information as text, graphics, audio, and so on.

Index

A

- absorb rules, 228
- access control, 71
- access control (LDAP), 95
- accessing circuit proxy configuration, 200
- accessing content scanning configuration, 288
- accessing destination access configuration, 43
- accessing encryption configuration, 265
- accessing FTP proxy configuration, 89
- accessing Gauntlet Firewall Manager, 20
- accessing Gopher proxy configuration, 172
- accessing HTTP proxy configuration, 171
- accessing Info Server configuration, 279
- accessing lp proxy configuration, 141
- accessing Microsoft SQL Server proxy, 103
- accessing multimedia configuration, 111
- accessing network configuration, 50
- accessing network groups configuration, 53
- accessing packet screening configuration, 232
- accessing plug proxy configuration, 192
- accessing POP3 proxy configuration, 220
- accessing report configuration, 244
- accessing rsh proxy configuration, 147
- accessing service group configuration, 31
- accessing service group rules, 36
- accessing SMTP proxy configuration, 215
- accessing SNMP agent configuration, 307
- accessing SNMP proxy configuration, 123
- accessing SSL proxy configuration, 171
- accessing Sybase configuration, 153
- accessing user configuration, 63
- accessing web services
 - non-proxy-aware browsers, 179
- accessing X11 proxy configuration, 183
- Access Key II, 60, 331
- access *See Also* destination access
- access to destinations, 39
- accounts
 - creating user accounts, 15
 - on the firewall, 16
- ACE/Serve, 62
- activating license, 325
- activities
 - specifying allowed activities, 3
- addfile program, 282
- adding files to Info Server, 281
- adding network groups, 49
- adding new service groups, 30
- adding packet screening rules, 233
- adding service group rules, 37
- addtext program, 281
- addUserAccount command, 15
- Administrator's Guide
 - audience, xxix
 - conventions, xxx
- A files, 277
- Alcohol, Beer, Wine & Tobacco category, 299

- alerts, 247
- allowed activities, 3
- Allow Password Change option, 34
- America Online *See* AOL
- anonymous FTP server, 94
- AOL, 6, 189, 190, 191
- APOP, 60, 219
 - authentication, 213
 - password, 222
- application-level proxies, 81
- Assign Access option (security rules), 37
- audience, xxix
- authenticated HTTP, 169
 - configuring, 174
 - creating authentication entries, 175
- authentication
 - and GUI FTP tools, 92
 - APOP, 213
 - authmgr program, 59
 - changing methods, 66
 - custom services, 197
 - entries, 175
 - overview, 57
 - See Also* strong authentication
 - token, 57
 - weak authentication, 178
- authentication information
 - login-sh program, 58
 - non-firewall services, 58
 - used by firewall, 58
- authentication management system
 - and FTP proxy, 90
 - groups, 59
 - HTTP, 170
 - logs, 240
 - rlogin, 162
 - TELNET, 162
 - users, 59
- authmgr program

- authentication information, 59
- AVI movie header, 282

B

- backdoors, 2
- backing up the firewall, 16
- Bellcore, 62

C

- Cancel menu command, 22, 24
- certificate management, 6
- changing authentication methods, 66
- changing group membership, 65
- changing passwords, 66
- changing user IDs, 65
- changing user names, 65
- checking source and destination, 13
- checksums, 253
- choosing host, 18
- circuit proxy, 6, 197
 - accessing configuration, 200
 - configuring firewall, 201
 - configuring settings, 202
 - database applications, 198
 - enabling, 204
 - financial applications, 198
 - groupware, 198
 - how it works, 199
 - planning, 201
 - strong authentication, 198
 - understanding, 198
 - using, 204
 - verifying setup, 204
 - vs. plug proxy, 197
- ck-gw daemon, 199

- clients
 - configuring for Sybase proxy, 156
 - SMTP proxy client (smap), 214
 - starting TN3270 client, 165
 - X client, 182
- color setting requirement, 20
- CompuServe, 6, 189, 190
- configurable access control (LDAP), 95
- configurable logging (LDAP), 95
- configuration information
 - upgrading, 328
- configurations
 - circuit proxy, 200
 - content scanning, 288
 - Cyber Patrol, 299
 - destination access, 43
 - Encryption Key, 265
 - ftp proxy, 89
 - Gopher proxy, 172
 - HTTP proxy, 171
 - Info Server, 279
 - LDAP proxy settings, 97
 - lp proxy, 141
 - Microsoft SQL Server, 103
 - multimedia proxy, 111
 - network, 50
 - network groups, 53
 - plug proxy, 192
 - POP3 proxy, 220
 - rlogin, 160
 - rsh proxy, 147
 - service group, 31
 - service group rules, 36
 - SMTP proxy, 215
 - SNMP agent, 307
 - SNMP proxy, 123
 - SSL proxy, 171
 - Sybase proxy, 153
 - TELNET, 160
 - URL filtering, 294
 - user restrictions, 73
 - users, 63
 - VNP, 268
 - X11, 183
- configuration sets
 - creating, 84
 - deleting, 85
 - modifying, 85
 - name change procedure, 85
 - planning, 83
- configurations files
 - editing with text-based interface, 16
- configuring additional logging, 241
- configuring circuit proxy settings, 202
- configuring content scanning, 290
- configuring Cyber Patrol, 299
- configuring firewall for plug proxy, 192
- configuring firewall for rsh proxy, 147
- configuring FTP proxy, 90
- configuring Gopher proxy, 174
- configuring HTTP proxy, 174
- configuring LDAP clients, 97
- configuring LDAP proxy settings, 97
- configuring log retention time, 241
- configuring lp proxy, 142
- configuring Microsoft SQL clients, 108
- configuring Microsoft SQL Server proxy, 104, 105
- configuring multimedia proxy, 113
- configuring News Feed settings, 131
- configuring News Reader settings, 133
- configuring news server, 134
- configuring non-proxy-aware web browsers, 179
- configuring plug proxy, 193
- configuring print client, 140
- configuring print server, 140
- configuring proxies, 82

- configuring remote system for rsh service, 149
- configuring report frequency, 245
- configuring report recipients, 244
- configuring rlogin proxy, 162
- configuring rsh proxy, 148
- configuring SMPT proxy, 216
- configuring SMPT proxy settings, 216
- configuring SNMP agent settings, 308
- configuring SNMP proxy, 125
- configuring SSL proxy, 174
- configuring Sybase proxy firewall, 154
- configuring Sybase proxy settings, 154
- configuring TELNET proxy, 162
- configuring the firewall, 19
- configuring URL filtering, 294
- configuring VNPs, 268
- configuring web browsers, 20, 176
- configuring X11 proxy, 184
- configuring your system, 20
- configuring reports, 244
- content scanning
 - accessing configuration, 288
 - configuring, 290
 - configuring firewall, 289
 - configuring scanning engines, 291
 - enabling, 290
 - FTP, 286
 - how it works, 287
 - HTTP, 286
 - Infected File Handling, 290
 - Java applets, 287
 - planning, 290
 - quarantine area, 290
 - SMPT, 286
 - understanding, 285
- Content Vectoring Protocol, 287, 291
- conventions, xxx
- creating authentication entries, 175
- creating configuration sets, 84
- creating destination access rules, 45
- creating encryption keys, 266
- creating groups, 69
- creating Info Server files, 280
- creating integrity database, 256
- creating logs, 240
- creating network groups, 54
 - Network Group Name option, 55
- creating networks, 51
 - Interface option, 52
 - Network IP Address option, 51
- creating packet screening rules, 233
- creating passthrough links, 270
- creating reports, 242
- creating service groups, 33
- creating trusted VNPs, 269
- creating user accounts, 15
- creating user restriction rules, 75
- creating users, 64
- CRYPTOCARD RB-I, 61, 334
- ctavi, 282
- ctgif, 282
- cthtml, 282
- ctjpg, 282
- ctps, 282
- ctqt, 282
- cttext, 282
- ctzip, 282
- custom services, 189
 - with authentication, 197
- CVP, 287, 291
- CyberNOT, 299
- CyberNOT database, 297
- Cyber Patrol, 293

Alcohol, Beer, Wine & Tobacco, 299
 configuring, 299
 Drugs & Drug Culture, 298
 enabling, 304
 Full Nudity, 297
 Gambling/Questionable/Illegal, 298
 Gross Depictions/Text, 298
 Intolerance, 298
 license, 302
 Militant/Extremist, 298
 overview, 297
 Partial Nudity & Art, 297
 Satanic or Cult, 298
 Search Engines, 299
 Sex Education, 298
 Sexual Acts/Text, 297
 Sports & Entertainment, 299
 Violence/Profanity, 297
 Work Time, 303

D

database applications and circuit proxy, 198
 databases
 Info Server, 276
 Data Encryption Standard, 264
 decrypting data, 264
 default destination access rules, 40
 default network groups, 48
 default service groups, 28
 Defender Security Server, 339
 deleting configuration sets, 85
 deleting destination access rules, 46
 deleting encryption keys, 267
 deleting groups, 70
 deleting network groups, 56
 deleting networks, 53
 deleting packet screening rules, 235

deleting service group rules, 38
 deleting service groups, 35
 deleting user restriction rules, 76
 deleting users, 69
 deleting VNPs, 272
 deny rules, 227
 DES, 264
 destination access
 accessing configuration, 43
 changing order, 46
 creating rules, 45
 default rules, 40
 deleting rules, 46
 how it works, 40
 modifying rules, 46
 understanding, 39
 unknown keyword, 42
 destination access rules
 changing order of precedence, 46
 order of precedence, 40
 planning, 45
 destination addresses
 specifying, 42
 destination IP address rules, 229
 destination port rules, 231
 destination restrictions for service groups, 34
 destinations
 checking, 13
 See also destination access, 13
 Digipass, 61, 339
 directories
 Info Server, 277
 disabled services, 5
 disabling groups, 70
 disabling users, 68
 dot (.) character and Info Server, 277
 Drugs & Drug Culture category, 298

dual-homed bastion host, 10

E

editing configuration files, 16

electronic mail, 6, 213

email, 6, 213

enabling circuit proxy, 204

enabling content scanning, 290

enabling Cyber Patrol, 304

enabling FTP proxy services, 90

enabling Gopher proxy, 175

enabling HTTP proxy, 175

enabling Info Server, 283

enabling Java access, 20

enabling LDAP proxy services, 99

enabling lp proxy, 144

enabling Microsoft SQL Server proxy, 107

enabling multimedia proxy, 113

enabling News proxy, 133

enabling plug proxy, 195

enabling POP3 proxy services, 222

enabling rlogin proxy, 162

enabling rsh proxy, 148

enabling SMTP proxy, 217

enabling SNMP agent, 308

enabling SNMP proxy, 126

enabling SSL proxy, 175

enabling Sybase proxy, 155

enabling TELNET proxy, 162

enabling users, 68

enabling X11 proxy, 185

encrypted information, 12

encrypting data, 264

encryption keys, 265

creating, 266

deleting, 267

modifying, 267

planning, 265

encryption policy configuration utility, 273

encryption through multiple firewalls, 263

events to ignore in exception reports, 245

examining packets, 13

exception reports, 242, 243

events to ignore, 245

example, 251

possible items of interest section, 250

reading, 249

security alerts section, 250

system warnings section, 250

exchanging encrypted information, 12

executable programs and Info Server, 278

exiting Gauntlet Firewall Manager, 22

F

field rules, 229

files

infected, 287

Info Server, 277

file transfer activity and FTP proxy, 87

file transfer services, 6

filtering URLs, 293

financial applications

circuit proxy, 198

finger, 6, 190

firewall

back up and restore, 16

circuit proxy, 201

configuring, 19

configuring as SNMP agent, 307

configuring for multimedia proxy, 112

- configuring for News proxy, 131
- configuring for plug proxy, 192
- configuring for SNMP service, 124
- configuring for X11 proxy, 183
- configuring rlogin proxy, 161
- configuring Sybase proxy, 154
- configuring TELENT proxy, 161
- definition, 1
- Info Server, 280
- Info Server files, 280
- logging, 239
- lp proxy, 142
- mail exchange host, 218
- mail forwarder, 218
- planning, 18
- planning for SNMP service, 124
- rsh proxy, 147
- standard configuration, 9
- text-based interface, 16
- transparency, 4
- trusted networks, 2
- untrusted networks, 3
 - using authentication information, 58
- visibility, 4
- firewall account, 16
- firewall rules, 28
- FTP
 - and content scanning, 285, 286
 - services, 6, 91
- ftpd, 88
- ftp-gw, 88
- FTP proxy
 - accessing configuration, 89
 - and authentication management system, 90
 - configuring, 89, 90
 - enabling, 90
 - for Web services, 170
 - GUI tools and authentication, 92
 - how it works, 88
 - planning settings, 90

- understanding, 87
 - using authentication, 91
 - verifying setup, 91
- FTP server, 94
- Full Nudity category, 297
- fwadmin user ID, 19

G

- Gambling/Questionable/Illegal category, 298
- gauntlet-admin command, 19
- Gauntlet Firewall
 - concepts, 1
 - described, 1
 - design philosophy, 1
 - dual-homed bastion host, 10
 - hardware components, 5
 - how it works, 8
 - IP screening facility, 7
 - maintenance, 15
 - management utilities, 8
 - managing, 15
 - managing options, 16
 - operating system, 5
 - processing packets, 12
 - proxies, 5
 - security perimeter, 2
 - software components, 5
 - trap, 310
 - user accounts, 2
- Gauntlet Firewall Manager, 8
 - accessing, 20
 - configuring web browser, 20
 - configuring your system, 20
 - described, 17
 - exiting, 22
 - interface, 17
 - online help, 24
 - saving your changes, 22

- server, 17
 - using, 18
 - using remotely, 16
- gauntlet license
 - activating, 325
- gauntlet-upgrade, 329
- G files, 278
- GIF image header, 282
- Gopher+, 6
- Gopher menu files, 283
- Gopher proxy
 - accessing configuration, 172
 - configuring, 174
 - enabling, 175
 - how it works, 168
 - understanding, 168
 - using Gopher services, 180
- gopher service proxy
 - planning, 173
- Gopher services, 6, 180
 - Info Service, 278
- graphical interface to Gauntlet. *See* Gauntlet Firewall Manager
- Gross Depictions/Text category, 298
- groups
 - authentication management system, 59
 - creating, 69
 - deleting, 70
 - disabling, 70
 - managing, 69
 - understanding, 59
 - See Also* service groups
- groups of systems, 47
- groupware, circuit proxy, 198

H

- hardened operating system, 5

- hardware components of Gauntlet Firewall, 5
- Help button, 24
- H files, 277
- host name
 - access, 39
 - unknown, 42
- hosts
 - choosing, 18
- HTML text header, 282
- HTTP, 6
 - and content scanning, 285
 - content scanning, 286
- HTTP access for untrusted service group, 29
- http-gw daemon, 168, 170
- HTTP proxy
 - accessing configuration, 171
 - and NetShow player, 115
 - authenticated HTTP, 169, 174
 - configuring, 174
 - default setup, 169
 - enabling, 175
 - how it works, 168
 - nontransparent access, 175
 - persistent connections, 176
 - planning, 173
 - strong or weak authentication, 168
 - transparent access, 176
 - understanding, 168

I

- ICMP and network monitoring, 225
- ICMP traffic, 126
- IETF authentication protocol (RADIUS), 62
- Infected File Handling option, 290
- infected files, 287
- info-gw daemon, 276

Info Server

- accessing configuration, 279
- adding files, 281
- addtext program, 281
- A files, 277
- creating files, 280
- database, 276
- data files, 277
- directories, 277
- dot (.) character, 277
- enabling, 283
- G files, 278
- Gopher menu, 278
- Gopher menu files, 283
- headers, 282
- H files, 277
- placing files on firewall, 280
- planning, 280
- Q files, 278
- queries, 278
- query files, 283
- sample headers, 282
- understanding, 275
- using, 284

Info server

- configuring firewall, 280
- how it works, 276

Info Server addfile program, 282**integrity database, 253**

- creating, 256
- disk space requirement, 256
- planning, 256
- protection, 257
- updating, 257

integrity *See Also* system integrity**Interface option (creating networks), 52****interface rules, 230****interfaces**

- using simultaneously, 17

interface to Gauntlet, 17**internal mail hub, 218****internal mail server, 223****Intolerance category, 298****IP address**

- access, 39
- modifying (not permitted), 52

ipfs utility, 236**IP screening facility, 7****IP-spoofing checks, 226****IRIX for Gauntlet Firewall, 5****IRIX sendmail program, 213****J****Java access, 20****Java applets, content scanning, 287****Java-based interface****JPEG image header, 282****K****keys for encryption, 265****L****LDAP, 6****LDAP proxy, 6**

- configuring clients, 97
- configuring settings, 97
- enabling services, 99
- how it works, 96
- planning, 97
- recommended configuration, 96
- understanding, 95

license

- activating, 325

- links,creating passthrough links, 270
- loading packet screening rules, 235
- log file, 248
- logging, 239
 - configuring additional logging, 241
 - configuring log retention time, 241
 - configuring logs, 241
 - configuring reports, 244
 - creating reports, 242
 - default, 241
 - LDAP, 95
 - log file, 248
 - security alerts, 247
 - understanding, 239
- logging and reporting
 - creating logs, 240
- Login shell
 - configuring, 312
 - how it works, 312
 - understanding, 311
 - using the program, 315
- login-sh program and authentication information, 58
- log retention time, 241
- logs
 - configuring, 241
 - creating, 240
 - reading, 248
- Lotus Notes, 7, 189, 190
- lp, 6
- lp proxy
 - accessing configuration, 141
 - configuring, 142
 - configuring firewall, 142
 - configuring print server, 140
 - configuring the print client, 140
 - default configuration, 138
 - enabling, 144
 - how it works, 138
 - planning, 142

- Print Manager, 140
- transparency, 140
- understanding, 137
- using, 144
- using lp services, 144

M

- mail exchange host, 218
- mail forwarding, 218
- mail setup,verifying, 218
- maintenance of Gauntlet Firewall, 15
- management
 - backing up and restoring, 16
 - creating user accounts, 15
 - tools, 16
- management information base (MIB), 309
- managing custom services, 189
- managing custom services with authentication, 197
- managing Gauntlet Firewall, 15
- managment utilities, 8
- manually adding packet screening rules, 233
- MD5 secure hash algorithm, 60
- MediaBase proxy, 207
- membership in service groups, 30
- MIB, 309
- MIB-II, 309
- Microsoft Internet Explorer, 135
 - proxy awareness, 176
- Microsoft SQL, 6
 - configuring clients, 108
 - replication server, 102
- Microsoft SQL Server proxy
 - accessing configuration, 103
 - configuring, 104, 105
 - configuring Microsoft SQL clients, 108
 - enabling services, 107

- how it works, 102
 - recommended configuration, 103
 - understanding, 101
 - verifying setup, 108
- Militant/Extremist category, 298
- modifying configuration sets, 85
- modifying destination access rules, 46
- modifying encryption keys, 267
- modifying network groups, 56
- modifying networks, 52
- modifying packet screening rules, 235
- modifying service group rules, 38
- modifying service groups, 34
- modifying user restriction rules, 76
- modifying users, 64
- modifying VNPs, 271
- monitoring tools, 225
- mssql-gw, 102
- multimedia proxy
 - accessing configuration, 111
 - configuring, 113
 - configuring firewall, 112
 - enabling, 113
 - how it works, 110
 - planning firewall, 112
 - understanding, 110
 - verifying setup, 113
- multimedia services, 6
- multiple configurations for a proxy, 82
- multiple firewalls, 263
- multiple POP3 servers, 224
- N**
- netacl daemon, 88, 159
- Netscape Navigator, 135
 - proxy awareness, 176
- NetShow, 6
 - TCP port 1755, 110
- NetShow player
 - and HTTP proxy, 115
 - configuring, 114
- NetShow proxy
 - configuring, 111
 - transparency, 114
 - using, 114
- network access control daemon, 88, 159
- Network Group Name option, 55
- network groups
 - accessing configuration, 53
 - creating, 54
 - default, 48
 - definition, 48
 - deleting, 56
 - modifying, 56
 - planning, 54
 - understanding, 47
 - when to add new, 49
- network interface cards, 5
- Network IP Address option, 51
- network management services (SNMP), 6
- network managers, SNMP, 309
- network management services, 121
- network monitoring services, 121
- network monitoring tools, 225
- Network News Transfer Protocol, 129
- networks
 - accessing configuration, 50
 - adding, 50
 - creating, 51
 - definition, 47
 - deleting, 53
 - modifying, 52
 - planning, 50
 - trusted, 2
 - types, 2

- understanding, 47
- unknown, 3
- untrusted, 3
- wildcard character, 48
- Network Source option (security rules), 37
- news (Usenet), 129
- news feed, informing, 134
- News proxy
 - configuring, 131
 - configuring firewall, 131
 - configuring news server, 134
 - enabling, 133
 - how it works, 130
 - informing news feed, 134
 - planning settings, 131
 - understanding, 130
 - using news, 135
- news proxy, using, 135
- News Reader settings, 133
- news server, configuring, 134
- NNTP
 - News Client, 190
 - News Server, 190
 - port, 130
 - protocol, 129
- nontransparent access
 - Web services, 175
- Norton AntiVirus for Firewalls, 290

O

- one-time password, 57, 62
- One-Time Password in Everything (OPIE), 62
- online help, 24
- operating system for Gauntlet Firewall, 5
- OPIE, 62
- order of precedence

- change for packet screening rules, 236
- changing, 46
- changing for service group rules, 38
- changing for user restriction rules, 77
- destination access rules, 40
- packet screening rules, 231
- service group rules, 35
- user restriction rules, 72

P

- packets
 - examined by Gauntlet Firewall, 13
 - processing, 12
 - receiving, 12
 - spoofed packets, 7
- packet screening, 225
 - absorb rules, 228
 - accessing configuration, 232
 - adding rules, 233
 - adding rules manually, 233
 - changing order of precedence, 236
 - changing rule order of precedence, 236
 - creating rules, 233
 - deleting rules, 235
 - deny rules, 227
 - destination IP address rules, 229
 - destination port rules, 231
 - field rules, 229
 - how it works, 227
 - how rules work, 227
 - interface rules, 230
 - modifying rules, 235
 - protocol rules, 230
 - source IP address rules, 229
 - source port rules, 231
 - understanding, 226
 - verifying configuration, 236
 - verifying your configuration, 236
- packet screening permit rules, 228

- packet screening rules, 227
 - deleting, 235
 - loading, 235
 - modifying, 235
 - order of precedence, 231
 - planning, 233
- Partial Nudity & Art category, 297
- passthrough link, creating, 270
- passwords
 - and untrusted service groups, 29
 - changing, 66
 - reusable, 63
- PC Extender for Windows 95, 18
- permit rules, 228
- persistent connections, 176
- ping, 228, 273
 - and SNMP, 308
- ping and SNMP proxy, 126
- planning circuit proxy firewall, 201
- planning configuration sets, 83
- planning content scanning, 290
- planning destination access rules, 45
- planning encryption keys, 265
- planning firewall for multimedia, 112
- planning firewall for SNMP, 124
- planning FTP proxy settings, 90
- planning Info Server, 280
- planning integrity database, 256
- planning LDAP proxy, 97
- planning lp proxy, 142
- planning network groups, 54
- planning networks, 50
- planning News settings, 131
- planning packet screening rules, 233
- planning plug proxy settings, 193
- planning POP3 proxy settings, 221
- planning rlogin proxy, 161
- planning rsh proxy settings, 147
- planning service groups, 32
- planning SMTP proxy, 216
- planning SNMP agent settings, 308
- planning Sybase proxy settings, 154
- planning TELENT proxy, 161
- planning the firewall, 18
- planning user restriction rules, 75
- planning VPNs, 268
- planning X11 proxy settings, 184
- plug-board proxy, 6
- plug-gw daemon, 191
- plug proxy
 - accessing configuration, 192
 - and NNTP traffic, 130
 - and TCP protocol, 81
 - configuring, 193
 - configuring firewall, 192
 - enabling, 195
 - how it works, 191
 - planning settings, 193
 - SSL plug proxy for SSL or SHTTP services, 170
 - TCP-based protocols, 193
 - UDP-based services, 189
 - understanding, 190
 - vs. circuit proxy, 197
 - wildcard characters, 194
- plug proxy risk assessment, 189
- plug proxy versions
 - AOL, 6
 - CompuServe, 6
 - finger, 6
 - Usenet news, 6
 - Web services, 6
 - whois, 6
- POP3, 6, 213
 - APOP authentication, 213

- pop3-gw daemon, 220
- POP3 proxy
 - accessing configuration, 220
 - authentication, 60
 - configuring, 221
 - configuring internal mail server, 223
 - configuring your internal POP3 mail server, 223
 - enabling services, 222
 - exchanging mail, 223
 - how it works, 220
 - internal mail server, 223
 - multiple servers, 224
 - planning settings, 221
 - understanding, 219
 - user authentication entries, 222
- Port option for service groups, 34
- possible items of interest section, 250
- Post Office Protocol Version 3, 213
- PostScript header, 282
- precedence
 - changing for service group rules, 38
 - packet screening rules, 231
- precedence for destination access rules, 40
- precedence for service group rules, 35
- precedence for user restriction rules, 72
- print client, configuring, 140
- printer port, 138
- Print Manager, 140
- print queue for lp proxy, 140
- print server
 - configuring, 140
- print services, 6, 137
- privacy without trust, 262
- privacy with trust, 262
- private links, 262
- processing packets
 - calling appropriate program, 13
 - checking request type, 13
 - checking source and destination, 13
 - overview, 12
 - process the request, 14
 - receive packet, 12
- processing requests, 14
- programs
 - called by firewall during processing, 13
- protecting integrity database, 257
- protocols
 - and proxies, 81
 - rules, 230
- proxies
 - and protocols, 81
 - circuit, 6
 - circuit proxy, 197
 - configuring, 82
 - creating configuration sets, 82
 - described, 5
 - FTP proxy, 87
 - Gopher, 168
 - logs created, 240
 - MediaBase, 207
 - multimedia, 109
 - plug, 6
 - plug proxy, 190
 - POP3, 219
 - rlogin, 158
 - SMTP, 213
 - SNMP, 121
 - strong user authentication, 7
 - supported, 6
 - Sybase, 151
 - TELNET, 158
 - understanding, 81
 - X11, 181
- proxy-aware web browsers, 176
- proxy services
 - user restrictions, 71

Q

Q files, 278
quarantine area, 290
queries to Info Server, 278
query files, 283
QuickTime movie header, 282
Quit menu command, 22
Quote of the Day service (in example), 192

R

RADIUS, 62, 356
random-challenge password, 60, 61
reading exception reports, 249
reading service summary reports, 249
RealAudio
 TCP port 7070, 110
RealAudio/RealVideo proxy
 RealPlayer, 116
RealAudio/RealVideo *See* RealPlayer
RealPlayer, 6, 116
RealPlayer proxy
 configuring, 112
 transparency, 115
 using, 115
Reboot menu command, 22, 23
receiving packets, 12
Remote Authentication Dial-In User Service, 62
remote execution, 6
remote print queue, 140
remote shell *See* rsh proxy
remote use of Gauntlet Firewall Manager, 16
replication server, 102
reporting
 security alerts, 247

 understanding, 239
report recipients, 244
reports
 accessing configuration, 244
 configuring, 244
 configuring frequency, 245
 configuring recipients, 244
 creating, 242
 events to ignore, 245
 exception reports, 242, 243
 frequency, 245
 log file, 248
 reading, 248
 service summary reports, 242, 243
requests
 checking type, 13
 processing, 14
request type, 13
resolution requirement, 20
restoring the firewall, 16
retention time for log, 241
reusable passwords, 63, 360
RFC 1213 (MIB-II), 309
risk assessment for plug proxy, 189
rlogin, 6, 159
rlogin-gw, 159
rlogin proxy
 accessing configuration, 160
 authentication management, 162
 configuring, 161, 162
 default rules, 159
 enabling, 162
 how it works, 159
 planning, 161
 understanding, 157, 158
 using terminal services, 163
 verifying setup, 162
 with authentication, 163
 without authentication, 163

- routing VNP packets, 265
 - rsh, 6
 - rsh-gw daemon, 146
 - rsh proxy
 - accessing configuration, 147
 - configuring, 148
 - configuring firewall, 147
 - configuring remote system, 149
 - default configuration, 146
 - enabling, 148
 - how it works, 146
 - planning settings, 147
 - time-out values, 148
 - understanding, 145
 - using, 149
 - verifying setup, 148
 - rules
 - and service groups, 4
 - creating destination access rules, 45
 - deleting destination access rules, 46
 - destination access, 40
 - modifying destination access rules, 46
 - order of precedence, 35, 40, 72
 - packet screening rules, 233
 - understanding service group rules, 35
 - user restriction, 72
 - running web servers, 180
- S**
- SafeWord, 61, 342
 - sample headers, 282
 - Satanic or Cult category, 298
 - Save, Apply, and Reboot menu command, 22, 23
 - Save and Apply menu command, 22, 23
 - Save menu command, 22, 23
 - saving changes, 22
 - scanning *See* content scanning
 - screen resolution requirement, 20
 - Search Engines category, 299
 - SecureNet Key, 346
 - SecurID, 62, 346
 - security, per-application, 5
 - security alerts, 247
 - security alerts section, 250
 - security breaches, reasons, 2
 - Security Dynamics, 62
 - security perimeter, 2
 - example, 10
 - trusted networks, 2
 - untrusted networks, 3
 - sendmail program, 213
 - bt, 219
 - servers
 - description, 17
 - running web servers, 180
 - X server, 182
 - service access, 39
 - Service Configurations option (security rules), 37
 - service group access, 39
 - service group rules
 - accessing configuration, 36
 - adding, 37
 - changing order, 38
 - deleting, 38
 - modifying, 38
 - order of precedence, 35
 - understanding, 35
 - service groups
 - accessing configuration, 31
 - Allow Password Change option, 34
 - configuring, 32
 - creating, 33
 - default, 4, 28
 - deleting, 35
 - described, 3

- destination restrictions, 34
- membership, 30
- modifying, 34
- planning, 32
- Port option, 34
- rules, 4
- trusted, 29
- understanding, 27
- untrusted, 29
- when to add new, 30
- services to be disabled, 5
- service summary reports, 242, 243
 - reading, 249
- Sex Education category, 298
- Sexual Acts/Text category, 297
- SHTTP services, 6, 170
- Simple Mail Transport Protocol, 213
- S/Key, 62, 351
- smap daemon, 214
- smapd daemon, 214
- smap SMTP client, 214
- SMTP proxy
 - configuring, 216
- SMTP, 6, 213
 - and content scanning, 285
- SMTP Mail
 - content scanning, 286
- SMTP proxy
 - accessing configuration, 215
 - client, 214
 - configuring, 216
 - configuring other settings, 218
 - enabling, 217
 - how it works, 214
 - planning, 216
 - understanding, 214
 - verifying, 218
- SNMP agent, 127
 - accessing configuration, 307
 - and MIB, 309
 - community, 309
 - configuring, 127
 - configuring firewall, 307
 - configuring settings, 308
 - configuring SNMP network managers, 309
 - enabling, 308
 - how it works, 306
 - object ID, 310
 - planning settings, 308
 - trap, 310
 - understanding, 306
- SNMP network managers
 - configuring, 309
- SNMP proxy
 - accessing configuration, 123
 - configuring, 125
 - configuring firewall, 124
 - configuring SNMP agent, 127
 - configuring SNMP agents, 127
 - enabling, 126
 - how it works, 122
 - ICMP traffic, 126
 - ping, 126
 - planning firewall, 124
 - trap requests, 123
 - understanding, 121
- SNMP requests, 122
- SNMP trap requests, 123
- software components of Gauntlet Firewall, 5
- source IP address rules, 229
- source port rules, 231
- sources
 - checking, 13
- specifying allowed activities, 3
- spoofing
 - IP-spoofing checks, 226
 - spoofed packets, 7

- Sports & Entertainment category, 299
 - SQL server proxy. *See* Microsoft SQL Server Proxy
 - SQL services, 6
 - SSL proxy
 - accessing configuration, 171
 - configuring, 174
 - enabling, 175
 - how it works, 168
 - transparent access, 176
 - SSL services, 170
 - standard firewall configuration, 9
 - standard printer port, 138
 - starting VNP, 272
 - stopping VNP, 273
 - StreamWorks, 6
 - configuring proxy, 112
 - player, 117
 - proxy, transparency, 117
 - StreamWorks, using proxy, 117
 - StreamWorks XDMA port (UDP port 1558), 110
 - strong authentication, 57
 - Access Key II, 60, 331
 - and proxies, 7
 - APOP, 60
 - circuit proxy, 198
 - CRYPTOCARD RB-I, 61, 334
 - Defender Security Server, 339
 - Digipass, 61, 339
 - HTTP proxy, 168
 - RADIUS, 62, 356
 - reusable passwords, 63, 360
 - SafeWord, 61, 342
 - SecureNet Key, 346
 - SecurID, 62, 346
 - S/Key, 62, 351
 - understanding, 60
 - web services, 178
 - supported proxies, 6
 - Sybase clients
 - configuring, 156
 - Sybase proxy
 - accessing configuration, 153
 - configuring, 154
 - configuring firewall, 154
 - configuring Sybase clients, 156
 - enabling, 155
 - how it works, 152
 - planning settings, 154
 - recommended configuration, 153
 - understanding, 151
 - verifying setup, 156
 - verifying your setup, 156
 - Sybase SQL, 6
 - syb-gw, 152
 - Symantec Norton AntiVirus for Firewalls, 291
 - system configuration, 20
 - system integrity
 - configuring checks, 254
 - creating an integrity database, 256
 - files to ignore in check, 255
 - how it works, 254
 - understanding, 253
 - updating the integrity database, 257
 - verifying, 257
 - System Manager
 - creating user accounts, 15
 - system warnings section, 250
- ## T
- TCP-based protocols for plug proxy, 193
 - TCP port 110, 220
 - TCP port 119, 130
 - TCP port 1755, 110
 - TCP port 25, 214
 - TCP port 44, 170
 - TCP port 513, 159

- TCP port 514, 146
 - TCP port 515, 138
 - TCP port 6000, 182, 183
 - TCP port 7, 170
 - TCP port 7000, 110
 - TCP port 7070, 110
 - TCP port 80, 170
 - TCP protocol and plug proxy, 81
 - TCP traffic
 - circuit proxy, 6
 - plug-board proxy, 6
 - TELNET, 6, 7
 - and circuit proxy, 199
 - and X11 proxy, 182
 - telnetd, 159
 - TELNET proxy
 - accessing configuration, 160
 - authentication management, 162
 - configuring, 162
 - configuring firewall, 161
 - default rules, 159
 - enabling, 162
 - how it works, 159
 - planning, 161
 - understanding, 157, 158
 - using terminal services, 163
 - verifying setup, 162
 - with authentication, 163
 - without authentication, 163
 - terminal services, 6, 157
 - testing VNPs, 272
 - text-based interface to Gauntlet Firewall, 16
 - Text header, 282
 - time-based password, 62
 - time-out values
 - rsh proxy, 148
 - TN3270
 - with authentication, 165
 - TN3270 requests, 158
 - tn-gw, 159
 - traceroute, 228
 - transparency
 - default for Gauntlet Firewall, 4
 - described, 4
 - HTTP access, 176
 - LDAP clients, 97
 - lp proxy, 140
 - NetShow proxy, 114
 - RealPlayer proxy, 115
 - SMTP proxy, 218
 - StreamWorks proxy (not supported), 117
 - VODLive proxy, 118
 - trap, 310
 - trap requests
 - SNMP proxy, 123
 - trusted links, 262
 - trusted network groups, 48
 - trusted networks, 2
 - and unknown networks, 3
 - transparency, 4
 - trusted service groups
 - default, 29
 - FTP, 88
 - trusted VNP
 - creating, 269
- ## U
- UDP and network monitoring, 225
 - UDP-based services and plug proxy, 189
 - UDP port 1558, 110
 - unknown host name, 42
 - Unknown keyword, 42
 - unknown networks, 3
 - untrusted network groups, 49

- untrusted networks, 3
 - and unknown networks, 3
 - transparency, 4
 - untrusted service groups
 - default, 29
 - HTTP access, 29
 - password, 29
 - updating integrity database, 257
 - upgrade information, 329
 - URL filtering
 - configuring, 294
 - understanding, 293
 - Usenet news, 7, 129
 - user accounts, 2
 - creating, 15
 - user authentication entries, 222
 - user authentication management system, 19, 57
 - user configuration, accessing, 63
 - user IDs, changing, 65
 - user names, changing, 65
 - user restriction rules, 72
 - order of precedence, 72
 - planning, 75
 - wildcard, 76
 - user restrictions, 71
 - accessing configuration, 73
 - changing order of precedence, 77
 - configuring rules, 75
 - creating rules, 75
 - deleting rules, 76
 - how they work, 72
 - modifying rules, 76
 - order of precedence, 72
 - supporting proxy services, 71
 - users
 - changing authentication methods, 66
 - changing group membership, 65
 - changing passwords, 66
 - configuring, 63
 - creating, 64
 - deleting, 69
 - disabling, 68
 - enabling, 68
 - in authentication management system, 59
 - modifying, 64
 - understanding, 59
 - using circuit proxy, 204
 - using FTP services, 91
 - using Gauntle Firewall Manager, 18
 - using Gopher services, 180
 - using Info Server, 284
 - using lp proxy, 144
 - using news, 135
 - using RealPlayer proxy, 115
 - using rsh services, 149
 - using StreamWorks proxy, 117
 - using VODLive proxy, 118
 - using web services, 175
 - using X11 services, 185
 - /usr/local/etc/infodb, 276
 - /usr/local/etc/ipe, 273
 - utilities
 - management utilities, 8
- ## V
- /var/adm/SYSLOG, 240
 - VASCO Data Security, 60
 - VDOLive
 - TCP port 7000, 110
 - VDOLive proxy
 - configuring, 112
 - using, 118
 - verifying circuit proxy setup, 204
 - verifying FTP proxy setup, 91
 - verifying Microsoft SQL Server proxy setup, 108

verifying multimedia proxy setup, 113
verifying packet screening configuration, 236
verifying rsh proxy setup, 148
verifying setup for rlogin proxy, 162
verifying setup for TELNET proxy, 162
verifying SMTP proxy, 218
verifying Sybase proxy setup, 156
verifying system integrity, 257
video
 MediaBase proxy, 207
Violence/Profanity category, 297
virtual private networks (VPN), 259
virtual X server, 182
virus checking *See* content scanning, 285
VODLive proxy
 transparency, 118
VPN (virtual private networks), 259
VPNs
 accessing encryption configuration, 265
 adding, 268
 configuring, 268
 deleting, 272
 deleting keys, 267
 encryption keys, 265
 how they work, 264
 modifying, 271
 modifying keys, 267
 passthrough link, 263
 planning, 268
 privacy without trust (private link), 262
 privacy with trust (trusted link), 262
 private link, 262
 routing packets, 265
 starting, 272
 stopping, 273
 testing, 272
 trusted link, 262
 understanding, 260

W

warnings
 packet screening, 225
 packet screening rules, 233
 using both interfaces, 17
 using proprietary protocols, 189
weak authentication, 178
web browser
 configuring, 20
web browsers
 configuring, 176
 configuring non-proxy-aware browsers, 179
 proxy aware, 176
web servers, running, 180
Gopher proxy
web services, 6
 accessing for non-proxy-aware browsers, 179
 accessing with authentication, 178
 accessing without authentication, 178
 Info Server, 275
 See Also HTTP proxy, 175
 using, 175
Web services (SSL), 7
whois, 7, 190
wildcard
 in user restriction rules, 76
wildcard character, 48
wildcard characters
 plug proxy, 194
wildcards
 in SNMP proxy configuration, 125
Work Time, 303

X

X11, 6
X11 proxy

- accessing configuration, 183
- and TELNET, 182
- configuring, 184
- configuring firewall, 183
- enabling, 185
- how it works, 182
- planning, 184
- understanding, 181
- using X11 services, 185
- X port, 183
- X client, 182
- X port, 183
- X server, 182
- xurl-encoded characters
 - filtering, 293
- X Window services, 6

Z

- ZIP header, 282

Tell Us About This Manual

As a user of Silicon Graphics products, you can help us to better understand your needs and to improve the quality of our documentation.

Any information that you provide will be useful. Here is a list of suggested topics:

- General impression of the document
- Omission of material that you expected to find
- Technical errors
- Relevance of the material to the job you had to do
- Quality of the printing and binding

Please send the title and part number of the document with your comments. The part number for this document is 007-2826-007.

Thank you!

Three Ways to Reach Us

- To send your comments by **electronic mail**, use either of these addresses:
 - On the Internet: techpubs@sgi.com
 - For UUCP mail (through any backbone site): *[your_site]!sgi!techpubs*
- To **fax** your comments (or annotated copies of manual pages), use this fax number: 650-932-0801
- To send your comments by **traditional mail**, use this address:

Technical Publications
Silicon Graphics, Inc.
2011 North Shoreline Boulevard, M/S 535
Mountain View, California 94043-1389