NetScreen Concepts & Examples

ScreenOS Reference Guide

Volume 1: Overview

ScreenOS 5.0.0

P/N 093-0924-000

Rev. B

Copyright Notice

NetScreen Technologies, Inc.
Building #3
805 11th Avenue
Sunnyvale, CA 94089
www.netscreen.com

FCC Statement

The following information is for FCC compliance of Class A devices: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. The equipment generates, uses, and can radiate radio-frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case users will be required to correct the interference at their own expense.

The following information is for FCC compliance of Class B devices: The equipment described in this manual generates and may radiate radio-frequency energy. If it is not installed in accordance with NetScreen's installation instructions, it may cause interference with radio and television reception. This equipment has been tested and found to comply with the limits for a Class B digital device in accordance with the specifications in part 15 of the FCC rules. These specifications are designed to provide reasonable protection against such interference in a residential installation. However, there is no guarantee that interference will not occur in a particular installation.

If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Consult the dealer or an experienced radio/TV technician for help.
- Connect the equipment to an outlet on a circuit different from that to which the receiver is connected.

Caution: Changes or modifications to this product could void the user's warranty and authority to operate this device.

Disclaimer

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR NETSCREEN REPRESENTATIVE FOR A COPY.

Contents

Volume 1: Overview

Contents i
Preface xxi
Concepts & Examples Organization xxiii
Conventions xxvii
CLI Conventions xxvii
WebUI Conventions xxviii
Illustration Conventions xxx
Naming Conventions and Character Types xxxi
NetScreen Documentation xxxii
Appendix A Glossary ..... A-I
Index ..... IX-I

Volume 2: Fundamentals

Contents i
Preface ix
Conventions x
CLI Conventions x
WebUI Conventions xi
Illustration Conventions xiii
Naming Conventions and Character Types xiv
NetScreen Documentation xv
Chapter 1 ScreenOS Architecture 1
Security Zones 2
Security Zone Interfaces 3
Physical Interfaces 3
Subinterfaces 4
Virtual Routers 5
Policies 6
VPNs 8
Virtual Systems 10
Packet Flow Sequence 11
Example (Part 1): Enterprise with Six Zones 14
Example (Part 2): Interfaces for Six Zones 16
Example (Part 3): Enterprise with Two Routing Domains 20
Example (Part 4): Policies for an Enterprise with Six Zones 22
Chapter 2 Routing Tables and Static Routing 29
Routing Essentials 30
Routing Methods 30
Static Routing 30
Dynamic Routing 30
Routing Tables 31

Routing with Static Routes 33
Virtual Routers on NetScreen Devices 35
When to Configure Static Routes 36
Configuring Static Routes 38
Example: Configuring Static Routes 39
Example: Static Route through a Tunnel Interface 43
Chapter 3 Zones 45
Security Zones 48
Global Zone 48
SCREEN Options 48
Tunnel Zones 49
Example: Binding a Tunnel Interface to a Tunnel Zone 50
Configuring Security Zones and Tunnel Zones 51
Creating a Zone 51
Modifying a Zone 52
Deleting a Zone 53
Function Zones 54
Null Zone 54
MGT Zone 54
HA Zone 54
Self Zone 54
VLAN Zone 54
Port Modes 55
Setting the Port Mode on NetScreen Appliances 59
Example: Setting Home-Work Port Mode 60
Home Zone/Work Zone 61
Example: Configuring Home and Work Zones 63
Chapter 4 Interfaces 65
Interface Types 66
Security Zone Interfaces 66
Physical 66
Subinterface 66
Aggregate Interfaces 67
Redundant Interfaces 67
Virtual Security Interfaces 68
Function Zone Interfaces 68
Management Interface 68
HA Interface 69
Tunnel Interfaces 69
Deleting Tunnel Interfaces 72
Example: Deleting a Tunnel Interface 72
Viewing Interfaces 74
Interface Table 74
Configuring Security Zone Interfaces 76
Binding an Interface to a Security Zone 76
Example: Binding an Interface 76
Defining an Address for a L3 Security Zone Interface 77
Public IP Addresses 77
Private IP Addresses 78
Example: Addressing an Interface 79
Unbinding an Interface from a Security Zone 80
Example: Unbinding an Interface 80
Modifying Interfaces 81
Example: Modifying Settings on an Interface 81
Creating Subinterfaces 82
Example: Creating a Subinterface in the Root System 82

Deleting Subinterfaces....83
Example: Deleting a Security Zone Interface....83
Secondary IP Addresses....84
Secondary IP Address Properties....84
Example: Creating a Secondary IP Address....85
Loopback Interfaces....86
Example: Creating a Loopback Interface....86
Using Loopback Interfaces....87
Example: Using the Loopback Interface to Manage a Device....87
Example: Enabling BGP on a Loopback Interface....88
Example: Configuring NSRP VSIs on a Loopback Interface....88
Example: Specifying a Loopback Interface as a Source Interface....89
Chapter 5 Interface Modes....91
Transparent Mode....92
Zone Settings....93
VLAN Zone....93
Predefined Layer 2 Zones....93
Traffic Forwarding....94
Unknown Unicast Options....95
Flood Method....96
ARP/Trace-Route Method....98
Example: VLAN1 Interface for Management....102
Example: Transparent Mode....105
NAT Mode....110
Inbound and Outbound NAT Traffic....112
Interface Settings....113
Example: NAT Mode....114
Route Mode....118
Interface Settings....119
Example: Route Mode....120
Chapter 6 Building Blocks for Policies....125
Addresses....126
Address Entries....127
Example: Adding Addresses....127
Example: Modifying Addresses....128
Example: Deleting Addresses....129
Address Groups....129
Example: Creating an Address Group....131
Example: Editing a Group Address Entry....132
Example: Removing an Address Group Member and a Group....133
Services....134
Predefined Services....134
Example: Setting a Predefined Service Timeout....136
Custom Services....136
Example: Adding a Custom Service....136
Example: Modifying a Custom Service....138
Example: Removing a Custom Service....138
ICMP Services....139
Example: Defining an ICMP Service....140
RSH ALG....140
H.323 Protocol for Voice-over-IP....141
Example: Gatekeeper in the Trust Zone (Transparent or Route Mode)....141
Example: Gatekeeper in the Trust Zone (NAT Mode)....143
Example: Gatekeeper in the Untrust Zone (Transparent or Route Mode)....148

Example: Gatekeeper in the Untrust Zone (NAT Mode) 151
SIP – Session Initiation Protocol 156
SIP Request Methods 157
Classes of SIP Responses 157
ALG – Application-Layer Gateway 159
SDP 160
Pinhole Creation 161
Session Inactivity Timeout 163
Example: Creating a Policy to Permit SIP 164
Example: Signaling and Media Inactivity Timeouts 166
Service Groups 167
Example: Creating a Service Group 168
Example: Modifying a Service Group 169
Example: Removing a Service Group 170
DIP Pools 171
Port Address Translation 172
Example: Creating a DIP Pool with PAT 172
Example: Modifying a DIP Pool 174
Sticky DIP Addresses 174
Extended Interface and DIP 175
Example: Using DIP in a Different Subnet 175
Loopback Interface and DIP 183
Example: DIP on a Loopback Interface 184
DIP Groups 189
Example: DIP Group 191
Schedules 193
Example: Recurring Schedule 193
Chapter 7 Policies 197
Basic Elements 199
Three Types of Policies 200
Interzone Policies 200
Intrazone Policies 201
Global Policies 201
Policy Set Lists 202
Policies Defined 203
Policies and Rules 203
Anatomy of a Policy 205
ID 206
Zones 206
Addresses 206
Services 206
Action 207
Application 207
Name 208
VPN Tunneling 208
L2TP Tunneling 209
Deep Inspection 209
Placement at the Top of the Policy List 209
Source Address Translation 210
Destination Address Translation 210
User Authentication 210
HA Session Backup 212
URL Filtering 213
Logging 213
Counting 213
Traffic Alarm Threshold 213
Schedules 214
Antivirus Scanning 214
Traffic Shaping 215
Policies Applied 216
Viewing Policies 216

Policy Icons 216
Creating Policies 217
Policy Location 218
Example: Interzone Policies for E-Mail Service 218
Example: Interzone Policy Set 223
Example: Intrazone Policies 231
Example: Global Policy 234
Entering a Policy Context 235
Multiple Items per Policy Component 236
Address Negation 237
Example: Destination Address Negation 237
Modifying and Disabling Policies 241
Policy Verification 242
Reordering Policies 243
Removing a Policy 244

Chapter 8 Address Translation 245
Introduction to Address Translation 246
Policy-Based Translation Options 253
Directional Nature of NAT-Src and NAT-Dst 257
Source Network Address Translation 259
NAT-Src from a DIP Pool with PAT Enabled 260
Example: NAT-Src with PAT Enabled 261
NAT-Src from a DIP Pool with PAT Disabled 264
Example: NAT-Src with PAT Disabled 264
NAT-Src from a DIP Pool with Address Shifting 267
Example: NAT-Src with Address Shifting 268
NAT-Src from the Egress Interface IP Address 273
Example: NAT-Src without DIP 273
Destination Network Address Translation 276
Packet Flow for Destination Translation 278
Routing for Destination Translation 282
Addresses Connected to the Same Interface 283
Addresses Connected to the Same Interface but Separated by a Router 284
Addresses Separated by an Interface 285
NAT-Dst: One-to-One Mapping 286
Example: One-to-One Destination Translation 287
Translating from One Address to Multiple Addresses 291
Example: One-to-Many Destination Translation 291
NAT-Dst: Many-to-One Mapping 295
Example: Many-to-One Destination Translation 295
NAT-Dst: Many-to-Many Mapping 300
Example: Many-to-Many Destination Translation 301
NAT-Dst with Port Mapping 305
Example: NAT-Dst with Port Mapping 305
NAT-Src and NAT-Dst in the Same Policy 310
Example: NAT-Src and NAT-Dst Combined 310
Mapped IP Addresses 331
MIP and the Global Zone 332
Example: Adding a MIP to an Untrust Zone Interface 333
Example: Reaching a MIP from Different Zones 336
Example: Adding a MIP to a Tunnel Interface 341
MIP-Same-as-Untrust 342
Example: MIP on the Untrust Interface 343
MIP and the Loopback Interface 346

Example: MIP for Two Tunnel Interfaces 347
Virtual IP Addresses 356
VIP and the Global Zone 359
Example: Configuring Virtual IP Servers 359
Example: Editing a VIP Configuration 362
Example: Removing a VIP Configuration 362
Example: VIP with Custom and Multiple-Port Services 363
Chapter 9 User Authentication 371
Authentication Servers 372
Local Database 374
Supported User Types and Features 374
Example: Setting the Local Database Timeout 375
External Auth Servers 376
Auth Server Object Properties 377
Auth Server Types 379
RADIUS 379
RADIUS Auth Server Object Properties 380
Supported User Types and Features 380
NetScreen Dictionary File 381
RADIUS Access-Challenge 382
SecurID 384
SecurID Auth Server Object Properties 385
Supported User Types and Features 385
LDAP 386
LDAP Auth Server Object Properties 387
Supported User Types and Features 387
Defining Auth Server Objects 388
Example: Defining an Auth Server Object for RADIUS 388
Example: Defining an Auth Server Object for SecurID 391
Example: Defining an Auth Server Object for LDAP 393
Defining Default Auth Servers 395
Example: Changing the Default Auth Servers 395
Authentication Types and Applications 397
Auth Users and User Groups 398
Referencing Auth Users in Policies 398
Referencing Auth User Groups in Policies 402
Example: Run-Time Authentication (Local User) 403
Example: Run-Time Authentication (Local User Group) 406
Example: Run-Time Authentication (External User) 409
Example: Run-Time Authentication (External User Group) 412
Example: Local Auth User in Multiple Groups 416
Example: WebAuth (Local User Group) 420
Example: WebAuth (External User Group) 423
Example: WebAuth + SSL (External User Group) 427
IKE Users and User Groups 431
Example: Defining IKE Users 432
Example: Creating an IKE User Group 434
Referencing IKE Users in Gateways 435
XAuth Users and User Groups 436
XAuth Users in IKE Negotiations 437
Example: XAuth Authentication (Local User) 440

Example: XAuth Authentication (Local User Group) 442
Example: XAuth Authentication (External User) 444
Example: XAuth Authentication (External User Group) 447
Example: XAuth Authentication and Address Assignments (Local User Group) 452
XAuth Client 458
Example: NetScreen Device as an XAuth Client 459
L2TP Users and User Groups 460
Example: Local and External L2TP Auth Servers 461
Admin Users 465
Multiple-Type Users 467
Group Expressions 468
Example: Group Expressions (AND) 470
Example: Group Expressions (OR) 472
Example: Group Expressions (NOT) 474
Banner Customization 476
Example: Customizing the WebAuth Success Message 476
Chapter 10 Traffic Shaping 477
Applying Traffic Shaping 478
Managing Bandwidth at the Policy Level 478
Example: Traffic Shaping 479
Setting Service Priorities 485
Example: Priority Queuing 486
Chapter 11 System Parameters 493
Domain Name System Support 495
DNS Lookup 496
DNS Status Table 497
Example: Defining DNS Server Addresses and Scheduling Lookups 498
Example: Setting a DNS Refresh Interval 499
DHCP 500
DHCP Server 502
Example: NetScreen Device as DHCP Server 502
DHCP Server in an NSRP Cluster 508
DHCP Server Detection 508
Example: Turning on DHCP Server Detection 509
Example: Turning off DHCP Server Detection 509
DHCP Relay Agent 510
Example: NetScreen Device as DHCP Relay Agent 511
DHCP Client 516
Example: NetScreen Device as DHCP Client 516
TCP/IP Settings Propagation 518
Example: Forwarding TCP/IP Settings 519
PPPoE 521
Example: Setting Up PPPoE 521
Example: Configuring PPPoE on Primary and Backup Untrust Interfaces 526
Downloading/Uploading Settings and Software 528
Saving and Importing Settings 528
Uploading and Downloading Software 530
Configuration Rollback 531
Last-Known-Good Configuration 531

Automatic and Manual Configuration Rollback 531
Loading a New Configuration File 533
Locking the Configuration File 534
Adding Comments to a Configuration File 535
License Keys 536
Example: Expanding User Capacity 537
Registration and Activation of Signature Services 538
Temporary Service 538
AV and DI Bundled with a New Device 538
AV Upgrade with DI 539
DI Upgrade Only 540
System Clock 541
Date and Time 541
Time Zone 541
NTP 542
Multiple NTP Servers 542
Maximum Time Adjustment 542
NTP and NSRP 543
Example: Configuring NTP Servers and a Maximum Time Adjustment Value 544
Secure NTP Servers 545
Index ..... IX-I

Volume 3: Administration

Contents i
Preface v
Conventions vi
CLI Conventions vi
WebUI Conventions vii
Illustration Conventions ix
Naming Conventions and Character Types x
NetScreen Documentation xi
Chapter 1 Administration 1
Management via the Web User Interface 3
WebUI Help 4
Copying the Help Files to a Local Drive 4
Pointing the WebUI to the New Help Location 4
HTTP 5
Session ID 5
Secure Sockets Layer 7
Management via the Command Line Interface 9
Telnet 9
Securing Telnet Connections 10
Secure Shell 11
Client Requirements 13
Basic SSH Configuration on the NetScreen Device 13
Authentication 15
SSH and Vsys 17
Host Key 18
Example: SSHv1 with PKA for Automated Logins 19
Secure Copy (SCP) 20
Serial Console 21
Modem Port 22
Management via NetScreen-Security Manager 23

Initiating Connectivity Between Agent and Management System 24
Enabling and Disabling the Agent 25
Example: Enabling the Security Manager Agent 25
Changing Management System Server Address 26
Example: Setting the Primary Server IP Address 26
Setting Report Parameters 26
Example: Enabling Alarm and Statistics Reporting 27
Controlling Administrative Traffic 29
MGT and VLAN1 Interfaces 30
Example: Administration through the MGT Interface 30
Example: Administration through the VLAN1 Interface 31
Administrative Interface 32
Example: Setting Administrative Interface Options 32
Manage IP 34
Example: Setting Manage IPs for Multiple Interfaces 34
Levels of Administration 37
Root Administrator 37
Read/Write Administrator 38
Read-Only Administrator 38
Virtual System Administrator 38
Virtual System Read-Only Administrator 39
Defining Admin Users 39
Example: Adding a Read-Only Admin 39
Example: Modifying an Admin 40
Example: Deleting an Admin 40
Example: Clearing an Admin's Sessions 41
Securing Administrative Traffic 42
Changing the Port Number 43
Example: Changing the Port Number 43
Changing the Admin Login Name and Password 44
Example: Changing an Admin User's Login Name and Password 45
Example: Changing One's Own Password 46
Setting the Minimum Length of the Root Admin Password 47
Resetting the Device to the Factory Default Settings 48
Restricting Administrative Access 49
Example: Restricting Administration to a Single Workstation 49
Example: Restricting Administration to a Subnet 50
Restricting the Root Admin to Console Access 50
VPN Tunnels for Administrative Traffic 51
Example: Administration through a Route-Based Manual Key VPN Tunnel 52
Example: Administration through a Policy-Based Manual Key VPN Tunnel 58
Chapter 2 Monitoring NetScreen Devices 65
Storing Log Information 66
Event Log 67
Viewing the Event Log 68
Example: Viewing the Event Log by Severity Level and Keyword 69
Sorting and Filtering the Event Log 70
Example: Sorting Event Log Entries by IP Address 70
Downloading the Event Log 71

Example: Downloading the Event Log 71
Example: Downloading the Event Log for Critical Events 71
Traffic Log 72
Viewing the Traffic Log 74
Example: Viewing Traffic Log Entries 74
Sorting and Filtering the Traffic Log 75
Example: Sorting the Traffic Log by Time 75
Downloading the Traffic Log 76
Example: Downloading a Traffic Log 76
Self Log 77
Viewing the Self Log 77
Sorting and Filtering the Self Log 78
Example: Filtering the Self Log by Time 79
Downloading the Self Log 80
Example: Downloading the Self Log 80
Asset Recovery Log 81
Example: Downloading the Asset Recovery Log 81
Traffic Alarms 82
Example: Policy-Based Intrusion Detection 83
Example: Compromised System Notification 84
Example: Sending E-mail Alerts 86
Syslog 87
Example: Enabling Multiple Syslog Servers 88
WebTrends 89
Example: Enabling Syslog and WebTrends for Notification Events 89
SNMP 91
Implementation Overview 94
Example: Defining a Read/Write SNMP Community 95
VPN Tunnels for Self-Initiated Traffic 97
Example: Self-Generated Traffic through a Route-Based Tunnel 99
Example: Self-Generated Traffic through a Policy-Based Tunnel 109
Counters 120
Example: Viewing Screen Counters 126
Appendix A SNMP MIB Files ..... A-I
Index ..... IX-I

Volume 4: Attack Detection and Defense Mechanisms

Contents i
Preface v
Conventions vi
CLI Conventions vi
WebUI Conventions vii
Illustration Conventions ix
Naming Conventions and Character Types x
NetScreen Documentation xi
Chapter 1 Protecting a Network 1
Stages of an Attack 2
Detection and Defense Mechanisms 3
Exploit Monitoring 5

Example: Monitoring Attacks from the Untrust Zone 6
Chapter 2 Reconnaissance Deterrence 7
IP Address Sweep 8
Port Scanning 10
Network Reconnaissance Using IP Options 12
Operating System Probes 16
SYN and FIN Flags Set 16
FIN Flag without ACK Flag 18
TCP Header without Flags Set 20
Evasion Techniques 22
FIN Scan 22
IP Spoofing 22
Example: L3 IP Spoof Protection 25
Example: L2 IP Spoof Protection 29
IP Source Route Options 31
Chapter 3 Denial-of-Service Attack Defenses 35
Firewall DoS Attacks 36
Session Table Flood 36
Source- and Destination-Based Session Limits 36
Example: Source-Based Session Limiting 39
Example: Destination-Based Session Limiting 40
Aggressive Aging 40
Example: Aggressively Aging Out Sessions 42
SYN-ACK-ACK Proxy Flood 43
Network DoS Attacks 45
SYN Flood 45
Example: SYN Flood Protection 52
ICMP Flood 59
UDP Flood 61
Land Attack 63
OS-Specific DoS Attacks 65
Ping of Death 65
Teardrop Attack 67
WinNuke 69
Chapter 4 Content Monitoring and Filtering 71
Fragment Reassembly 72
Malicious URL Protection 72
Application Layer Gateway 73
Example: Blocking Malicious URLs in Packet Fragments 74
Antivirus Scanning 76
Internal AV Scanning 77
Enabling Internal AV Scanning 81
Updating the Pattern File Automatically or Semi-Automatically 82
Example: Automatic Pattern Update 83
Example: Semi-Automatic Pattern Update 83
Configuring Content Processing 84
Example: Internal AV Scanning for SMTP 84
Example: Internal AV Scanning for SMTP and HTTP 85
Configuring Decompression and Maximum Content Size 85
Example: Dropping Large Files 86
Applying Internal AV Scanning 87
Example: Internal AV Scanning (POP3) 87
External AV Scanning 90
Defining AV Objects 93
Example: Defining Three AV Objects 99

Applying External AV Scanning 102
Example: Antivirus with One AV Object 103
Example: Antivirus with Two AV Objects 106
URL Filtering 113
Example: URL Filtering Configuration 119
Chapter 5 Deep Inspection 123
Deep Inspection Overview 124
Attack Object Database Server 128
Example: Immediate Update 129
Example: Automatic Updates 130
Example: Automatic Notification and Immediate Update 132
Example: Manual Update 134
Attack Objects and Groups 136
Stateful Signatures 138
TCP Stream Signatures 139
Protocol Anomalies 139
Attack Object Groups 140
Changing Severity Levels 140
Attack Actions 142
Example: Attack Actions – Close Server, Close, Close Client 143
Mapping Custom Services to Applications 152
Example: Mapping an Application to a Custom Service 153
Customized Attack Objects and Groups 156
User-Defined Stateful SIgnature Attack Objects 156
Contexts 156
Signatures 157
Example: User-Defined Stateful Signature Attack Objects 160
TCP Stream Signature Attack Objects 164
Example: User-Defined Stream Signature Attack Object 164
Granular Blocking of HTTP Components 167
ActiveX Controls 167
Java Applets 168
EXE Files 168
ZIP Files 168
Example: Blocking Java Applets and .exe Files 169
Chapter 1 Suspicious Packet Attributes 1
ICMP Fragments 2
Large ICMP Packets 4
Bad IP Options 6
Unknown Protocols 8
IP Packet Fragments 10
SYN Fragments 12
Index ..... IX--I

Volume 5: VPNs

Contents i
Preface v
Conventions vi
CLI Conventions vi
WebUI Conventions vii

Illustration Conventions ix
Naming Conventions and Character Types x
NetScreen Documentation xi
Chapter 1 IPSec 1
Introduction to VPNs 2
IPSec Concepts 3
Modes 4
Transport Mode 4
Tunnel Mode 5
Protocols 7
AH 7
ESP 8
Key Management 9
Manual Key 9
AutoKey IKE 9
Security Association 10
Tunnel Negotiation 11
Phase 1 11
Main Mode and Aggressive Mode 12
The Diffie-Hellman Exchange 13
Phase 2 13
Perfect Forward Secrecy 14
Replay Protection 14
Chapter 2 Public Key Cryptography 15
Introduction to Public Key Cryptography 16
PKI 18
Certificates and CRLs 21
Obtaining a Certificate Manually 22
Example: Requesting a Certificate Manually 23
Example: Loading Certificates and CRLs 26
Example: Configuring CRL Settings for a CA Certificate 28
Obtaining a Local Certificate Automatically 30
Example: Requesting a Local Certificate Automatically 31
Automatic Certificate Renewal 34
Key Pair Generation 35
Checking for Revocation Using OCSP 36
Configuring for OCSP 37
Specifying either CRL or OCSP for Revocation Checking 37
Displaying Certificate Revocation Status Attributes 37
Specifying the URL of an OCSP Responder for a Certificate 38
Removing Certificate Revocation Check Attributes 38
Chapter 3 VPN Guidelines 39
Cryptographic Options 40
Site-to-Site Cryptographic Options 41
Dialup VPN Options 50
Route- and Policy-Based Tunnels 58
Packet Flow: Site-to-Site VPN 60
Tunnel Configuration Tips 67
Chapter 4 Site-to-Site VPNs 69
Site-to-Site VPN Configurations 70
Site-to-Site Tunnel Configuration Steps 71
Example: Route-Based Site-to-Site VPN, AutoKey IKE 77
Example: Policy-Based Site-to-Site VPN, AutoKey IKE 91

Example: Route-Based Site-to-Site VPN, Dynamic Peer 102
Example: Policy-Based Site-to-Site VPN, Dynamic Peer 117
Example: Route-Based Site-to-Site VPN, Manual Key 131
Example: Policy-Based Site-to-Site VPN, Manual Key 142
FQDN for Dynamic IKE Gateways 151
Aliases 152
Example: AutoKey IKE Peer with FQDN 153
VPN Sites with Overlapping Addresses 168
Example: Tunnel Interface with NAT-Src and NAT-Dst 171
Transparent Mode VPN 186
Example: Transparent Mode, Policy-Based AutoKey IKE VPN 187
Chapter 5 Dialup VPNs 199
Dialup VPNs 200
Example: Policy-Based Dialup VPN, AutoKey IKE 201
Example: Route-Based Dialup VPN, Dynamic Peer 209
Example: Policy-Based Dialup VPN, Dynamic Peer 220
Bidirectional Policies for Dialup VPN Users 229
Example: Bidirectional Dialup VPN Policies 230
Group IKE ID 237
Group IKE ID with Certificates 238
Wildcard and Container ASN1-DN IKE ID Types 240
Example: Group IKE ID (Certificates) 243
Group IKE ID with Preshared Keys 250
Example: Group IKE ID (Preshared Keys) 252
Shared IKE IDs 259
Example: Shared IKE ID (Preshared Keys) 260
Chapter 6 L2TP 269
Introduction to L2TP 270
Packet Encapsulation and Decapsulation 274
Encapsulation 274
Decapsulation 275
L2TP Parameters 276
Example: Configuring an IP Pool and L2TP Default Settings 277
L2TP and L2TP-over-IPSec 279
Example: Configuring L2TP 280
Example: Configuring L2TP-over-IPSec 286
Chapter 7 Advanced VPN Features 299
IPSec NAT Traversal 301
Traversing a NAT Device 302
UDP Checksum 303
The Keepalive Frequency Value 303
IPSec NAT-Traversal and Initiator/Responder Symmetry 304
Example: Enabling NAT-Traversal 305
VPN Monitoring 307
Rekey and Optimization Options 307
Source Interface and Destination Address 308
Policy Considerations 310
Configuring the VPN Monitoring Feature 310
Example: Specifying Source and Destination Addresses for VPN Monitoring 312

Security Consideration for a Route-Based VPN Design 323
SNMP VPN Monitoring Objects and Traps 325
Multiple Tunnels per Tunnel Interface 326
Route-to-Tunnel Mapping 327
Remote Peers' Addresses 328
Manual and Automatic Table Entries 330
Manual Table Entries 330
Automatic Table Entries 331
Example: Multiple VPNs on One Tunnel Interface to Overlapping Subnets 333
Example: Automatic Route and NHTB Table Entries 364
Redundant VPN Gateways 382
VPN Groups 383
Monitoring Mechanisms 384
IKE Heartbeats 384
IKE Recovery Procedure 385
TCP SYN-Flag Checking 388
Example: Redundant VPN Gateways 389
Back-to-Back VPNs 401
Example: Back-to-Back VPNs 402
Hub-and-Spoke VPNs 412
Example: Hub-and-Spoke VPNs 413
Index ..... IX-I

Volume 6: Dynamic Routing

Contents i
Preface v
Conventions vi
CLI Conventions vi
WebUI Conventions vii
Illustration Conventions ix
Naming Conventions and Character Types x
NetScreen Documentation xi
Chapter 1 Virtual Routers 1
Virtual Routers on NetScreen Devices 3
Using Two VRs 3
Forwarding Traffic between VRs 4
Configuring Two Virtual Routers 4
Example: Binding a Zone to the untrust-vr 5
Custom Virtual Routers 7
Example: Creating a Custom Virtual Router 7
Example: Removing a Custom Virtual Router 8
Virtual Routers and Virtual Systems 9
Example: Creating a Custom Virtual Router in a vsys 10
Example: Defining a Route with a Shared Virtual Router as the Next-Hop 11
Modifying Virtual Routers 12
Virtual Router ID 12
Example: Assigning a Virtual Router ID 13
Maximum Number of Routing Table Entries 14
Example: Limiting the Maximum Number of Routing Table Entries 14
Route Selection 15
Route Preference 15

Example: Setting a Route Preference 16
Route Metric 17
Source-Based Routing 17
Example: Source-Based Routing 19
Route Redistribution 21
Configuring a Route Map 22
Route Filtering 24
Access Lists 24
Example: Configuring an Access List 25
Example: Redistributing BGP Routes into OSPF 26
Exporting and Importing Routes between VRs 28
Example: Configuring a Route Export Rule 29
Chapter 2 Open Shortest Path First (OSPF) 33
Overview of OSPF 34
Areas 34
Router Classification 35
Hello Protocol 35
Network Types 36
Broadcast Networks 36
Point-to-Point Networks 36
Link State Advertisements 37
Basic OSPF Configuration 38
Creating an OSPF Routing Instance in a Virtual Router 39
Example: Creating an OSPF Routing Instance 39
Example: Removing an OSPF Routing Instance 40
Defining an OSPF Area 41
Example: Creating an OSPF Area 41
Assigning Interfaces to an OSPF Area 42
Example: Assigning Interfaces to OSPF Areas 42
Example: Configuring an Area Range 43
Enabling OSPF on Interfaces 44
Example: Enabling OSPF on Interfaces 44
Example: Disabling OSPF on an Interface 45
Verifying the Configuration 46
Redistributing Routes 49
Example: Redistributing a BGP Route into OSPF 49
Summarizing Redistributed Routes 50
Example: Summarizing Redistributed Routes 50
Global OSPF Parameters 51
Example: Advertising the Default Route 52
Virtual Links 53
Example: Creating a Virtual Link 54
Example: Creating an Automatic Virtual Link 56
OSPF Interface Parameters 57
Example: Setting OSPF Interface Parameters 59
Security Configuration 60
Authenticating Neighbors 60
Example: Configuring the Clear-Text Password Authentication Method 60
Example: Configuring the MD5 Password Authentication Method 61
Filtering OSPF Neighbors 62
Example: Configuring a Neighbor List 62
Rejecting Default Routes 63
Example: Removing the Default Route from the Route Table 63
Protecting against Flooding 64
Example: Configuring the Hello Threshold 64

Example: Configuring the LSA Threshold 65

Chapter 3 Routing Information Protocol (RIP) 67

Overview of RIP 68

Basic RIP Configuration 69

Creating a RIP Routing Instance in a Virtual Router 70

Example: Creating a RIP Routing Instance 70

Example: Removing a RIP Routing Instance 71

Enabling RIP on Interfaces 72

Example: Enabling RIP on Interfaces 72

Example: Disabling RIP on an Interface 73

Redistributing Routes 73

Example: Redistributing Routes into RIP 74

Global RIP Parameters 76

Example: Advertising the Default Route to RIP Neighbors 77

RIP Interface Parameters 78

Example: Setting RIP Interface Parameters 79

Security Configuration 80

Authenticating Neighbors 80

Example: Configuring the MD5 Password Authentication Method 81

Filtering RIP Neighbors 82

Example: Configuring Trusted Neighbors 82

Rejecting Default Routes 83

Example: Rejecting Default Routes 83

Protecting Against Flooding 84

Example: Configuring an Update Threshold 84

Example: RIP on Tunnel Interfaces 85

Chapter 4 Border Gateway Protocol (BGP) 87

Overview of BGP 88

Types of BGP Messages 89

Path Attributes 89

External and Internal BGP 90

Basic BGP Configuration 91

Creating and Enabling a BGP Routing Instance in a Virtual Router 92

Example: Creating a BGP Routing Instance 92

Example: Removing a BGP Routing Instance 93

Enabling BGP on Interfaces 94

Example: Enabling BGP on Interfaces 94

Example: Disabling BGP on Interfaces 94

Configuring a BGP Peer 95

Example: Configuring a BGP Peer 97

Example: Configuring an IBGP Peer-Group 98

Verifying the BGP Configuration 100

Security Configuration 102

Authenticating Neighbors 102

Example: Configuring MD5 Authentication for BGP Peers 102

Rejecting Default Routes 103

Example: Rejecting Default Routes 103

Optional BGP Configurations 104

Redistributing Routes 105

Example: Redistributing an OSPF Route into BGP 105

AS-Path Access List 106

Example: Configuring an AS-Path Access List 106

Route Reflection 107

Example: Configuring the Virtual Router as a Route Reflector 108
Confederations 110
Example: Configuring a Confederation 111

Volume 7: Virtual Systems

Contents i
Preface iii
Conventions iv
CLI Conventions iv
WebUI Conventions v
Illustration Conventions vii
Naming Conventions and Character Types viii
NetScreen Documentation ix
Chapter 1 Virtual Systems 1
Creating a Vsys Object 3
Example: Vsys Objects and Admins 3
Virtual Routers 6
Zones 7
Interfaces 8
Traffic Sorting 10
Traffic Destined for the NetScreen Device 10
Through Traffic 11
Dedicated and Shared Interfaces 15
Dedicated Interfaces 15
Shared Interfaces 15
Importing and Exporting Physical Interfaces 18
Example: Importing a Physical Interface to a Virtual System 18
Example: Exporting a Physical Interface from a Virtual System 19
VLAN-Based Traffic Classification 21
VLANs 22
Defining Subinterfaces and VLAN Tags 23
Example: Defining Three Subinterfaces and VLAN Tags 25
Communicating between Virtual Systems 28
Example: InterVsys Communication 28
IP-Based Traffic Classification 33
Example: Configuring IP-Based Traffic Classification 35
Logging On as a Vsys Admin 38
Example: Logging On and Changing Your Password 38
Index ..... IX-I

Volume 8: High Availability

Contents i
BGP Communities 113
Index ..... IX-I
Preface v

Conventions vi
CLI Conventions vi
WebUI Conventions vii
Illustration Conventions ix
Naming Conventions and Character Types x
NetScreen Documentation xi
Chapter 1 NSRP 1
NSRP Overview 3
NSRP and NetScreen Operational Modes 8
Basic Active/Passive NSRP Configuration 8
Default Settings 9
Example: NSRP for an Active/Passive Configuration 10
NSRP Clusters 15
Cluster Name 17
Example: Creating an NSRP Cluster 18
Run-Time Objects 21
RTO Mirror States 22
VSD Groups 23
Preempt Option 23
VSD Group Member States 24
Heartbeat Messages 25
Example: Creating Two VSD Groups 26
VSIs and Static Routes 28
Example: Trust and Untrust Zone VSIs 29
Synchronization 33
Synchronizing Configurations 33
Synchronizing Files 34
Synchronizing RTOs 34
Example: Manually Resynchronizing RTOs 35
Example: Adding a Device to an Active NSRP Cluster 36
Synchronizing System Clocks 37
Dual HA Interfaces 38
Control Messages 39
Data Messages (Packet Forwarding) 40
Dynamic Routing Advisory 41
Dual HA Link Probes 42
Example: Sending Link Probes Manually 43
Example: Sending Link Probes Automatically 44
Setup Procedure 45
Cabling for a Full-Mesh Configuration 45
Active/Active NSRP Configuration 49
Example: NSRP for an Active/Active Configuration 49
Chapter 2 NSRP-Lite 57
Introduction to NSRP-Lite 59
Clusters and VSD Groups 60
Default Settings 61
Cluster 62
Cluster Name 63
Authentication and Encryption 64
VSD Group 65
VSD Group Member States 65
Heartbeat Messages 66
Preempt Option 67
Cabling and Configuring NSRP-Lite 68
Example: Configuring NSRP-Lite 69
Configuration and File Synchronization 76
Synchronizing Configurations 76

Synchronizing Files 77
Example: Adding a Device to an Active NSRP Cluster 77
Disabling Configuration and File Synchronization 78
Path Monitoring 79
Setting Thresholds 80
Weighting Tracked IP Addresses 80
IP Tracking for VPN Tunnel Failover 81
Example: IP Tracking through a VPN Tunnel 82
Chapter 3 Interface Redundancy 93
Redundant Interfaces 94
Example: Creating Redundant Interfaces for VSIs 96
Aggregate Interfaces 101
Example: Configuring an Aggregate Interface 102
Dual Untrust Interfaces 103
Interface Failover 104
Example: Manually Forcing Traffic from the Primary to the Backup Interface 104
Example: Manually Forcing Traffic from the Backup to the Primary Interface 104
Example: Automatically Switching Traffic between the Primary and Backup Interface 105
Determining Interface Failover 105
Interface Failover with IP Tracking 106
Example: Configuring Automatic Failover with IP Tracking 107
Interface Failover with VPN Tunnel Monitoring 111
Example: Configuring Automatic Failover with VPN Tunnel Monitoring 112
Serial Interface 118
Modem Settings 119
Example: Configuring Modem Settings 120
ISP Configuration 121
Example: Configuring ISP Information 122
Serial Interface Failover 123
Example: Configuring Dial Backup in the Trust-Untrust Mode 124
Example: Deleting a Default Route for the Serial Interface 127
Example: Adding a Default Route for the Serial Interface 127
Example: Specifying a Policy as Inactive for Serial Interface Failover 128
Chapter 4 Failover 129
Device Failover (NSRP) 130
VSD Group Failover (NSRP) 131
Configuring Object Monitoring for Device or VSD Group Failover 132
Configuring Monitored Objects 134
Physical Interface Objects 134
Example: Monitoring an Interface 134
Zone Objects 135
Example: Monitoring an Interface 135
Tracked IP Objects 136
Example: Track IP for Device Failover 139
Virtual System Failover 144
Example: VSIs for Inter-Virtual System Load Sharing 144
Index ..... IX-I

Preface

NetScreen devices are ASIC-based, ICSA-certified[1] Internet security appliances and security systems that integrate firewall, virtual private networking (VPN), and traffic-shaping features to provide flexible protection for security zones such as the internal local area network (LAN) or demilitarized zone (DMZ) when connecting to the Internet.

- **Firewall:** A firewall screens traffic crossing the boundary between a private LAN and the public network, such as the Internet.
- **VPN:** A VPN provides a secure communications channel between two or more remote network appliances.
- **Integrated Networking Functions:** Dynamic routing protocols learn reachability and advertise dynamically changing network topologies. In addition, traffic shaping functionality allows administrative monitoring and control of traffic passing across the NetScreen firewall to maintain a network's quality-of-service (QoS) level.
- **Redundancy:** High availability of interfaces, routing paths, NetScreen devices, and—on high-end NetScreen devices—power supplies and fans, to avoid a single point of failure in any of these areas.

Note: *For information on NetScreen compliance with Federal Information Processing Standards (FIPS) and for instructions on setting a FIPS-compliant NetScreen device in FIPS mode, see the platform-specific* NetScreen Cryptographic Module Security Policy *document on the NetScreen documentation CD-ROM*.

1. The Internet Computer Security Association (ICSA) is an organization focused on all types of network security for Internet-connected companies. Among its many functions, ICSA provides product certification for several kinds of security products such as virus protection, firewall, PKI, intrusion detection, IPSec, and cryptography. ICSA has certified all NetScreen products for firewall and IPSec.

Key Features in ScreenOS

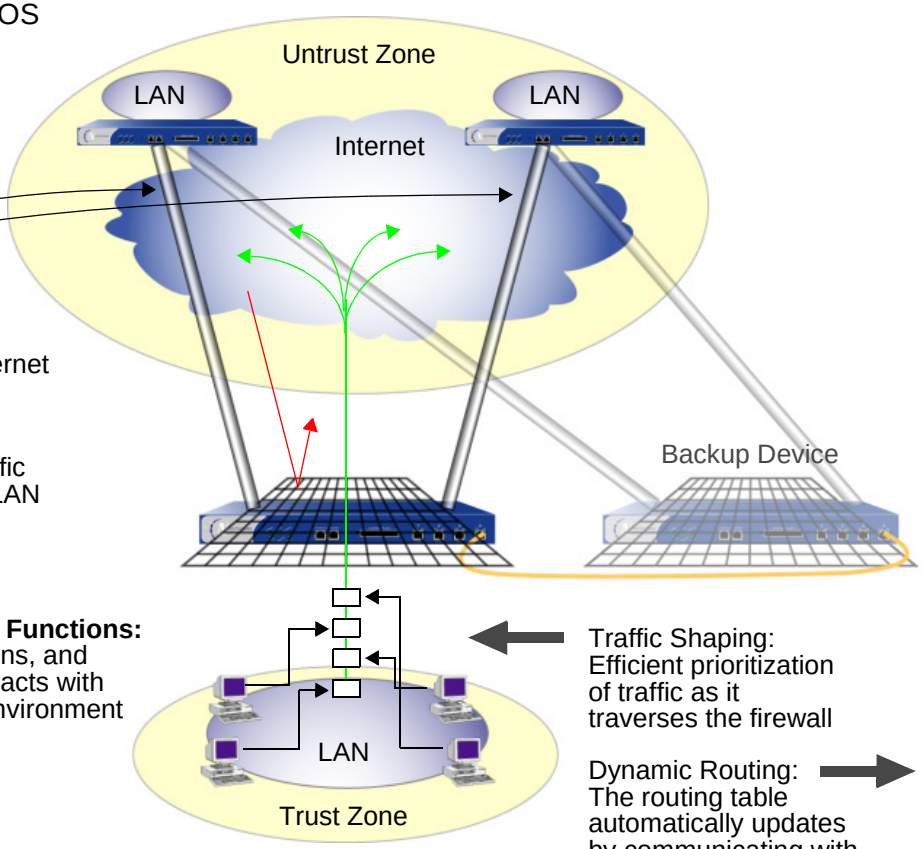

VPNs: Secure communication tunnels between sites for traffic passing through the Internet

Firewall: Screening traffic between the protected LAN and the Internet

Integrated Networking Functions: Performs routing functions, and communicates and interacts with routing devices in the environment

Redundancy: The backup device maintains identical configuration and sessions as those on the primary device to assume the place of the primary device if necessary. (Note: Interfaces, routing paths, power supplies, and fans can also be redundant.)

Traffic Shaping: Efficient prioritization of traffic as it traverses the firewall

Dynamic Routing: The routing table automatically updates by communicating with dynamic routing peers.

Dst	Use
0.0.0.0/0	1.1.1.250
1.1.1.0/24	eth3
1.2.1.0/24	eth2
10.1.0.0/16	trust-vr
10.2.2.0/24	tunnel.1
10.3.3.0/24	tunnel.2

NetScreen ScreenOS is the operating system that provides all the features needed to set up and manage any NetScreen security appliance or system. The *NetScreen Concepts & Examples ScreenOS Reference Guide* provides a useful reference guide for configuring and managing a NetScreen appliance through the ScreenOS.

CONCEPTS & EXAMPLES ORGANIZATION

The *NetScreen Concepts & Examples ScreenOS Reference Guide* is a multiple-volume set of documents. The following information outlines and summarizes the material in each volume:

Volume 1, "Overview"

- "Contents" contains a master table of contents for all volumes in the *NetScreen Concepts & Examples ScreenOS Reference Guide*.
- Appendix A, "Glossary" provides definitions for all the key terms used throughout all volumes in the *NetScreen Concepts & Examples ScreenOS Reference Guide*.
- "Index" is a master index encompassing all volumes in the *NetScreen Concepts & Examples ScreenOS Reference Guide*.

Volume 2, "Fundamentals"

- Chapter 1, "ScreenOS Architecture" presents the fundamental elements of USGA—the architecture in the NetScreen ScreenOS—and concludes with a four-part example illustrating an enterprise-based configuration incorporating most of those elements. In this and all subsequent chapters, each concept is accompanied by illustrative examples.
- Chapter 2, "Routing Tables and Static Routing" describes the ScreenOS routing table, the basic routing process on the NetScreen device, and how to configure static routes on NetScreen devices.
- Chapter 3, "Zones" explains security zones, tunnel zones, and function zones.
- Chapter 4, "Interfaces" describes the various physical, logical, and virtual interfaces on NetScreen devices, and includes information on various firewall attacks and the attack blocking options that NetScreen provides.
- Chapter 5, "Interface Modes" explains the concepts behind Transparent, Network Address Translation (NAT), and Route interface operational modes.
- Chapter 6, "Building Blocks for Policies" discusses the elements used for creating policies and virtual private networks (VPNs): addresses (including VIP addresses), users, and services. It also presents several example configurations support for the H.323 protocol.
- Chapter 7, "Policies" explores the components and functions of policies and offers guidance on their creation and application.

- Chapter 8, “Address Translation” explains the different methods for source address translation and destination address translation.
- Chapter 9, “User Authentication” details the various authentication methods and uses that NetScreen supports.
- Chapter 10, “Traffic Shaping” explains how you can manage bandwidth at the interface and Policy levels and prioritize services.
- Chapter 11, “System Parameters” presents the concepts behind Domain Name System (DNS) addressing; using Dynamic Host Configuration Protocol (DHCP) to assign or relay TCP/IP settings; downloading and uploading system configurations and software; and setting the system clock.

Volume 3, “Administration”

- Chapter 1, “Administration” explains the different means available for managing a NetScreen device both locally and remotely. This chapter also explains the privileges pertaining to each of the four levels of network administrators that can be defined. Finally, it explains how to secure local and remote administrative traffic.
- Chapter 2, “Monitoring NetScreen Devices” explains various monitoring methods and provides guidance in interpreting monitoring output.
- Appendix A, “SNMP MIB Files” lists and briefly describes the Management Information Base (MIB) files available for MIB compilers.

Volume 4, “Attack Detection and Defense Mechanisms”

- Chapter 1, “Protecting a Network” outlines the basic stages of an attack and the firewall options available to combat the attacker at each stage.
- Chapter 2, “Reconnaissance Deterrence” describes the options available for blocking IP address sweeps, port scans, and attempts to discover the type of operating system (OS) of a targeted system.
- Chapter 3, “Denial-of-Service Attack Defenses” explains firewall, network, and OS-specific DoS attacks and how NetScreen mitigates such attacks.
- Chapter 4, “Content Monitoring and Filtering” describes how to protect Hypertext Transfer Protocol (HTTP) and File Transfer Protocol (FTP) users from malicious uniform resource locators (URLs) and how to configure the NetScreen device to work with third party products to provide antivirus scanning and URL filtering.

- Chapter 5, “Deep Inspection” describes how to configure the NetScreen device to obtain IDP attack object updates, how to create user-defined attack objects and attack object groups, and how to apply IDP at the policy level.
- Chapter 1, “Suspicious Packet Attributes” explains a number of SCREEN options that block potentially dangerous packets.

Volume 5, “VPNs”

- Chapter 1, “IPSec” provides background information about IPSec, presents a flow sequence for Phase 1 in IKE negotiations in Aggressive and Main modes, and concludes with information regarding NAT-Traversal.
- Chapter 2, “Public Key Cryptography” provides information on how to obtain and load digital certificates and certificate revocation lists (CRLs).
- Chapter 3, “VPN Guidelines” offers some useful information to help in the selection of the available VPN options. It also presents a packet flow chart to demystify VPN packet processing.
- Chapter 4, “Site-to-Site VPNs” provides extensive examples VPN configurations connecting two private networks.
- Chapter 5, “Dialup VPNs” provides extensive examples of client-to-LAN communication using AutoKey IKE. It also details group IKE ID and shared IKE ID configurations.
- Chapter 6, “L2TP” explains the Layer 2 Tunneling Protocol (L2TP), its use alone and in conjunction with IPSec (L2TP-over-IPSec).
- Chapter 7, “Advanced VPN Features” contains information and examples for the more advanced VPN configurations, such as VPN monitoring, binding multiple tunnels to a single tunnel interface, and hub-and-spoke and back-to-back tunnel designs

Volume 6, “Dynamic Routing”

- Chapter 1, “Virtual Routers” explains how to configure virtual routers on NetScreen devices and how to redistribute routing table entries between protocols or between virtual routers.
- Chapter 2, “Open Shortest Path First (OSPF)” describes how to configure the OSPF dynamic routing protocol on NetScreen devices.
- Chapter 3, “Routing Information Protocol (RIP)” describes how to configure the RIP dynamic routing protocol on NetScreen devices.
- Chapter 4, “Border Gateway Protocol (BGP)” describes how to configure the BGP dynamic routing protocol on NetScreen devices.

Volume 7, “Virtual Systems”

- Chapter 1, “Virtual Systems” presents the concepts of virtual systems, dedicated and shared interfaces, and VLAN-based and IP-based traffic classification. It also explains how to set up virtual systems and create virtual system administrators.

Volume 8, “High Availability”

- Chapter 1, “NSRP” explains how to cable, configure, and manage NetScreen devices in a redundant group to provide high availability using the NetScreen Redundancy Protocol (NSRP).
- Chapter 2, “NSRP-Lite” explains how to configure NetScreen devices that support NSRP-Lite.
- Chapter 3, “Interface Redundancy” describes the various ways in which NetScreen devices provide interface redundancy.
- Chapter 4, “Failover” describes the configuration for the failover of a device, virtual security device (VSD) group, and virtual system. It also explains how to monitor certain objects to determine the failover of a device or VSD group.

CONVENTIONS

This document contains several types of conventions, which are introduced in the following sections:

- "CLI Conventions"
- "WebUI Conventions" on page xxviii
- "Illustration Conventions" on page xxx
- "Naming Conventions and Character Types" on page xxxi

CLI Conventions

The following conventions are used when presenting the syntax of a command line interface (CLI) command:

- Anything inside square brackets [ ] is optional.
- Anything inside braces { } is required.
- If there is more than one choice, each choice is separated by a pipe (|). For example,

  ```
  set interface { ethernet1 | ethernet2 | ethernet3 } manage
  ```

  means "set the management options for the ethernet1, ethernet2, or ethernet3 interface".
- Variables appear in *italic.* For example:

  ```
  set admin user name password
  ```

When a CLI command appears within the context of a sentence, it is in **bold** (except for variables, which are always in *italic*). For example: "Use the **get system** command to display the serial number of a NetScreen device."

Note: *When typing a keyword, you only have to type enough letters to identify the word uniquely. For example, typing* ***set adm u joe j12fmt54*** *is enough to enter the command* ***set admin user joe j12fmt54****. Although you can use this shortcut when entering commands, all the commands documented here are presented in their entirety.*

WebUI Conventions

Throughout this book, a chevron (>) is used to indicate navigation through the WebUI by clicking menu options and links. For example, the path to the address configuration dialog box is presented as **Objects > Addresses > List > New**. This navigational sequence is shown below.

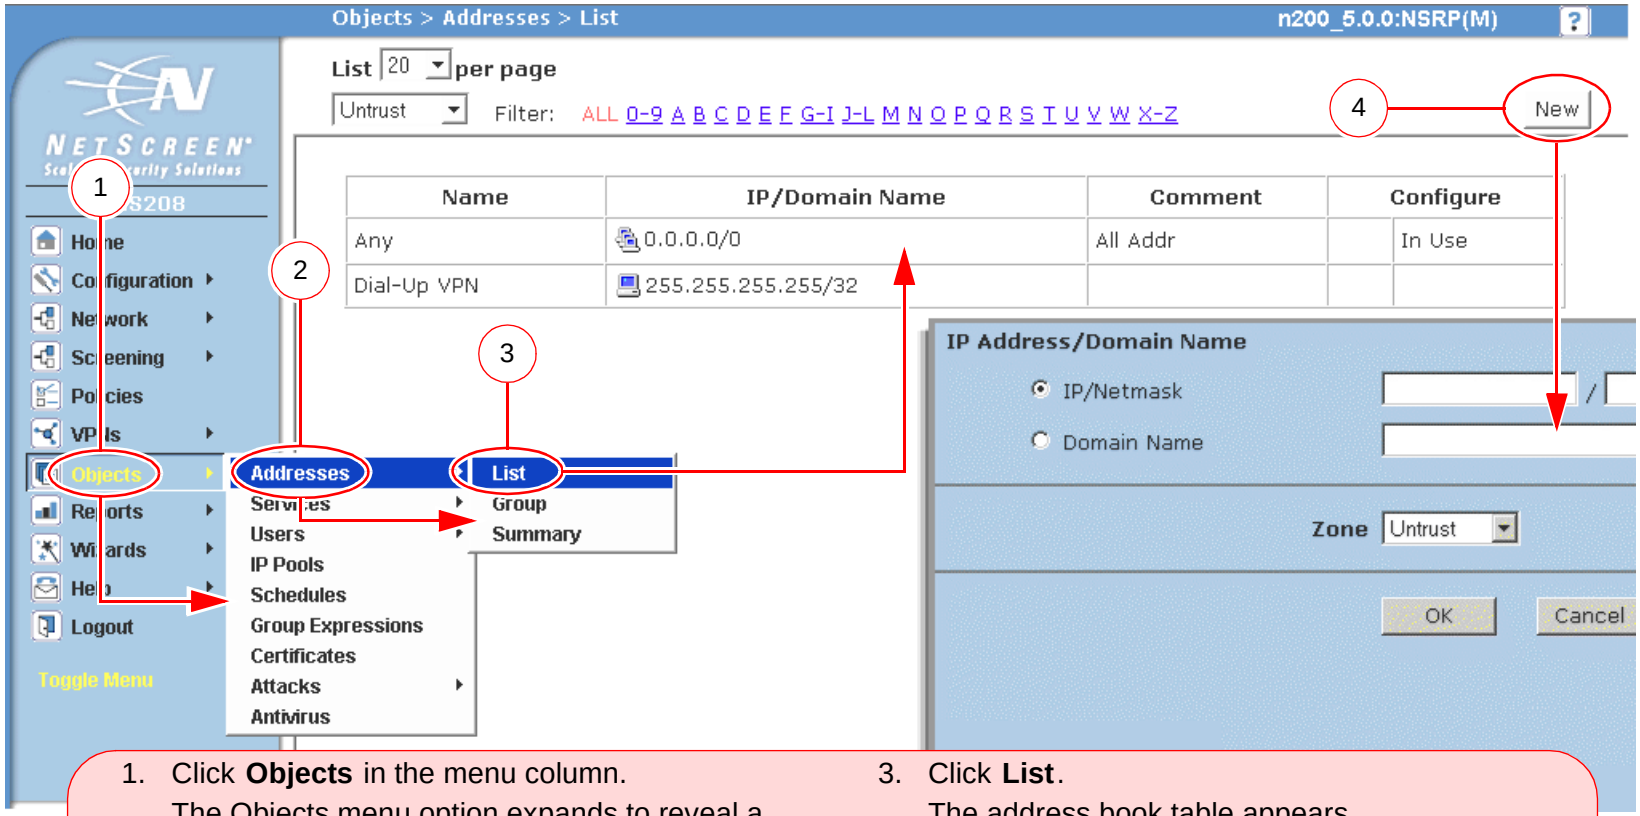

1. Click **Objects** in the menu column.
   The Objects menu option expands to reveal a subset of options for Objects.
2. (Applet menu) Hover the mouse over **Addresses**.
   (DHTML menu) Click **Addresses**.
   The Addresses option expands to reveal a subset of options for Addresses.
3. Click **List**.
   The address book table appears.
4. Click the **New** link.
   The new address configuration dialog box appears.

To perform a task with the WebUI, you must first navigate to the appropriate dialog box where you can then define objects and set parameters. The set of instructions for each task is divided into two parts: a navigational path and configuration details. For example, the following set of instructions includes the path to the address configuration dialog box and the settings for you to configure:

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: addr_1

IP Address/Domain Name:

IP/Netmask: (select), 10.2.2.5/32

Zone: Untrust

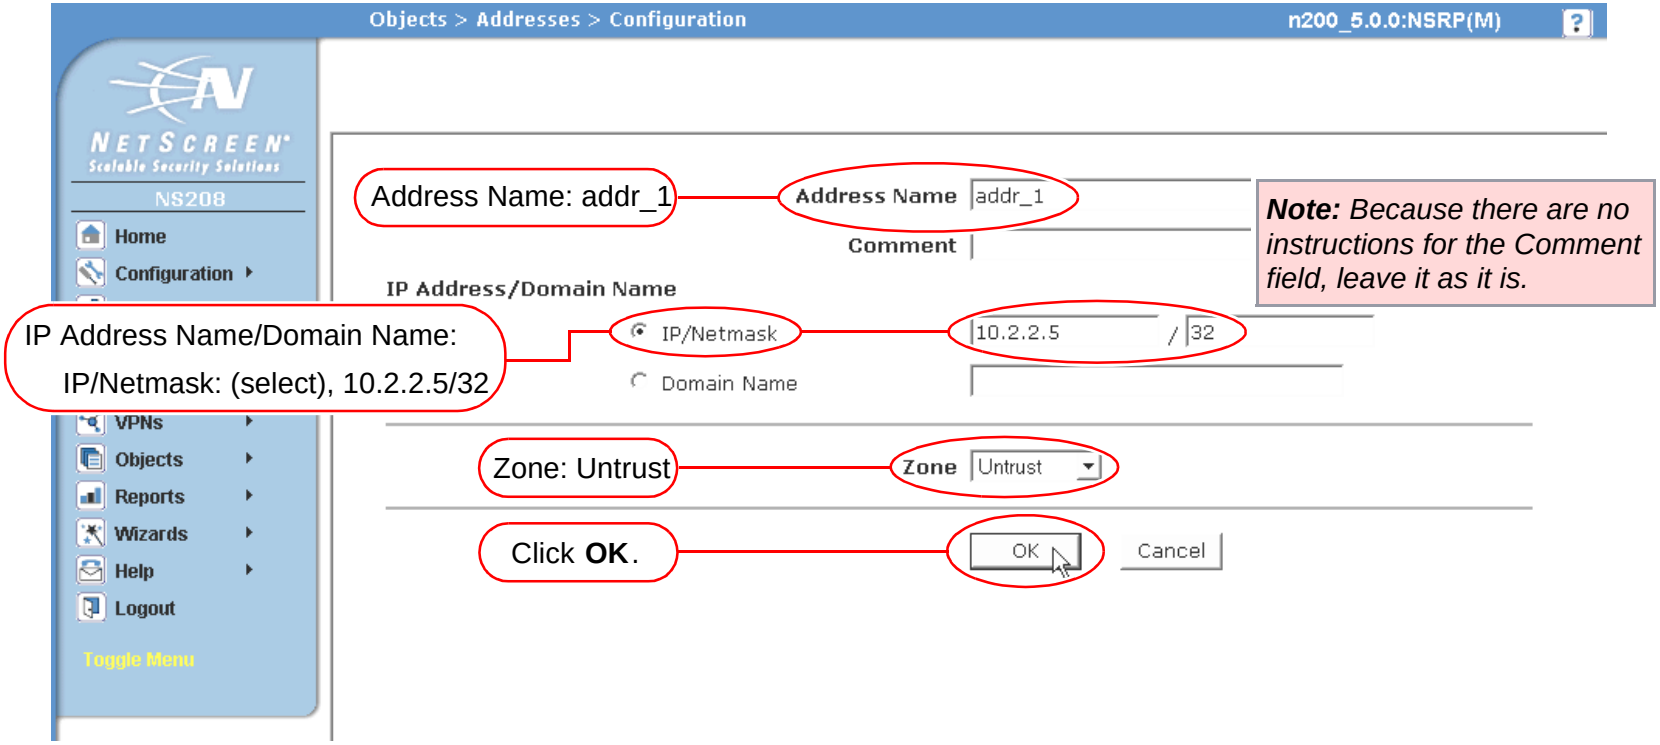

Illustration Conventions

The following graphics make up the basic set of images used in illustrations throughout this book:

Generic NetScreen Device

Virtual Routing Domain

Security Zone

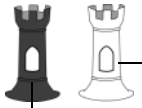

Security Zone Interfaces

White = Protected Zone Interface (example: Trust Zone)

Black = Outside Zone Interface (example: Untrust Zone)

Tunnel Interface

VPN Tunnel

Router Icon

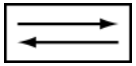

Switch Icon

Local Area Network (LAN) with a Single Subnet

(example: 10.1.1.0/24)

Internet

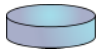

Dynamic IP (DIP) Pool

Desktop Computer

Laptop Computer

Generic Network Device

(examples: NAT server, Access Concentrator)

Server

Naming Conventions and Character Types

ScreenOS employs the following conventions regarding the names of objects—such as addresses, admin users, auth servers, IKE gateways, virtual systems, VPN tunnels, and zones—defined in ScreenOS configurations.

- If a name string includes one or more spaces, the entire string must be enclosed within double quotes ("); for example, **set address trust "local LAN" 10.1.1.0/24**.
- NetScreen trims any spaces leading or trailing text within a set of double quotes; for example, **" local LAN "** becomes **"local LAN"**.
- NetScreen treats multiple consecutive spaces as a single space.
- Name strings are case sensitive, although many CLI key words are case insensitive. For example, **"local LAN"** is different from **"local lan"**.

ScreenOS supports the following character types:

- Single-byte character sets (SBCS) and multiple-byte character sets (MBCS). Examples of SBCS are ASCII, European, and Hebrew. Examples of MBCS—also referred to as double-byte character sets (DBCS)—are Chinese, Korean, and Japanese.

  Note: *A console connection only supports SBCS. The WebUI supports both SBCS and MBCS, depending on the character sets that your Web browser supports.*

- ASCII characters from 32 (0x20 in hexidecimals) to 255 (0xff), except double quotes ("), which have special significance as an indicator of the beginning or end of a name string that includes spaces.

NETSCREEN DOCUMENTATION

To obtain technical documentation for any NetScreen product, visit www.netscreen.com/resources/manuals/.

To obtain the latest software version, visit www.netscreen.com/services/download_soft. Select a category of software product from the dropdown list, then follow the displayed instructions. (You must be a registered user to download NetScreen software.)

If you find any errors or omissions in the following content, please contact us at the e-mail address below:

techpubs@netscreen.com

A

Glossary

10BaseT: The most common form of ethernet is called 10BaseT, which denotes a peak transmission speed of 10 Mbps using copper twisted-pair cable. Ethernet is a standard for connecting computers into a local area network (LAN). The maximum cable distance is 100 meters (325 feet), the maximum devices per segment is 1, and the maximum devices per network are 1024. See also *100BaseT* and *Unshielded Twisted Pair (UTP).*

100BaseT: Another term for fast ethernet, an upgraded standard for connecting computers into a local area network (LAN). 100BaseT ethernet works just like regular ethernet except that it can transfer data at a peak rate of 100 Mbps. It is also more expensive and less common than its slower 10BaseT sibling. See also *10BaseT.*

Access List: To restrict the routing information that the router learns or advertises, you can filter based on routing updates to or from a particular neighbor. The filter consists of an access list that is applied to updates to or from a neighbor. The filtering of routing information can be applied on a per-neighbor or per-peer-group basis.

Access-Challenge: An additional condition required for a successful Telnet login by an authentication user via a RADIUS server.

Adjacencies: When two routers can exchange routing information with one another, they are considered to have constructed an adjacency. Point-to-point networks have only two routers so those routers automatically form an adjacency. But point-to-multipoint networks are a series of several point-to-point networks. When routers pair in this more complex networking scheme, they are considered to be adjacent to one another.

Advertisement: A method a router uses to announce itself to other devices on the network, transmitting basic information including IP address, network mask, and other data.

Aggregate State: A router is in an aggregate state when it is one of multiple virtual BGP routing instances bundled into one address.

Aggregation: The process of combining several different routes in such a way that only a single route advertises itself. This technique minimizes the size of the routing table for the router.

Aggregator: An object used to bundle multiple routes under one common route generalized according to the value of the network mask.

Aggressive Aging: A mechanism to accelerate the timeout process when the number of sessions in the session table surpasses a specified high-watermark threshold. When the number of sessions in the table dips below a specified low-watermark threshold, the timeout process returns to normal.

Area: The most fundamental ordering method in the OSPF routing protocol. An OSPF area divides the internetwork into smaller, more manageable constituent pieces. This technique reduces the amount of information that each router must store and maintain about all the other routers. When a router in the area needs information about another device in or out of the area, it contacts a special router that stores this information. This router is called the Area Border Router (ABR) and contains all essential device information. In addition, the ABR area border router filters all information coming into the area to avoid bogging down other routers in the area with information they may not need.

Area Range: A sequence of IP addresses defined by a lower limit and upper limit that indicates a series of addresses of devices that exist within an area.

Area Border Router: A router with at least one interface in Area 0 and at least one interface in another area.

AS: See *Autonomous System*.

AS Number: The identification number of the local autonomous system mapped to a BGP routing instance. The ID number can be any valid integer.

AS Path Access List: An access list used by a BGP routing instance to permit or deny packets sent by neighbor routing instances to the current virtual routing instance.

AS Path Attribute Class: The BGP provides four classes of path attributes: Well-Known Discretionary, Optional Transitive, and Optional Non-Transitive.

AS Path String: A string that acts as an identifier for an AS path. It is configured alongside an AS Path access list ID.

Atomic Aggregate: An object used by a BGP router to inform other BGP routers that the local system selected a generalized route.

Attack Objects: Stateful signatures and protocol anomalies that a NetScreen device with Deep Inspection functionality uses to detect attacks aimed at compromising one or more hosts on a network.

Authentication Header (AH): See *ESP/AH*.

Authentication: Authentication ensures that digital data transmissions are delivered to the intended receiver. Authentication also assures the receiver of the integrity of the message and its source (where or whom it came from). The simplest form of authentication requires a user name and password to gain access to a particular account. Authentication protocols can also be based on secret-key encryption, such as DES or 3DES, or on public-key systems using digital signatures.

Autonomous System (AS): An AS is a set of routers set off from the rest of the network and governed by a single technical administration. This router group uses an interior gateway protocol (IGP) or several IGPs and common metrics to route packets within the group. The group also uses an exterior gateway protocol (EGP) to route packets to other ASs. Each AS has a routing plan that indicates what destinations are reachable through it. This plan is called the Network Layer Reachability Information (NLRI) object. BGP routers generate and receive NLRI updates periodically.

Autonomous System Boundary Router: A router that connects an AS running one routing protocol to another AS running a different protocol.

Autonomous System Path: A list of all the autonomous systems that a router update has traveled through in the current transmission.

BGP: An inter-autonomous system routing protocol. BGP routers and autonomous systems exchange routing information for the Internet.

Bridge: A device that forwards traffic between network segments based on data link layer information. These segments share a common network layer address space.

Broadcast Network: A broadcast network is a network that supports many routers with the capability to communicate directly with one another. Ethernet is an example of a broadcast network.

Circuit-level Proxy: Proxy or Proxy Server is a technique used to cache information on a Web server and acts as an intermediary between a Web client and that Web server. It basically holds the most commonly and recently used content from the World Wide Web for users in order to provide quicker access and to increase server security. This is common for an ISP especially if they have a slow link to the Internet. On the Web, a proxy first attempts to find data locally, and if it is not there, fetches it from the remote server where the data resides permanently. Proxy servers are also constructs that allow direct Internet access from behind a firewall. They open a socket on the

server, and allow communication via that socket to the Internet. For example, if your computer is inside a protected network, and you want to browse the Web using Netscape, you can set up a proxy server on a firewall. You can configure the proxy server to allow HTTP requests to port 80 from your computer, and it then redirects all requests to the proper places.

Classless Routing: Support for interdomain routing, regardless of the size or class of the network. Network addresses are divided into three classes, but these are transparent in BGP, giving the network greater flexibility.

Cluster: A group of routers in a BGP AS where one is established as a route reflector and the others are clients to the reflector. The reflector is responsible for informing the clients of route and address information it learns from devices in another AS.

Note: *The term "cluster" has another meaning in regards to high availability. See "NetScreen Redundancy Protocol (NSRP)".*

Cluster List: A list of paths recorded as a packet travels through a BGP route reflector cluster.

Community: A community is a grouping of BGP destination. By updating the community, you automatically update its member destinations with new attributes.

Confederation: An object inside a BGP AS that is a subset of routing instances in the AS. By grouping devices into confederations inside a BGP AS, you reduce the complexity associated with the matrix of routing connections, known as a mesh, within the AS.

Connection States: When a packet sent from one router arrives at another router, a negotiation occurs between the source and destination routers. The negotiation goes through six states: Idle, Connect, Active, OpenSent, OpenConnect, and Establish.

Data Encryption Standard (DES): A 40- and 56-bit encryption algorithm that was developed by the National Institute of Standards and Technology (NIST). DES is a block encryption method originally developed by IBM. It has since been certified by the U.S. government for transmission of any data that is not classified top secret. DES uses an algorithm for private-key encryption. The key consists of 64 bits of data, which are transformed and combined with the first 64 bits of the message to be sent. To apply the encryption, the message is broken up into 64-bit blocks so that each can be combined with the key using a complex 16-step process. Although DES is fairly weak, with only one iteration, repeating it using slightly different keys can provide excellent security.

Data Encryption Standard-Cipher Block Chaining (DES-CBC): Until recently, the most significant use of triple-DES (3DES) was for the encryption of single DES keys, and there was really no need to consider how one might implement various block cipher modes when the block cipher in question is actually one derived from multiple encryption. However, as DES nears the end of its useful lifetime, more thought is being given to an increasingly widespread use of triple-DES. In particular, there are two obvious ways to implement the CBC mode for triple-DES. With single-DES in CBC mode, the ciphertext is exclusive-ored with the plaintext before encryption. With triple-DES however, we might use feedback around all three DES operations from the ciphertext to the plaintext, something which is called outer-CBC. Alternatively, we might run the feedback around each individual encryption component, thereby making, in effect, triple-(DES-CBC). This is referred to as inner-CBC, since there are internal feedbacks that are never seen by the crypto-analyst. Performance-wise, there can be some advantages to use the inner-CBC option, but research has established that outer-CBC is in fact more secure. Outer-CBC is the recommended way for using triple-DES in the CBC mode.

De-Militarized Zone (DMZ): From the military term for an area between two opponents where fighting is prevented. DMZ ethernets connect networks and computers controlled by different bodies. They may be external or internal. External DMZ ethernets link regional networks with routers.

Dead Interval: The amount of time that elapses before a routing instance determines another routing instance is not running.

Distance Vector: A routing strategy that relies on an algorithm that works by having routers sporadically broadcast entire copies of their own routing table to all directly connected neighbors. This update identifies the networks each router knows about, and the distance between each of those networks. The distance is measured in hop counts or the number of routing domains that a packet must traverse between its source device and the device it attempts to reach.

Dynamic Routing: A routing method which adjusts to changing network circumstances by analyzing incoming routing update messages. If the message indicates that a network change has occurred, the routing software recalculates routes and send s out new routing update messages. These messages populate the network, directing routers to rerun their algorithms and change their routing tables accordingly. There are two common forms of dynamic routing, including Distance Vector Routing and Link State Routing.

Encryption: Encryption is the process of changing data into a form that can be read only by the intended receiver. To decipher the message, the receiver of the encrypted data must have the proper decryption key. In traditional encryption schemes, the sender and the receiver use the same key to encrypt and decrypt data. Public-key encryption schemes use two keys: a public key, which anyone may use, and a corresponding private key, which is possessed only by the person who created it. With this method, anyone may send a message encrypted with the owner's public key, but only the owner has the private key necessary to decrypt it. DES (Data Encryption Standard) and 3DES (Triple DES) are two of the most popular public-key encryption schemes.

ESP/AH: The IP level security protocols, AH and ESP, were originally proposed by the Network Working Group focused on IP security mechanisms, IPSec. The term IPSec is used loosely here to refer to packets, keys, and routes that are associated with these protocols. The IP Authentication Header (AH) protocol provides authentication. The Encapsulating Security Protocol (ESP) provides both authentication and encryption.

Ethernet: A local area network technology invented at the Xerox Corporation, Palo Alto Research Center. Ethernet is a best-effort delivery system that uses CSMA/CD technology. Ethernet can be run over a variety of cable schemes, including thick coaxial, thin coaxial, twisted pair, and fiber optic cable. Ethernet is a standard for connecting computers into a local area network (LAN). The most common form of ethernet is called 10BaseT, which denotes a peak transmission speed of 10 Mbps using copper twisted-pair cable.

External Neighbors: Two BGP routers that are peers that reside in two different autonomous systems.

Extranet: The connecting of two or more intranets. An intranet is an internal Web site that allows users inside a company to communicate and exchange information. An extranet connects that virtual space with the intranet of another company, thus allowing these two (or more) companies to share resources and communicate over the Internet in their own virtual space. This technology greatly enhances business-to-business communications.

Filter List: A list of IP addresses permitted to send packets to the current routing domain.

Filtering, Dynamic: An IP service that can be used within VPN tunnels. Filters are one way some NetScreen devices control traffic from one network to another. When TCP/IP sends data packets to the firewall, the filtering function in the firewall looks at the header information in the packets and directs them accordingly. The filters operate on criteria such as IP source or destination address range, TCP ports, UDP, Internet Control Message Protocol (ICMP), or TCP responses. See also *Tunneling* and *Virtual Private Network (VPN).*

Firewall: A device that protects and controls the connection of one network to another, for traffic both entering and leaving. Firewalls are used by companies that want to protect any network-connected server from damage (intentional or otherwise) by those who log in to it. This could be a dedicated computer equipped with security measures or it could be a software-based protection.

Gateway: Also called a router, a gateway is a program or a special-purpose device that transfers IP datagrams from one network to another until the final destination is reached.

GBIC: A Gigabit Interface Connector (GBIC) is the kind of interface module card used on some NetScreen devices for connecting to a fiber optic network.

Hello Interval: The amount of time that elapses between instances of Hello Packets.

Hello Packet: A packet that advertises information, such as its presence and availability, to the network about the router that generated the packet.

Hold Time: In OSPF, the maximum amount of time between instances of initiating Shortest Path First (SPF) computations. In BGP, the maximum amount of time that elapses between message transmissions between a BGP speaker and its neighbor.

Hub: A hub is a hardware device used to link computers together (usually over an ethernet connection). It serves as a common wiring point so that information can flow through a central location to any other computer on the network. A hub repeats signals at the physical ethernet layer. A hub retains the behavior of a standard bus type network (such as Thinnet), but produces a star topology with the hub at the center of the star. This configuration enables centralized management.

Internet Control Message Protocol (ICMP): Occasionally a gateway or destination host uses ICMP to communicate with a source host, for example, to report an error in datagram processing. ICMP uses the basic support of IP as if it were a higher level protocol, however, ICMP is actually an integral part of IP, and must be implemented by every IP module. ICMP messages are sent in several situations: for example, when a datagram cannot reach its destination, when the gateway does not have the buffering capacity to forward a datagram, and when the gateway can direct the host to send traffic on a shorter route. The Internet Protocol is not designed to be absolutely reliable. The purpose of these control messages is to provide feedback about problems in the communication environment, not to make IP reliable.

Internet: Also known as “the Net.” Originally designed by the U.S. Defense Department so that a communication signal could withstand a nuclear war and serve military institutions worldwide. The Internet was first known as the ARPAnet. A system of linked computer networks, international in scope, that facilitates data communication services such as remote login, file transfer, electronic mail, and newsgroups. The Internet is a way of connecting existing computer networks that greatly extends the reach of each participating system.

Internet Key Exchange (IKE). The method for exchanging keys for encryption and authentication over an unsecured medium, such as the Internet.

Internet Protocol (IP): An Internet standard protocol that defines a basic unit of data called a datagram. A datagram is used in a connectionless, best-effort, delivery system. The Internet protocol defines how information gets passed between systems across the Internet.

IP Address: Each node on a TCP/IP network usually has an IP address. The IP address has a network number portion and a host number portion, as shown in the following table of IP address classes and formats:

Class	Number of Nodes	Address Format
A	> 32,768	nnn.hhh.hhh.hhh
B	256–32,768	nnn.nnn.hhh.hhh
C	< 256	nnn.nnn.nnn.hhh

This format is called decimal-dot format. The “n” represents a digit of a network number and “h” represents a digit of a host number; for example, 128.11.2.30. If you are sending data outside of your network, such as to the Internet, you need to obtain the network number from a central authority, currently the Network Information Center. See also *Netmask* and *Subnet Mask*.

IP Gateway: Also called a router, a gateway is a program or a special-purpose device that transfers IP datagrams from one network to another until the final destination is reached.

IP Security (IPSec): Security standard produced by the Internet Engineering Task Force (IETF). It is a protocol suite that provides everything you need for secure communications—authentication, integrity, and confidentiality—and makes key exchange practical even in larger networks. See also *DES-CBC,* and *ESP/AH.*

ISAKMP: The Internet Security Association and Key Management Protocol (ISAKMP) provides a framework for Internet key management and provides the specific protocol support for negotiation of security attributes. By itself, it does not establish session keys, however it can be used with various session key establishment protocols to provide a complete solution to Internet key management.

Intranet: A play on the word Internet, an intranet is a restricted-access network that works like the Web, but isn't on it. Usually owned and managed by a corporation, an intranet enables a company to share its resources with its employees without confidential information being made available to everyone with Internet access.

Keepalive: The amount of time in seconds that elapses between keepalive packets which ensures that the TCP connection between the local BGP router and a neighbor router is up. This value is equal to one-third of the hold time. The default is 60 seconds.

Key Management: The only reasonable way to protect the integrity and privacy of information is to rely upon the use of secret information in the form of private keys for signing and/or encryption. The management and handling of these pieces of secret information is generally referred to as "key management." This includes the activities of selection, exchange, storage, certification, expiration, revocation, changing, and transmission of keys. Most of the work in managing information security systems lies in the key management.

Link State: Link state routing protocols operate using an algorithm commonly called the Shortest Path First (SPF) algorithm. Instead of relying on rumored information from directly connected neighbors as in distance vector protocols, each router in a link state system maintains a complete topology of the network and computes SPF information based on the topology.

Link State Advertisement: The conveyance that enables OSPF routers to make device, network, and routing information available for the link state database. Each router retrieves information from the LSAs sent by other routers on the network to construct a picture of the entire internetwork from which an individual routing instance distills path information to use in its routing table.

Load balancing: Load balancing is the mapping (or re-mapping) of work to two or more processors, with the intent of improving the efficiency of a concurrent computation.

Local Area Network (LAN): Any network technology that interconnects resources within an office environment, usually at high speeds, such as ethernet. A local area network is a short-distance network used to link a group of computers together within a building. 10BaseT ethernet is the most commonly used form of LAN. A hardware device called a hub serves as the common wiring point, enabling data to be sent from one machine to another over the network. LANs are typically limited to distances of less than 1,640 feet (500 meters) and provide low-cost, high-bandwidth networking capabilities within a small geographical area.

Local Preference: To provide better information than the Multi-Exit Discriminator (MED) value provides for a packet's path selection, BGP provides an attribute known as the LOCAL_PREF or local preference value. You can configure the LOCAL_PREF attribute so that it has a higher value for prefixes received from a router that provides a desired path to be higher than prefixes heard on the router that provides a less desirable path. The higher the value, the more preferred the route. The LOCAL_PREF attribute is the metric most often used in practice to express preferences for one set of paths over another.

Mapped IP Address: A MIP is a direct one-to-one mapping of traffic destined for one IP address to another IP address.

MD5: Message Digest (version) 5, an algorithm that produces a 128-bit message digest (or hash) from a message of arbitrary length. The resulting hash is used, like a "fingerprint" of the input, to verify authenticity.

MED Comparison: The Multi Exit Discriminator (MED) attribute is used to determine an ideal link to reach a particular prefix in or behind the current Autonomous System (AS). The MED contains a metric expressing a degree of preference for entry into the AS. You can establish precedence for one link over others by configuring a MED value for one link that is lower than other links. The lower the MED value, the higher priority the link has. The way this occurs is that one AS sets the MED value and the other AS uses the value in deciding which path to choose.

Media Access Control (MAC) Address: An address that uniquely identifies the network interface card, such as an ethernet adapter. For ethernet, the MAC address is a 6 octet address assigned by IEEE. On a LAN or other network, the MAC address is a computer's unique hardware number. (On an ethernet LAN, it's the same as the ethernet address.) When you're connected to the Internet from your computer (or host as the Internet protocol thinks of it), a correspondence table relates your IP address to your computer's physical (MAC) address on the LAN. The MAC address is used by the Media Access Control sublayer of the Data-Link Control (DLC) layer of telecommunication protocols. There is a different MAC sublayer for each physical device type.

Multi Exit Discriminator: A BGP attribute that determines the relative preference of entry points into an Autonomous System.

Member AS: The name of the autonomous system being included in a BGP confederation.

Neighbor: To begin configuring a BGP network, you need to establish a connection between the current device and a counterpart, adjacent device known as a *neighbor* or *peer*. While this counterpart device may seem like unneeded information at first, it is actually central to the way BGP works. Unlike RIP or OSPF, you now have to configure two devices, both the current router and its neighbor, for BGP to work. While this requires more effort, it enables networking to occur on a larger scale as BGP eludes deploying the limited advertising techniques inherent to interior networking standards.

There are two types of BGP neighbors: **internal neighbors** which are in the same autonomous system and **external neighbors** which are in different autonomous systems. A reliable connection is required between neighbors and is achieved by creating a TCP connection between the two. The handshake that occurs between the two prospect neighbors evolves through a series of phases or *states* before a true connection can be made. See *Connection States*.

Netmask: A netmask indicates which part of an IP address indicates network identification and which part indicates the host identification. For example, the IP address and netmask 10.20.30.1 255.255.255.0 (or 10.20.30.1/24) refers to all the hosts in the 10.20.30.0 subnet. The IP address and netmask 10.20.30.1 255.255.255.255 (or 10.20.30.1/32) refers to a single host. See also *IP Address* and *Subnet Mask.*

NetScreen Redundancy Protocol (NSRP): A proprietary protocol that provides configuration and run time object (RTO) redundancy and a device failover mechanism for NetScreen units in a high availability (HA) cluster.

Network Address Translation (NAT): A standard for translating secure IP addresses to temporary, external, registered IP address from the address pool. This allows Trusted networks with privately assigned IP addresses to have access to the Internet. This also means that you don’t have to get a registered IP address for every machine in your network.

Network Layer Reachability Information: Each AS has a routing plan that indicates what destinations are reachable through it. This routing plan is called the Network Layer Reachability Information (NLRI) object. BGP routers generate and receive NLRI updates periodically. Each NLRI update contains information on the list of ASs that reachability information capsules traverse. Common values described by an NLRI update include: a network number, a list of ASs that the information passed through, and a list of other path attributes.

Peer: See *Neighbor*

Policies: Policies provide the initial protection mechanism for the firewall, allowing you to determine which traffic passes across it based on IP session details. You can use policies to protect the resources in a security zone from attacks from another zone (interzone policies) or from attacks from within a zone (intrazone policies). You can also use policies to monitor traffic attempting to cross your firewall.

Prefix: An IP address that represents a route.

Redistribution: The process of importing a route into the current routing domain from another part of the network that uses another routing protocol. When this occurs, the current domain has to translate all the information, particularly known routes, from the other protocol. For example, if you are on an OSPF network and it connects to a BGP network, the OSPF domain has to import all the routes from the BGP network to inform all of its devices about how to reach all the devices on the BGP network. The receipt of all the route information is known as route redistribution.

Redistribution List: A list of routes the current routing domain imported from another routing domain using a different protocol.

RJ-45: Resembling a standard phone connector, an RJ-45 connector is twice as wide (with eight wires) and is used for hooking up computers to local area networks (LANs) or phones with multiple lines.

Route Flap Damping: BGP provides a technique to block the advertisement of the route somewhere close to the source until the route becomes stable. This method is called *flap damping*. Route flap damping allows routing instability to be contained at an AS border router adjacent to the region where instability is occurring. The impact of limiting the unnecessary propagation is to maintain reasonable route change convergence time as a routing topology grows.

Route Map: Route maps are used with BGP to control and modify routing information and to define the conditions by which routes are redistributed between routing domains. A route map contains a list of route map entries, each containing a sequence number and a match and a set value. The route map entries are evaluated in the order of an incrementing sequence number. Once an entry returns a matched condition, no further route maps are evaluated. Once a match has been found, the route map carries out a permit or deny operation for the entry. If the route map entry is not a match, then the next entry is evaluated for matching criteria.

Route Redistribution: The exporting of route rules from one virtual router to another.

Route Reflector: A router whose BGP configuration enables readvertising of routes between Interior BGP (IBGP) neighbors or neighbors within the same BGP AS. A route reflector client is a device that uses a route reflector to readvertise its routes to the entire AS. It also relies on that route reflector to learn about routes from the rest of the network.

Router: A hardware or virtual (in a NetScreen environment) device that distributes data to all other routers and receiving points in or outside of the local routing domain. Routers also act as filters, allowing only authorized devices to transmit data into the local network so that private information can remain secure. In addition to supporting these connections, routers also handle errors, keep network usage statistics, and handle security issues.

Routing Table: A list in a virtual router's memory that contains a real-time view of all the connected and remote networks to which a router is currently routing packets.

Run Time Object (RTO): A code object created dynamically in memory during normal operation. Some examples of RTOs are session table entries, ARP cache entries, certificates, DHCP leases, and IPSec Phase 2 security associations (SAs).

Security Association: An SA is a unidirectional agreement between the VPN participants regarding the methods and parameters to use in securing a communication channel. For bidirectional communication, there must be at least two SAs, one for each direction. The VPN participants negotiate and agree to Phase 1 and Phase 2 SAs during an AutoKey IKE negotiation. See also *Security Parameters Index.*

Security Parameters Index: (SPI) is a hexadecimal value which uniquely identifies each tunnel. It also tells the NetScreen device which key to use to decrypt packets.

Security Zone: A security zone is a collection of one or more network segments requiring the regulation of inbound and outbound traffic via policies.

SHA-1: Secure Hash Algorithm-1, an algorithm that produces a 160-bit hash from a message of arbitrary length. (It is generally regarded as more secure than MD5 because of the larger hashes it produces.)

Static Routing: User-defined routes that cause packets moving between a source and a destination to take a specified path. Static routing algorithms are table mappings established by the network administrator prior to the beginning of routing. These mappings do not change unless the network administrator alters them. Algorithms that use static routes are simple to design and work well in environments where network traffic is relatively predictable and where network design is relatively simple.

The software remembers static routes until you remove them. However, you can override static routes with dynamic routing information through judicious assignment of administrative distance values. To do this, you must ensure that the administrative distance of the static route is higher than that of the dynamic protocol.

Subinterface: A subinterface is a logical division of a physical interface that borrows the bandwidth it needs from the physical interface from which it stems. A subinterface is an abstraction that functions identically to an interface for a physically present port and is distinguished by 802.1Q VLAN tagging.

Subnet Mask: In larger networks, the subnet mask lets you define subnetworks. For example, if you have a class B network, a subnet mask of 255.255.255.0 specifies that the first two portions of the decimal dot format are the network ID, while the third portion is a subnet ID. The fourth portion is the host ID. If you do not want to have a subnet on a class B network, you would use a subnet mask of 255.255.0.0. A network can be subnetted into one or more physical networks which form a subset of the main network. The subnet mask is the part of the IP address which is used to represent a subnetwork within a network. Using subnet masks allows you to use network address space which is normally unavailable and ensures that network traffic does not get sent to the whole network unless intended. See also *IP address* and *Netmask.*

Three-Way Handshake: A TCP connection is established with a triple exchange of packets known as a three-way handshake. The procedure transpires as follows:

1. The initiator sends a SYN (synchronize/start) packet.
2. The recipient replies with a SYN/ACK (synchronize/acknowledge) packet.
3. The initiator responds with an ACK (acknowledge) packet.
4. At this point, the two endpoints of the connection have been established and data transmission can commence.

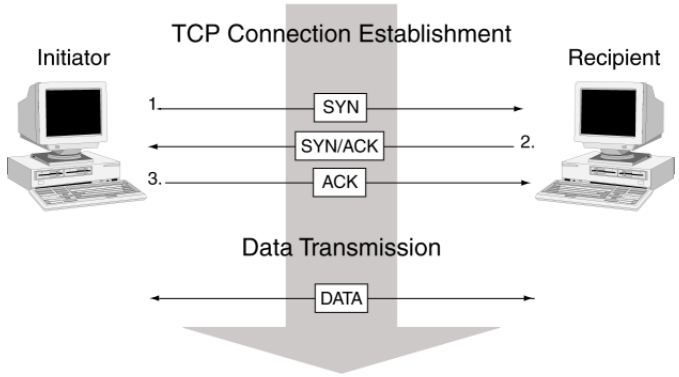

Transmission Control Protocol/Internet Protocol (TCP/IP): TCP/IP is a set of communication protocols that support peer-to-peer connectivity functions for both local and wide area networks. (A communication protocol is a set of rules that allow computers with different operating systems to communicate with each other.) TCP/IP controls how data is transferred between computers on the Internet.

Trunk Port: A trunk port allows a switch to bundle traffic from several VLANs through a single physical port, sorting the various packets by the VLAN identifier (VID) in their frame headers.

Trust: One of two NetScreen zones that enables packets to be secured from being seen by devices external to your current NetScreen domain.

Tunneling: A method of data encapsulation.With VPN tunneling, a mobile professional dials into a local Internet Service Provider's Point of Presence (POP) instead of dialing directly into their corporate network. This means that no matter where mobile professionals are located, they can dial a local Internet Service Provider that supports VPN

tunneling technology and gain access to their corporate network, incurring only the cost of a local telephone call. When remote users dial into their corporate network using an Internet Service Provider that supports VPN tunneling, the remote user as well as the organization knows that it is a secure connection. All remote dial-in users are authenticated by an authenticating server at the Internet Service Provider's site and then again by another authenticating server on the corporate network. This means that only authorized remote users can access their corporate network, and can access only the hosts that they are authorized to use.

Tunnel Interface: A tunnel interfaces is the opening, or doorway, through which traffic to or from a VPN tunnel passes. A tunnel interface can be numbered (that is, assigned an IP address) or unnumbered. A numbered tunnel interface can be in either a tunnel zone or security zone. An unnumbered tunnel interface can only be in a security zone that contains at least one security zone interface. The unnumbered tunnel interface borrows the IP address from the security zone interface.

Tunnel Zone: A tunnel zone is a logical segment that hosts one or more tunnel interfaces. A tunnel zone is associated with a security zone that acts as its carrier.

User Datagram Protocol (UDP): A protocol in the TCP/IP protocol suite, the User Datagram Protocol or UDP allows an application program to send datagrams to other application programs on a remote machine. Basically UDP is a protocol that provides an unreliable and connectionless datagram service where delivery and duplicate detection are not guaranteed. It does not use acknowledgments, or control the order of arrival.

Uniform Resource Locator (URL): A standard way developed to specify the location of a resource available electronically. Also referred to as a location or address, URLs specify the location of files on servers. A general URL has the syntax protocol://address. For example, http://www.netscreen.com/support/manuals.html specifies that the protocol is HTTP and the address is www.netscreen.com/support/manuals.html.

Unshielded Twisted Pair (UTP): Also known as 10BaseT. This is the standard cabling used for telephone lines. It is also used for ethernet connections. See also *10BaseT.*

Untrust: One of two NetScreen zones that enables packets to be seen by devices external to your current NetScreen domain.

Virtual Adapter: The TCP/IP settings (IP address, DNS server addresses, and WINS server addresses) that a NetScreen device assigns to a remote XAuth user for use in a VPN connection.

Virtual IP Address: A VIP address maps traffic received at one IP address to another address based on the destination port number in the packet header.

Virtual Local Area Network (VLAN): A logical rather than physical grouping of devices that constitute a single broadcast domain. VLAN members are not identified by their location on a physical subnetwork but through the use of tags in the frame headers of their transmitted data. VLANs are described in the IEEE 802.1Q standard.

Virtual Private Network (VPN): A VPN is an easy, cost-effective and secure way for corporations to provide telecommuters and mobile professionals local dial-up access to their corporate network or to another Internet Service Provider (ISP). Secure private connections over the Internet are more cost-effective than dedicated private lines. VPNs are possible because of technologies and standards such as tunneling, screening, encryption, and IPSec.

Virtual Router: A virtual router is the component of ScreenOS that performs routing functions. By default, a NetScreen device supports two virtual routers: Untrust-VR and Trust-VR.

Virtual Security Device (VSD): A single logical device composed by a set of physical NetScreen devices.

Virtual Security Interface (VSI): A logical entity at layer 3 that is linked to multiple layer 2 physical interfaces in a VSD group. The VSI binds to the physical interface of the device acting as master of the VSD group. The VSI shifts to the physical interface of another device in the VSD group if there is a failover and it becomes the new master.

Virtual System: A virtual system (vsys) is a subdivision of the main system that appears to the user to be a stand-alone entity. Virtual systems reside separately from each other in the same NetScreen device. Each one can be managed by its own virtual system administrator.

Windows Internet Naming Service (WINS): WINS is a service for mapping IP addresses to NetBIOS computer names on Windows NT server-based networks. A WINS server maps a NetBIOS name used in a Windows network environment to an IP address used on an IP-based network.

Zone: A zone can be a segment of network space to which security measures are applied (a security zone), a logical segment to which a VPN tunnel interface is bound (a tunnel zone), or either a physical or logical entity that performs a specific function (a function zone).

Index

Symbols

100BaseT, defined **1**-A-I
10BaseT, defined **1**-A-I
3DES **5**-8

A

access list **1**-A-I
access list for routes **6**-24
access policies
    See policies
access-challenge **1**-A-I
ActiveX controls, blocking **4**-167
address book
    adding addresses **2**-127
    editing group entries **2**-132
    entries **2**-127
    groups **2**-129
    modifying addresses **2**-128
    removing addresses **2**-133
    See also addresses
address groups **2**-129, **2**-206
    creating **2**-131
    editing **2**-132
    options **2**-130
    removing entries **2**-133
address negation **2**-237
address sweep **4**-8
address translation
    See NAT, NAT-dst, and NAT-src
addresses
    address book entries **2**-127
    defined **2**-206
    in policies **2**-206
    private **2**-78
    public **2**-77
adjacencies **1**-A-I
admin users **2**-465– **2**-466
    auth process **2**-466
    privileges from RADIUS **2**-465
    server support **2**-372
    timeout **2**-378
administration
    CLI (Command Line Interface) **3**-9
    restricting **3**-49, **3**-50
    vsys admin **7**-38
    WebUI **3**-3
administrative traffic **3**-30
advertisement **1**-A-I
AES (Advanced Encryption Standard) **5**-8
aggregate interfaces **2**-67, **8**-101
aggregate state **1**-A-I
aggregation **1**-A-I
aggregator **1**-A-I
aggressive aging **4**-40– **4**-42
    defined **1**-A-II
Aggressive Mode **5**-12
AH **5**-3, **5**-7
alarms
    E-mail alert **3**-82
    reporting to NSM **3**-26
    thresholds **2**-213, **3**-82
    traffic **3**-82– **3**-86
ALG **2**-159, **4**-73
    for custom services **2**-207
anti-replay checking **5**-46, **5**-54
antivirus objects
    See AV objects
antivirus scanning
    policies **2**-214
    See AV scanning
application layer gateway
    See ALG
application, in policies **2**-207
area **1**-A-II
area border router **1**-A-II
area range **1**-A-II
ARP **2**-95, **8**-56, **8**-136
    broadcasts **8**-18
    ingress IP address **2**-98
    path monitoring **8**-79
AS number **1**-A-II
AS path access list **1**-A-II
AS path attribute class **1**-A-II
AS path string **1**-A-II
asset recovery log **3**-81
atomic aggregate **1**-A-II
attack actions **4**-142– **4**-151
    close **4**-142
    close client **4**-142
    close server **4**-142
    drop **4**-142
    drop packet **4**-142
    ignore **4**-143
    none **4**-143
attack object database **4**-128– **4**-135
    auto notification and manual update **4**-128, **4**-132
    automatic update **4**-128, **4**-130
    changing the default URL **4**-134
    immediate update **4**-128, **4**-129
    manual update **4**-129, **4**-134
attack object groups **4**-140
    changing severity **4**-140
    severity levels **4**-140
attack objects **4**-125
    defined **1**-A-II
    protocol anomalies **4**-139
    stateful signatures **4**-138
    TCP stream signatures **4**-164
attack protection
    policy level **4**-5
    security zone level **4**-5
attacks
    common objectives **4**-1
    detection and defense options **4**-3– **4**-5
    ICMP flood **4**-59
    ICMP fragments **4**-2
    IP packet fragments **4**-10
    Land attack **4**-63
    large ICMP packets **4**-4
    Ping of Death **4**-65
    Replay **5**-14
    stages of **4**-2

SYN flood **4**-45– **4**-51
SYN fragments **4**-12– **4**-13
Teardrop **4**-67
UDP flood **4**-61
unknown MAC addresses **4**-51
unknown protocols **4**-8
WinNuke **4**-69
auth servers **2**-372
address **2**-377
authentication process **2**-376
backup servers **2**-377
default **2**-395
defining **2**-388– **2**-396
external **2**-376
feature support **2**-372
ID number **2**-377
in IKE gateways **2**-396
in policies **2**-396
LDAP **2**-386– **2**-387
LDAP, defining **2**-393
maximum number **2**-373
multiple user types **2**-373
object name **2**-377
object properties **2**-377
RADIUS **2**-379– **2**-381
RADIUS, defining **2**-388
RADIUS, user type support **2**-380
SecurID **2**-384– **2**-385
SecurID, defining **2**-391
timeout **2**-377
types **2**-377
user type support **2**-372
XAuth queries **2**-437
auth users **2**-398– **2**-430
groups **2**-398, **2**-402
in policies **2**-398
point of authentication **2**-397
pre-policy auth **2**-212, **2**-400
run-time (external user group) **2**-412
run-time (external user) **2**-409
run-time (local user group) **2**-406
run-time (local user) **2**-403
run-time auth process **2**-211, **2**-399
run-time authentication **2**-211, **2**-399
server support **2**-372
timeout **2**-377
WebAuth **2**-212, **2**-400
WebAuth (external user group) **2**-423
WebAuth (local user group) **2**-420
WebAuth + SSL (external user group) **2**-427
authentication
algorithms **5**-7, **5**-44, **5**-49, **5**-53, **5**-57
Allow Any **2**-212
IPSec **1**-A-III
NSRP **8**-7, **8**-18
NSRP-Lite **8**-64
policies **2**-210
users **2**-210, **2**-371– **2**-476
WebAuth **2**-400
Authentication Header
See AH
authentication, users **2**-371– **2**-476
accounts **2**-371
admin **2**-465
auth servers **2**-372
auth users **2**-398
IKE users **2**-372, **2**-431
L2TP users **2**-460
local database **2**-374– **2**-375
Manual Key users **2**-372
multiple-type **2**-467
point of authentication **2**-397
profiles **2**-371
types and applications **2**-397– **2**-467
user types **2**-372
WebAuth **2**-372
with different logins **2**-467
XAuth users **2**-436
AutoKey IKE VPN **3**-51, **3**-98, **5**-9
management **5**-9
autonomous system boundary router **1**-A-III
autonomous system path **1**-A-III
autonomous systems
defined **1**-A-III
AV objects **4**-93– **4**-101
port number **4**-94
states **4**-93
timeout **4**-94
AV scanning **4**-76– **4**-112
application **4**-102
AV objects **4**-93– **4**-101
decompression **4**-85
external AV scanner **4**-90– **4**-112
external, CSP resources **4**-95
external, HTTP **4**-92
external, SMTP **4**-91
fail-mode **4**-95
fail-mode threshold **4**-96
HTTP keep-alive **4**-96
HTTP trickling **4**-97
HTTP webmail **4**-80
internal AV scanner **4**-77– **4**-89
internal, HTTP **4**-79
internal, POP3 **4**-78
internal, SMTP **4**-77
internal, subscription **4**-81
InterScan VirusWall **4**-90
max TCP connections **4**-94
multiple AV objects **4**-106

B

back store **3**-122
bandwidth **2**-215
default priority **2**-485
guaranteed **2**-215, **2**-478, **2**-486
managing **2**-478
maximum **2**-215, **2**-486
maximum specification **2**-478
priority levels **2**-485
priority queues **2**-485
unlimited maximum **2**-478
banners, customizing **2**-476
BGP **1**-A-III
AS-path access list **6**-106
authenticating neighbors **6**-102
communities **6**-113
confederations **6**-110
configuration steps **6**-91
configuring peer group **6**-95
configuring peers **6**-95
creating instance in VR **6**-92
enabling in VR **6**-92
enabling on interface **6**-94
external BGP **6**-90
internal BGP **6**-90
message types **6**-89
parameters **6**-104

path attributes **6**-89
protocol overview **6**-88
redistributing routes **6**-105
regular expressions **6**-106
rejecting default routes **6**-103
route reflection **6**-107
security configuration **6**-102
verifying configuration **6**-100
bit stream **3**-121
bridges **1**-A-III
broadcast networks **1**-A-III
browser requirements **3**-3

C

CA certificates **5**-18, **5**-22
cables, serial **3**-21
certificates **5**-10
CA **5**-18, **5**-22
loading **5**-26
local **5**-22
requesting **5**-23
revocation **5**-21, **5**-36
via e-mail **5**-22
Challenge Handshake Authentication Protocol
See CHAP
CHAP **2**-453, **5**-273, **5**-276
character types, ScreenOS supported **1**-xxxi, **2**-xiv, **3**-x, **4**-x, **5**-x, **6**-x, **7**-viii, **8**-x
classless routing **1**-A-IV
CLI **3**-9, **3**-30, **3**-31
conventions **1**-xxvii, **2**-x, **3**-vi, **4**-vi, **5**-vi, **6**-vi, **7**-iv, **8**-vi
set arp always-on-dest **8**-56
set vip multi -port **2**-358
clock, system **2**-541– **2**-545
See also system clock
cluster list **1**-A-IV
cluster name, NSRP **8**-17, **8**-63
clusters **8**-16– **8**-20, **8**-49, **8**-60– **8**-63
defined **1**-A-IV
command line interface
See CLI
common name **2**-387
community **1**-A-IV
CompactFlash **3**-66
confederation **1**-A-IV
configuration
adding comments **2**-535
LKG **2**-531
loading **2**-533
locking **2**-534
rollback **2**-531– **2**-532, **2**-533
configuration settings
browser requirements **3**-3
downloading **2**-528
uploading **2**-528
connection states **1**-A-IV
connectors
GBIC, definition **1**-A-VII
RJ-45, definition **1**-A-XII
console **3**-66
container **5**-242
content filtering **4**-71– **4**-121
Content Scanning Protocol
See CSP
control messages **8**-38
HA messages **8**-40
HA physical link heartbeats **8**-39
RTO heartbeats **8**-40
VSD heartbeats **8**-40
conventions
CLI **1**-xxvii, **2**-x, **3**-vi, **4**-vi, **5**-vi, **6**-vi, **7**-iv, **8**-vi
illustration **1**-xxx, **2**-xiii, **3**-ix, **4**-ix, **5**-ix, **6**-ix, **7**-vii, **8**-ix
names **1**-xxxi, **2**-xiv, **3**-x, **4**-x, **5**-x, **6**-x, **7**-viii, **8**-x
WebUI **1**-xxviii, **2**-xi, **3**-vii, **4**-vii, **5**-vii, **6**-vii, **7**-v, **8**-vii
counting **2**-213
creating
address groups **2**-131
keys **3**-7
MIP addresses **2**-333
service groups **2**-168
zones **2**-51
CRL (Certificate Revocation List) **5**-20, **5**-36
loading **5**-20
cryptographic options **5**-40– **5**-57
anti-replay checking **5**-46, **5**-54
authentication algorithms **5**-44, **5**-49, **5**-53, **5**-57
authentication types **5**-42, **5**-51
certificate bit lengths **5**-43, **5**-51
dialup **5**-50– **5**-57
dialup VPN recommendations **5**-57
Diffie-Hellman groups **5**-43, **5**-46, **5**-52, **5**-55
encryption algorithms **5**-44, **5**-48, **5**-52, **5**-57
ESP **5**-48, **5**-56
IKE ID **5**-44– **5**-46, **5**-53– **5**-54
IPSec protocols **5**-47, **5**-56
key methods **5**-42
PFS **5**-46, **5**-55
Phase 1 modes **5**-42, **5**-51
site-to-site **5**-41– **5**-49
site-to-site VPN recommendations **5**-49
Transport mode **5**-56
Tunnel mode **5**-56
CSP **4**-90

D

Data Encryption Standard
See DES
data messages **8**-40
DDoS **4**-35
dead interval **1**-A-V
decompression, AV scanning **4**-85
Deep Inspection **4**-140– **4**-163
attack actions **4**-142– **4**-151
attack object database **4**-128– **4**-135
attack object groups **4**-140
attack objects **4**-125
attack objects, defined **1**-A-II
changing severity **4**-140
context **4**-156
custom attack objects **4**-156
custom services **4**-152– **4**-155
custom signatures **4**-157– **4**-163
protocol anomalies **4**-139
regular expressions **4**-157– **4**-159
stateful signatures **4**-138

defining
subinterfaces **7**-25
zones **2**-51
Denial-of-Service
See DoS
DES **5**-8
defined **1**-A-IV
DES-CBC, defined **1**-A-V
device failover **8**-130
DHCP **2**-115, **2**-121, **2**-521
client **2**-500
HA **2**-508
relay agent **2**-500
server **2**-500
dictionary file **2**-465
Diffie-Hellman exchange **5**-13
Diffie-Hellman groups **5**-13, **5**-43, **5**-46, **5**-52, **5**-55
DiffServ **2**-215
See DS Codepoint Marking
digital signature **5**-16
DIP **2**-119, **2**-171– **2**-174, **3**-124
fix-port **2**-173
groups **2**-189– **2**-192
modifying a DIP pool **2**-174
PAT **2**-172
pools **2**-210
DIP pools
address considerations **2**-259
extended interfaces **5**-168
NAT for VPNs **5**-168
NAT-src **2**-246
size **2**-259
distance vector **1**-A-V
distinguished name **2**-387
DMZ, definition **1**-A-V
DN (distinguished name) **5**-237
DNS **2**-495
L2TP settings **5**-276
lookup **2**-496
server **2**-523
status table **2**-497
Domain name system
See DNS
DoS **4**-35– **4**-70
firewall **4**-36– **4**-44
network **4**-45– **4**-63
OS-specific **4**-65– **4**-70
session table flood **4**-36
drop-no-rpf-route **4**-23
DS Codepoint Marking **2**-478, **2**-487, **2**-488
DSL **2**-517, **2**-522
dual Untrust interfaces **8**-103
Dynamic IP
See DIP
Dynamic IP pools
See DIP pools
dynamic packet filtering **4**-3
dynamic routing **1**-A-V, **2**-30

E

editing
address groups **2**-132
policies **2**-241
zones **2**-52
e-mail alert notification **3**-86, **3**-89, **3**-90
Encapsulating Security Payload
See ESP
encryption
algorithms **5**-8, **5**-44, **5**-48, **5**-52, **5**-57
definition **1**-A-VI
NSRP **8**-7, **8**-18
NSRP-Lite **8**-64
ESP **5**-3, **5**-7, **5**-8
authenticate only **5**-48
encrypt and authenticate **5**-48, **5**-56
encrypt only **5**-48
Ethernet, definition **1**-A-VI
evasion **4**-22– **4**-33
event log **3**-67
exe files, blocking **4**-168
exploits
See attacks
exporting routes **6**-28
external neighbors **1**-A-VI
extranet, definition **1**-A-VI

F

fail/pass mode, URL filtering **4**-116
fail-mode **4**-95
threshold **4**-96
failover
device **8**-130
dual Untrust interfaces **8**-104, **8**-105
object monitor **8**-132
serial interface **8**-123
virtual system **8**-144
VSD group **8**-131
filter list **1**-A-VI
filter source route **3**-125
filtering, packets **1**-A-VI
FIN scan **4**-22
FIN without ACK flag **4**-18
FIPS **1**-xxi
firewall, definition **1**-xxi, **1**-A-VII
fragment reassembly **4**-72– **4**-75
full-mesh configuration **8**-144
Function Zone Interfaces **2**-68
HA Interface **2**-69
Management Interface **2**-68

G

gatekeeper devices **2**-141
gateway
routing **1**-A-VIII
gateway (router) **1**-A-VII
global zone **2**-359
graphs, historical **2**-213
group
addresses **2**-129
services **2**-167
group expressions **2**-468– **2**-475
operators **2**-468
other group expressions **2**-469
server support **2**-372
user groups **2**-468
users **2**-468
group IKE ID
certificates **5**-238– **5**-249
preshared key **5**-250– **5**-258

group IKE ID user **5**-237– **5**-258
certificates **5**-238
preshared key **5**-250

H

H.323 protocol **2**-141
HA
active/active failover **8**-6
active/passive failover **8**-4
aggregate interfaces **8**-101
cabling **8**-45– **8**-48
cabling for dedicated HA interfaces **8**-45
cabling network interfaces as HA links **8**-47
control link **8**-38
data link **8**-41
DHCP **2**-508
dual Untrust interfaces **8**-103
HA LED **8**-25
IP tracking **8**-79, **8**-136
link probes **8**-42
messages **8**-40
path monitoring **8**-79
redundant interfaces **8**-94
secondary path **8**-25
serial interface **8**-118
Virtual HA Interface **2**-69
See also NSRP
hash-based message authentication code
See HMAC
hello interval **1**-A-VII
hello packet **1**-A-VII
High Availability
See HA
high-watermark threshold **4**-41
historical graphs **2**-213
HMAC **5**-7
hold time **1**-A-VII
Home zone **2**-61
HTTP **3**-5
blocking components **4**-167– **4**-169
keep-alive **4**-96
session ID **3**-5
session timeout **4**-41
trickling **4**-97
hubs, definition **1**-A-VII
Hypertext Transfer Protocol
See HTTP

I

ICMP
definition **1**-A-VII
fragments **4**-2
large packets **4**-4
ICMP flood **4**-59
ICMP services **2**-139
message code **2**-139
message type **2**-139
icons
defined **2**-216
policy **2**-216
Ident-Reset **3**-29
idle session timeout **2**-377
IEEE 802.1Q VLAN standard **7**-21
IKE **5**-9, **5**-77, **5**-91, **5**-201
group IKE ID user **5**-237– **5**-258
group IKE ID, container **5**-242
group IKE ID, wildcard **5**-241
heartbeats **5**-384
hello messages **5**-384
IKE ID **2**-431, **2**-452, **5**-44– **5**-46, **5**-53– **5**-54
IKE ID recommendations **5**-68
IKE ID, Windows 200 **5**-288
ISAKMP **1**-A-IX
key management **1**-A-IX
local ID, ASN1-DN **5**-240
Phase 1 proposals, predefined **5**-11
Phase 2 proposals, predefined **5**-14
proxy IDs **5**-14
redundant gateways **5**-382– **5**-400
remote ID, ASN1-DN **5**-240
shared IKE ID user **5**-259– **5**-267
user groups, defining **2**-434
users **2**-431– **2**-435
users, defining **2**-432
users, groups **2**-431
IKE users
IKE ID **2**-397, **2**-431
server support **2**-372
with other use types **2**-467
illustration
conventions **1**-xxx, **2**-xiii, **3**-ix, **4**-ix, **5**-ix, **6**-ix, **7**-vii, **8**-ix
importing routes **6**-28
inactive SA **3**-125
in-short error **3**-122
interfaces
addressing **2**-77
aggregate **2**-67, **8**-101
binding to zone **2**-76
dedicated **7**-15, **7**-33
default **2**-79
DIP **2**-171
dual Untrust **8**-103
exporting from vsys **7**-19
extended **5**-168
HA **2**-69
HA, dual **8**-38– **8**-41
importing to vsys **7**-18
L3 security zones **2**-77
loopback **2**-86
manageable **3**-34
management options **3**-29
MGT **2**-68
MIP **2**-331
modifying **2**-81
monitoring **8**-18
physical **2**-3
redundant **2**-67, **8**-94
secondary IP address **2**-84
serial **8**-118
shared **7**-15, **7**-33
tunnel **2**-49, **2**-69, **2**-70– **2**-73
tunnel, definition **1**-A-XVI
unbinding from zone **2**-80
viewing interface table **2**-74
VIP **2**-356
Virtual HA **2**-69, **8**-47
VSI **2**-68
VSIs **8**-28
internal flash storage **3**-66
Internet Key Exchange
See IKE
Internet, definition **1**-A-VIII
InterScan VirusWall **4**-90
intranet, definition **1**-A-IX

IP
definition **1**-A-VIII
packet fragments **4**-10
IP addresses
defining for each port **2**-127
definition **1**-A-VIII
extended **5**-168
host ID **2**-78
L3 security zones **2**-77– **2**-78
manage IP **3**-34
network ID **2**-78
NSM servers **3**-26
private **2**-77
private address ranges **2**-78
public **2**-77
secondary **2**-84
virtual **2**-356
IP options **4**-12– **4**-14
attributes **4**-12– **4**-14
incorrectly formatted **4**-6
loose source route **4**-13, **4**-31– **4**-33
record route **4**-13, **4**-14
security **4**-13, **4**-14
source route **4**-31
stream ID **4**-13, **4**-14
strict source route **4**-14, **4**-31– **4**-33
timestamp **4**-14
IP pools
See DIP pools
IP Security
See IPSec
IP spoofing **4**-22– **4**-30
drop-no-rpf-route **4**-23
Layer 2 **4**-24, **4**-29
Layer 3 **4**-23, **4**-25
IP tracking **8**-79, **8**-136
device failover threshold **8**-80
ping and ARP **8**-79, **8**-136
tracked IP failure threshold **8**-80, **8**-133
tunnel failover **8**-81
weights **8**-80
IP-based traffic classification **7**-33
IPSec **5**-3
AH **5**-2, **5**-47, **5**-56
AH, defined **1**-A-VI
authentication **1**-A-III
definition **1**-A-VIII
digital signature **5**-16
encryption **1**-A-VI
ESP **5**-2, **5**-47, **5**-56
ESP, defined **1**-A-VI
SAs **1**-A-XIII, **5**-2, **5**-10, **5**-11, **5**-13
SPI **5**-2
SPI, definition **1**-A-XIII
transport mode **5**-4, **5**-273, **5**-279, **5**-286
tunnel **5**-2
tunnel mode **5**-5
tunnel negotiation **5**-11
ISAKMP **1**-A-IX
ISP configuration for serial interface **8**-121

J

Java applets, blocking **4**-168

K

keep alive, BGP **1**-A-IX
keepalive
frequency, NAT-T **5**-303
L2TP **5**-283
keys
creating **3**-7
management **1**-A-IX

L

L2TP **5**-269– **5**-298
access concentrator, *See* LAC
address assignment **2**-460
compulsory configuration **5**-270
decapsulation **5**-275
default parameters **5**-276
encapsulation **5**-274
external auth server **2**-461
hello signal **5**-284
Keep Alive **5**-283, **5**-284
L2TP-only on Windows 2000 **5**-273
local database **2**-461
network server, *See* LNS
operational mode **5**-273
policies **2**-209
RADIUS server **5**-276
ScreenOS support **5**-273
SecurID server **5**-276
tunnel **5**-279
user authentication **2**-460
voluntary configuration **5**-270
Windows 2000 **5**-291
Windows 2000 tunnel authentication **5**-283
L2TP users **2**-460– **2**-464
point of authentication **2**-397
server support **2**-372
with XAuth **2**-467
L2TP-over-IPSec **5**-4, **5**-279, **5**-286
tunnel **5**-279
LAC **5**-270
NetScreen-Remote 5.0 **5**-270
Windows 2000 **5**-270
LAN, definition **1**-A-X
Land attack **4**-63
Last-Known-Good configuration
See LKG configuration
Layer 2 Tunneling Protocol
See L2TP
LDAP **2**-386– **2**-387
auth server object **2**-393
common name identifier **2**-387
distinguished name **2**-387
server port **2**-387
structure **2**-386
user types supported **2**-387
LED indicators, HA **8**-25
license keys **2**-536– **2**-537
Lightweight Directory Access Protocol
See LDAP
link state **1**-A-IX
link state advertisement **1**-A-IX
LKG (last-known-good) **2**-531
LKG configuration **2**-531
LNS **5**-270
load balancing
definition **1**-A-IX
load sharing **8**-144
local certificate **5**-22

local database **2**-374– **2**-375
IKE users **2**-431
timeout **2**-375
user types supported **2**-374
local preference **1**-A-X
logging **2**-213, **3**-66– **3**-81
asset recovery log **3**-81
CompactFlash (PCMCIA) **3**-66
console **3**-66
e-mail **3**-66
event log **3**-67
internal **3**-66
NSM reporting **3**-26
self log **3**-77
SNMP **3**-66, **3**-91
syslog **3**-66, **3**-87
WebTrends **3**-66, **3**-89
logging in
root admin **3**-50
Telnet **3**-10
vsys **7**-33, **7**-38
loopback interfaces **2**-86
loose source route IP option **4**-13, **4**-31– **4**-33
low-watermark threshold **4**-41

M

MAC address
definition **1**-A-X
Main Mode **5**-12
malicious URL protection **4**-72– **4**-75
manage IP **3**-34
VSD group 0 **8**-8
management client IP addresses **3**-49
Management information base II
See MIB II
Management interface
See MGT interface
management methods
CLI **3**-9
console **3**-21
SSL **3**-7
Telnet **3**-9
WebUI **3**-3
management options **3**-29
manageable **3**-34
NSM **3**-29
ping **3**-29
SCS **3**-29
SNMP **3**-29
SSL **3**-29
Telnet **3**-29
Transparent mode **3**-30
WebUI **3**-29
Manual Key **5**-131, **5**-142
management **5**-9
VPNs **3**-51, **3**-98
mapped IP
See MIP
MD5 **5**-7
definition **1**-A-X
MED comparison **1**-A-X
Message Digest version 5
See MD5
messages
alert **3**-67
critical **3**-67
debug **3**-67
emergency **3**-67
error **3**-67
info **3**-67
notice **3**-67
warning **3**-67
WebTrends **3**-90
MGT interface **2**-68
management options **3**-30
MIB files **3**-A-I
MIB files, importing **5**-325
MIB folders
primary **3**-A-II
MIB II **3**-29, **3**-91
MIP **2**-12, **2**-331
address range **2**-335
bidirectional translation **2**-252
creating addresses **2**-333
creating on tunnel interface **2**-341
creating on zone interface **2**-333
default netmask **2**-335
default virtual router **2**-335
definition **1**-A-X, **2**-252
global zone **2**-332
reachable from other zones **2**-336
same-as-untrust interface **2**-342– **2**-345
to zone with interface-based NAT **2**-112
virtual systems **7**-10
VPNs **5**-168
modem configuration for serial interface **8**-119
modem port **3**-22
modulus **5**-13
multi exit discriminator **1**-A-X
multimedia sessions, SIP **2**-156
multiple-type users **2**-467

N

names
conventions **1**-xxxi, **2**-xiv, **3**-x, **4**-x, **5**-x, **6**-x, **7**-viii, **8**-x
NAT
definition **1**-A-XI, **2**-246
IPSec and NAT **5**-301
NAT servers **5**-301
NAT-src with NAT-dst **2**-310– **2**-330
NAT mode **2**-110– **2**-117
interface settings **2**-113
traffic to Untrust zone **2**-91, **2**-112
NAT vector error **3**-125
NAT-dst **2**-276– **2**-330
address range **2**-250
address range to address range **2**-256, **2**-300
address range to single IP **2**-256, **2**-295
address shifting **2**-251, **2**-277, **2**-300
one-to-many translation **2**-291
one-to-one translation **2**-286
packet flow **2**-278– **2**-281
port mapping **2**-249, **2**-276, **2**-305
route considerations **2**-277, **2**-282– **2**-285
single IP with port mapping **2**-255
single IP, no port mapping **2**-255
unidirectional translation **2**-252, **2**-257
VPNs **5**-168
with MIPs or VIPs **2**-248
NAT-src **2**-246, **2**-259– **2**-275
address shifting **2**-267– **2**-272
address shifting, range considerations **2**-267

DIP pool with address shifting **2**-254
DIP pool with PAT **2**-253, **2**-260– **2**-263
DIP pool, fixed port **2**-253
DIP pools **2**-246
egress interface **2**-254, **2**-273– **2**-275
fixed port **2**-259, **2**-264– **2**-266
interface based **2**-247
port address translation **2**-247
Route mode Route mode
NAT-src **2**-118
unidirectional translation **2**-252, **2**-257
VPNs **5**-171
NAT-T **5**-301
enabling **5**-305
keepalive frequency **5**-303
NAT-Traversal
See NAT-T
negation, address **2**-237
neighbor **1**-A-XI
NetInfo **2**-501
netmasks **2**-206
definition **1**-A-XI, **1**-A-XIV
uses of **2**-78
NetScreen dictionary file **2**-381
NetScreen Redundancy Protocol
See NSRP
NetScreen Reliable Transport Protocol
See NRTP
NetScreen Security Manager
See NSM
NetScreen-Remote
AutoKey IKE VPN **5**-201
dynamic peer **5**-209, **5**-220
NAT-T option **5**-301
Network Address Translation (NAT) **3**-124
network layer reachability information **1**-A-XI
network, bandwidth **2**-478
NHTB table **5**-326– **5**-331
addressing scheme **5**-328
automatic entries **5**-331
manual entries **5**-330
mapping routes to tunnels **5**-327
NRTP **8**-33, **8**-76
NSM
Agent **3**-23, **3**-26
definition **3**-23
enabling the Agent **3**-25
initial connectivity setup **3**-24
management options **3**-29
Management System **3**-23, **3**-26
reporting events **3**-26, **3**-27
UI **3**-23
NSRP
ARP **8**-56
ARP broadcasts **8**-18
backup **8**-4
cabling **8**-45– **8**-48
clear cluster command **8**-16, **8**-63
cluster name **8**-17, **8**-63
clusters **8**-16– **8**-20, **8**-49
config sync **8**-33
configuration rollback **2**-533
control link **8**-38
control messages **8**-38, **8**-39
data link **8**-41
data messages **8**-40
debug cluster command **8**-16, **8**-63
default settings **8**-9, **8**-61
DHCP **2**-508
DIP groups **2**-189– **2**-192
files, sync **8**-34
full-mesh configuration **8**-45, **8**-144
HA cabling, dedicated interfaces **8**-45
HA cabling, network interfaces **8**-47
HA interfaces **8**-39
HA LED **8**-25
HA ports, redundant interfaces **8**-94
HA session backup **2**-212, **8**-21
hold-down time **8**-51, **8**-55
interface monitoring **8**-18
load sharing **8**-144
manage IP **8**-80, **8**-136
master **8**-4
NAT and Route modes **8**-8
NSRP, defined **1**-A-XI
NTP synchronization **2**-543, **8**-37
overview **8**-3
packet forwarding and dynamic routing **8**-41
port failover **8**-94
port monitoring **8**-134
preempt mode **8**-23
priority numbers **8**-23
redundant interfaces **2**-67
redundant ports **8**-38
RTO states **8**-22
RTOs **1**-A-XIII, **8**-21– **8**-22, **8**-49
RTOs, sync **8**-34
secondary path **8**-18, **8**-25
secure communications **8**-7, **8**-18
synchronization, PKI **8**-34
Transparent mode **8**-8
virtual systems **8**-144– **8**-150
VSD groups **8**-5, **8**-23– **8**-27, **8**-49, **8**-79
VSD, defined **1**-A-XVII
VSI **1**-A-XVII
VSIs **2**-68, **8**-5
VSIs, static routes **8**-28, **8**-99, **8**-100
NSRP-Lite **8**-57– **8**-78
cabling **8**-68
clusters **8**-60– **8**-63
config synchronization **8**-76
disabling synchronization **8**-78
file synchronization **8**-77
port monitoring **8**-79
preempt mode **8**-67
secure communications **8**-64
VSD groups **8**-65– **8**-67
NTP **2**-542– **2**-545
authentication types **2**-545
max time adjustment **2**-542
maximum time adjustment **2**-542
multiple servers **2**-542
NSRP synchronization **2**-543, **8**-37
secure servers **2**-545
server configuration **2**-544
servers **2**-542

O

object monitoring **8**-132
OCSP (Online Certificate Status Protocol) **5**-36
client **5**-36
responder **5**-36
operating system **3**-9
OSPF
areas **6**-34
assigning interface to area **6**-42
authenticating neighbors **6**-60

backup designated router **6**-36
broadcast network **6**-36
configuration steps **6**-38
creating instance in VR **6**-39
defining area **6**-41
designated router **6**-36
enabling on interface **6**-44
filtering neighbors **6**-62
global parameters **6**-51
hello protocol **6**-35
interface parameters **6**-57
link-state advertisements **6**-34, **6**-37
not so stubby area **6**-35
point-to-point network **6**-36
protecting against flooding **6**-64
redistributing routes **6**-49
rejecting default routes **6**-63
router adjacency **6**-35
router types **6**-35
security configuration **6**-60
stub area **6**-35
summarizing redistributed routes **6**-50
virtual links **6**-53

P

packet filtering **1**-A-VI
packet flow **2**-11– **2**-13
inbound VPN **5**-63– **5**-64
NAT-dst **2**-278– **2**-281
outbound VPN **5**-61– **5**-62
policy-based VPN **5**-65– **5**-66
route-based VPN **5**-60– **5**-64
packets **3**-125
address spoofing attack **3**-123
collision **3**-122
denied **3**-125
dropped **3**-124, **3**-125
fragmented **3**-125
incoming **3**-122
Internet Control Message Protocol (ICMP) **3**-120, **3**-123
IPSec **3**-123
land attack **3**-124
Network Address Translation (NAT) **3**-124
Point to Point Tunneling Protocol (PPTP) **3**-123
received **3**-121, **3**-122, **3**-123, **3**-125
transmitted underrun **3**-122
unreceivable **3**-122
unroutable **3**-124
PAP **5**-273, **5**-276
parent connection **3**-124
password
forgetting **3**-44
root admin **3**-47
vsys admin **7**-38
Password Authentication Protocol
See PAP
PAT **2**-172, **2**-259
path monitoring **8**-79
tunnel failover **8**-81
PC card **2**-528, **2**-530
PCMCIA **3**-66
peer **1**-A-XI
Perfect Forward Secrecy
See PFS
PFS **5**-14, **5**-46, **5**-55
Phase 1 **5**-11
proposals **5**-11
proposals, predefined **5**-11
Phase 2 **5**-13
proposals **5**-13
proposals, predefined **5**-14
ping
management options **3**-29
Ping of Death **4**-65
pinholes **2**-161
PKI **5**-18
encryption **1**-A-VI
key **3**-7
Point-to-Point Protocol
See PPP
Point-to-Point Tunneling Protocol (PPTP) **3**-123
policies **2**-3
actions **2**-207
address groups **2**-206
address negation **2**-237
addresses **2**-206
addresses in **2**-206
alarms **2**-213
antivirus scanning **2**-214
application **2**-207
authentication **2**-210
bidirectional VPNs **2**-208, **2**-216, **5**-143
changing **2**-241
context **4**-127
core section **4**-126
counting **2**-213
Deep Inspection **2**-209
definition **1**-A-XII
deny **2**-207
DIP groups **2**-190
disabling **2**-241
enabling **2**-241
functions of **2**-197
global **2**-201, **2**-217, **2**-234
HA session backup **2**-212
icons **2**-216
ID **2**-206
internal rules **2**-203
interzone **2**-200, **2**-217, **2**-218, **2**-223
intrazone **2**-201, **2**-217, **2**-231
L2TP **2**-209
L2TP tunnels **2**-209
location **2**-218
lookup sequence **2**-202
management **2**-216
managing bandwidth **2**-478
maximum limit **2**-130
multiple items per component **2**-236
name **2**-208
NAT-dst **2**-210
NAT-src **2**-210
order **2**-243
permit **2**-207
policy context **2**-235
policy set lists **2**-202
policy verification **2**-242
position at top **2**-209, **2**-243
removing **2**-244
reordering **2**-243
required elements **2**-199
root system **2**-203
schedules **2**-214
security zones **2**-206
service book **2**-134

service groups **2**-167
services **2**-206
services in **2**-134, **2**-206
shadowing **2**-242
traffic logging **2**-213
traffic shaping **2**-215
tunnel **2**-207
types **2**-200– **2**-201
URL filtering **2**-213, **4**-118
virtual systems **2**-203
VPN dialup user groups **2**-206
VPNs **2**-208
policy-based NAT
See NAT-dst and NAT-src
tunnel interfaces **2**-69
policy-based VPNs **5**-58
Port Address Translation
See PAT
port mapping **2**-249, **2**-276
port modes **2**-55
port scan **4**-10
ports
modem **3**-22
monitoring **8**-79, **8**-134
port failover **8**-94
port numbers **2**-366
primary trusted and untrusted **8**-94
redundant **8**-38
secondary trusted and untrusted **8**-94
trunk **7**-23
trunk ports **1**-A-XV
PPP **5**-271
preempt mode **8**-23, **8**-67
prefixes
defined **1**-A-XII
preshared key **5**-9, **5**-201
priority queuing **2**-485
private addresses **2**-78
probes
network **4**-8
open ports **4**-10
operating systems **4**-16– **4**-20
proposals
Phase 1 **5**-11, **5**-67
Phase 2 **5**-13, **5**-67
protocol anomalies **4**-139
protocol distribution
reporting to NSM **3**-26
protocols
CHAP **5**-273
NRTP **8**-33, **8**-76
NSRP **8**-1, **8**-57
PAP **5**-273
PPP **5**-271
VRRP **8**-79, **8**-136
proxy IDs **5**-14
matching **5**-67
VPNs and NAT **5**-168– **5**-169
proxy servers **1**-A-III
public addresses **2**-77
Public key infrastructure
See PKI
Public/private key pair **5**-19

Q

QoS **1**-xxi, **2**-478
Quality-of-service
See QoS

R

RADIUS **2**-379– **2**-381, **3**-44
access-challenge **1**-A-I
auth server object **2**-388
L2TP **5**-276
NetScreen dictionary file **2**-465
object properties **2**-380
port **2**-380
retry timeout **2**-380
shared secret **2**-380
reconnaissance **4**-7– **4**-33
address sweep **4**-8
FIN scan **4**-22
IP options **4**-12
port scan **4**-10
SYN and FIN flags set **4**-16
TCP packet without flags **4**-20
record route IP option **4**-13, **4**-14
redistribution **1**-A-XII
redistribution list **1**-A-XII
redundant gateways **5**-382– **5**-400
recovery procedure **5**-385
TCP SYN flag checking **5**-388
regular expressions **4**-157– **4**-159
rekey option, VPN monitoring **5**-308
Remote authentication dial in user service
See RADIUS
replay protection **5**-14
reset to factory defaults **3**-48
RFCs
1349, "Type of Service in the Internet Protocol Suite" **2**-215
1777, "Lightweight Directory Access Protocol" **2**-386
1918, “Address Allocation for Private Internets” **2**-78
2474, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers" **2**-215
RIP
authenticating neighbors **6**-80
configuration steps **6**-69
creating instance in VR **6**-70
enabling on interface **6**-72
filtering neighbors **6**-82
global parameters **6**-76
interface parameters **6**-78
protecting against flooding **6**-84
protocol overview **6**-68
redistributing routes **6**-73
rejecting default routes **6**-83
security configuration **6**-80
rollback, configuration **2**-531– **2**-532
route filtering **6**-24
route flap dampening **1**-A-XII
route map **1**-A-XII, **6**-22
route metric **6**-17
Route mode **2**-118– **2**-123
interface settings **2**-119
route redistribution **1**-A-XII, **6**-21
route reflector **1**-A-XIII
route-based VPNs **5**-58
routers, definition **1**-A-XIII
routing **2**-30
between secondary IP addresses **2**-84
route preference **6**-15

route selection **6**-15
routing table **1**-A-XIII, **2**-31
route selection **6**-15
RSH ALG **2**-140
RTOs **8**-21– **8**-22
operational states **8**-22
RTO peer **8**-24
rules, derived from polices **2**-203
run-time authentication **2**-211, **2**-399
run-time objects
See RTOs

S

SA policy **3**-125
SAs **5**-10, **5**-11, **5**-13
check in packet flow **5**-62
definition **1**-A-XIII
SCEP (Simple Certificate Enrollment Protocol) **5**-30
schedules **2**-193, **2**-214
SCREEN
address sweep **4**-8
bad IP options, drop **4**-6
drop unknown MAC addresses **4**-51
FIN with no ACK **4**-22
FIN without ACK flag, drop **4**-18
ICMP flood **4**-59
ICMP fragments, block **4**-2
IP options **4**-12
IP packet fragments, block **4**-10
IP spoofing **4**-22– **4**-30
Land attack **4**-63
large ICMP packets, block **4**-4
loose source route IP option, detect **4**-33
MGT zone **2**-48
Ping of Death **4**-65
port scan **4**-10
source route IP option, deny **4**-33
strict source route IP option, detect **4**-33
SYN and FIN flags set **4**-16
SYN flood **4**-45– **4**-51
SYN fragments, detect **4**-12– **4**-13
SYN-ACK-ACK proxy flood **4**-43
TCP packet without flags, detect **4**-20
Teardrop **4**-67
UDP flood **4**-61
unknown protocols, drop **4**-8
VLAN and MGT zones **4**-3
WinNuke attack **4**-69
ScreenOS **1**-xxii
function zones **2**-54
global zone **2**-48
Home-Work zone **2**-61
interfaces physical **2**-3
overview **2**-1– **2**-27
packet flow **2**-11– **2**-13
policies **2**-3
port modes **2**-55
security zone interfaces **2**-3
security zones **2**-2, **2**-48
security zones, global **2**-2
security zones, predefined **2**-2
subinterfaces **2**-4
tunnel zones **2**-49
updating **2**-530
virtual systems **2**-10
virtual systems, VRs **7**-6
virtual systems, zones **7**-7
zones **2**-45– **2**-54
SCS **3**-29
SDP **2**-159– **2**-160
secondary IP addresses **2**-84
secondary path **8**-18, **8**-25
Secure Hash Algorithm-1
See SHA-1
Secure Sockets Layer
See SSL
SecurID **2**-384– **2**-385
ACE server **2**-384
auth server object **2**-391
authentication port **2**-385
authenticator **2**-384
client retries **2**-385
client timeout **2**-385
duress **2**-385
encryption type **2**-385
L2TP **5**-276
token code **2**-384
user type support **2**-385
security association
See SAs
Security Associations (SA) **3**-124
security IP option **4**-13, **4**-14
security zones **1**-A-XIII, **2**-2
destination zone determination **2**-13
global **2**-2
interfaces **2**-3, **2**-66
physical interfaces **2**-66
predefined **2**-2
See zones
source zone determination **2**-12
subinterfaces **2**-66
self log **3**-77
serial cables **3**-21
serial interface **8**-118
failover **8**-123
ISP configuration **8**-121
modem configuration **8**-119
service book
adding service **2**-136
custom service **2**-134
custom service (CLI) **2**-136
modifying entries (CLI) **2**-138
modifying entries (Web UI) **2**-169
pre-configured services **2**-134
removing entries (CLI) **2**-138
service groups (Web UI) **2**-167
service groups **2**-167– **2**-170
creating **2**-168
deleting **2**-170
modifying **2**-169
services **2**-134
custom **4**-152
custom ALGs **2**-207
defined **2**-206
drop-down list **2**-134
ICMP **2**-139
in policies **2**-206
modifying timeout **2**-136
timeout threshold **2**-135
session ID **3**-5
Session Initiation Protocol
See SIP
session limits **4**-36– **4**-40
destination based **4**-37, **4**-40
source based **4**-36, **4**-39
session table flood **4**-36

session timeout
HTTP **4**-41
idle timeout **2**-377
TCP **4**-41
UDP **4**-41
settings
downloading **2**-528
importing **2**-528
saving **2**-528
uploading **2**-528
SHA-1 **5**-7
definition **1**-A-XIII
shadowed policies **2**-242
SIP **2**-156– **2**-166
ALG **2**-159, **2**-163
connection information **2**-160
defined **2**-156
inactivity timeouts **2**-163
media announcements **2**-160
media inactivity timeout **2**-163, **2**-166
messages **2**-156
multimedia sessions **2**-156
pinholes **2**-159
request method types **2**-157
Request Methods **2**-157
response codes **2**-158
response types **2**-157
responses **2**-157
RTCP **2**-160
RTP **2**-160
SDP **2**-159– **2**-160
session inactivity timeout **2**-163
signaling **2**-159
signaling inactivity timeout **2**-163, **2**-166
SMTP server IP **3**-86
SNMP **3**-29, **3**-91
cold start trap **3**-91
community, private **3**-95
community, public **3**-95
configuration **3**-95
encryption **3**-94, **3**-97
implementation **3**-94
management options **3**-29
MIB files **3**-A-I
MIB files, importing **5**-325
MIB folders, primary **3**-A-II
system alarm traps **3**-91
traffic alarm traps **3**-91
trap types **3**-92
traps **3**-91
VPN monitoring **5**-325
SNMP traps
100, hardware problems **3**-92
200, firewall problems **3**-92
300, software problems **3**-92
400, traffic problems **3**-92
500, VPN problems **3**-92
allow or deny **3**-94
software
key, vsys **7**-15
updating **2**-530
uploading and downloading **2**-530
source route **3**-125
source-based routing **6**-17
SPI
definition **1**-A-XIII
SSH **3**-11– **3**-17
authentication method priority **3**-17
automated logins **3**-19
connection procedure **3**-12
forcing PKA authentication only **3**-17
host key **3**-12
loading public keys, CLI **3**-16
loading public keys, TFTP **3**-16, **3**-19
loading public keys, WebUI **3**-16
password authentication **3**-15
PKA **3**-15
PKA authentication **3**-15
PKA key **3**-12
server key **3**-12
session key **3**-12
SSL **3**-7
management options **3**-29
with WebAuth **2**-427
SSL Handshake Protocol
See SSLHP
SSLHP **3**-7
stateful inspection **4**-3
stateful signatures **4**-138
definition **4**-138
static routing **1**-A-XIV, **2**-30, **2**-33– **2**-44
configuring **2**-38
using **2**-36
statistics
reporting to NSM **3**-27
stream ID IP option **4**-13, **4**-14
strict source route IP option **4**-14, **4**-31– **4**-33
subinterfaces **2**-4, **7**-23
configuring (vsys) **7**-23
creating (root system) **2**-82
creating (vsys) **7**-23
defined **1**-A-XIV
defining **7**-25
deleting **2**-83
multiple subinterfaces per vsys **7**-23
subnet masks
definition **1**-A-XIV
subscriptions
registration and activation **2**-538– **2**-540
temporary service **2**-538
support certificate **2**-539, **2**-540
SYN and FIN flags set **4**-16
SYN flood **4**-45– **4**-51
alarm threshold **4**-49
attack **4**-45
attack threshold **4**-49
destination threshold **4**-50
drop unknown MAC addresses **4**-51
queue size **4**-51
source threshold **4**-50
threshold **4**-46
timeout **4**-51
SYN fragments **4**-12– **4**-13
SYN-ACK-ACK proxy flood **4**-43
synchronization
configuration **8**-33
files **8**-34
PKI objects **8**-34
RTOs **8**-34
syslog **3**-66
encryption **3**-97
facility **3**-88, **3**-90, **3**-101, **3**-112
host **3**-87
host name **3**-88, **3**-89, **3**-90, **3**-101, **3**-112

messages **3**-87
port **3**-88, **3**-101, **3**-112
security facility **3**-88, **3**-90, **3**-101, **3**-112
system clock **2**-541– **2**-545
date & time **2**-541
sync with client **2**-541
time zone **2**-541
system, parameters **2**-493– **2**-544

T

TCP
max simultaneous connections **4**-94
packet without flags **4**-20
proxy **3**-125
session timeout **4**-41
stream signatures **4**-164
SYN flag checking **5**-388
three-way handshake **1**-A-XV
TCP/IP, definition **1**-A-XV
Teardrop attack **4**-67
Telnet **3**-9, **3**-29
TFTP server **2**-528, **2**-530
three-way handshake **1**-A-XV, **4**-45
time zone **2**-541
timeout
admin user **2**-378
auth user **2**-377
timestamp IP option **4**-14
token code **2**-384
trace-route **2**-98, **2**-101
traffic
alarms **3**-82– **3**-86
classification, IP-based **7**-33
classification, VLAN-based **7**-21
counting **2**-213
logging **2**-213
priority **2**-215
shaping **2**-478
through traffic, vsys sorting **7**-11– **7**-14
traffic shaping **1**-xxi, **2**-477– **2**-491
automatic **2**-478
interface requirement **2**-478
service priorities **2**-485
Transparent mode **2**-92– **2**-109
ARP/trace-route **2**-96
blocking non-ARP traffic **2**-94
blocking non-IP traffic **2**-94
broadcast traffic **2**-94
drop unknown MAC addresses **4**-51
flood **2**-96
management options **3**-30
routes **2**-94
unicast options **2**-96
transport mode **5**-4, **5**-273, **5**-279, **5**-286
Triple DES
See 3DES
trunk ports **7**-23
defined **7**-22
definition **1**-A-XV
manually setting **7**-22
trust **1**-A-XV
tunnel interfaces **2**-69
definition **1**-A-XVI, **2**-69
policy-based NAT **2**-69
tunnel mode **5**-5
tunnel zones
definition **1**-A-XVI

U

UDP
checksum **5**-303
definition **1**-A-XVI
NAT-T encapsulation **5**-301
session timeout **4**-41
UDP flood **4**-61
unknown protocols **4**-8
unknown unicast options **2**-95– **2**-101
ARP **2**-98– **2**-101
flood **2**-96– **2**-97
trace-route **2**-98, **2**-101
untrust **1**-A-XVI
URL filtering **2**-213, **4**-113– **4**-121
blocked URL message type **4**-117
communication timeout **4**-116
device-level activation **4**-117
fail/pass mode **4**-116
NetScreen blocked URL message **4**-117
policy-level application **4**-118
routing **4**-119
server status **4**-118
servers per vsys **4**-115
Websense server name **4**-116
Websense server port **4**-116
URL, definition **1**-A-XVI
user authentication
See authentication, users
users
group IKE ID **5**-237– **5**-258
groups, server support **2**-372
IKE **2**-431– **2**-435
IKE, groups **2**-434
multiple administrative users **3**-37
shared IKE ID **5**-259– **5**-267
users, admin **2**-465– **2**-466
auth process **2**-466
timeout **2**-378
users, IKE
defining **2**-432
groups **2**-431
IKE ID **2**-431
users, L2TP **2**-460– **2**-464
users, XAuth **2**-436– **2**-458

V

Valicert **5**-36
vendor-specific attributes
See VSAs
Verisign **5**-36
VIP **2**-12
bidirectional translation **2**-252
configuring **2**-359
custom and multi-port services **2**-363– **2**-369
custom services, low port numbers **2**-357
definition **1**-A-XVI, **2**-252
editing **2**-362
global zone **2**-359
reachable from other zones **2**-359
removing **2**-362
required information **2**-357
to zone with interface-based NAT **2**-112
virtual systems **7**-10
virtual adapter **2**-436
definition **1**-A-XVI

Virtual HA interface **2**-69, **8**-47
Virtual IP
See VIP
virtual private network
See VPNs
virtual routers
See VRs
virtual security device groups
See VSD groups
virtual security interface
See VSI
virtual system **2**-10, **7**-1– **7**-39
admin types **7**-3
administrators **3**-38
admins **7**-iii, **7**-1
basic functional requirements **7**-3
changing admin's password **7**-3, **7**-38
creating a vsys object **7**-3
definition **1**-A-XVII
exporting a physical interface **7**-19
failover **8**-144
importing a physical interface **7**-18
interfaces **7**-8
IP-based traffic classification **7**-33– **7**-37
load sharing **8**-144
manageability and security **7**-34
MIP **7**-10
NSRP **8**-144
overlapping address ranges **7**-25, **7**-34
overlapping subnets **7**-25
read-only admins **3**-38
shared VR **7**-15
shared zone **7**-15
software key **7**-15
traffic sorting **7**-10– **7**-17
Transparent mode **7**-22
VIP **7**-10
VLAN-based traffic classification **7**-21– **7**-32
VRs **7**-6
zones **7**-7
VLAN zone **2**-93
VLAN1
Interface **2**-93, **2**-102
management options **3**-30
Zones **2**-93
VLANs
communicating with another VLAN **7**-28– **7**-32
creating **7**-25– **7**-27
definition **1**-A-XVII
subinterfaces **7**-23
tag **7**-23, **7**-24
tags **1**-XIV, **2**-4
Transparent mode **7**-22, **7**-23
trunking **7**-22
VLAN-based traffic classification **7**-21
voice-over IP communication **2**-141
VPN monitoring **5**-307– **5**-322
destination address **5**-308– **5**-311
destination address, XAuth **5**-309
ICMP echo requests **5**-325
outgoing interface **5**-308– **5**-311
policies **5**-310
rekey option **5**-308, **5**-331
routing design **5**-323
SNMP **5**-325
status changes **5**-307, **5**-310
VPNs **1**-xxi
Aggressive mode **5**-12
AutoKey IKE **3**-51, **3**-98, **5**-9
configuration tips **5**-67– **5**-68
cryptographic options **5**-40– **5**-57
definition **1**-A-XVII
Diffie-Hellman exchange **5**-13
Diffie-Hellman groups **5**-13
for administrative traffic **3**-97
FQDN aliases **5**-152
FQDN for gateway **5**-151– **5**-167
idletime **2**-439
Main mode **5**-12
Manual Key **3**-51, **3**-98
MIP **5**-168
multiple tunnels per tunnel interface **5**-326– **5**-381
NAT for overlapping addresses **5**-168– **5**-185
NAT-dst **5**-168
NAT-src **5**-171
packet flow **5**-60– **5**-66
Phase 1 **5**-11
Phase 2 **5**-13
policies **2**-208
proxy IDs, matching **5**-67
redundant gateways **5**-382– **5**-400
redundant groups, recovery procedure **5**-385
replay protection **5**-14
route- vs policy-based **5**-58
SAs **5**-10
to zone with interface-based NAT **2**-112
tunnel always up **5**-308
tunnel zones **2**-49
tunneling, definition **1**-A-XV
VPN groups **5**-382
VPN monitoring and rekey **5**-308
VRRP **8**-79, **8**-136
VRs **2**-35, **6**-3– **6**-28
access lists **6**-24
BGP **6**-91– **6**-101
creating a shared VR **7**-16
custom **6**-7
definition **1**-A-XVII
exporting routes **6**-28
forwarding traffic between **2**-5, **6**-4
importing routes **6**-28
introduction **2**-5
maximum routing table entries **6**-14
modifying **6**-12
on vsys **6**-9
OSPF **6**-38– **6**-65
predefined **6**-3
RIP **6**-69– **6**-86
route filtering **6**-24
route map **6**-22
route metric **6**-17
route preference **6**-15
route redistribution **6**-21
route selection **6**-15
router ID **6**-12
shared **7**-15
source-based routing **6**-17
using two VRs **6**-3, **6**-4
VSAs **2**-381
attribute name **2**-381
attribute number **2**-381
attribute type **2**-381
vendor ID **2**-381

VSD groups **8**-5, **8**-23– **8**-27, **8**-65– **8**-67
failover **8**-131
heartbeats **8**-18, **8**-25, **8**-66
hold-down time **8**-51, **8**-55
member states **8**-24, **8**-65– **8**-66, **8**-79
priority numbers **8**-23
VSD, defined **1**-A-XVII
VSIs **8**-5, **8**-23, **8**-65
defined **1**-A-XVII
multiple VSIs per VSD group **8**-144
static routes **8**-28

W

Web browser requirements **3**-3
Web user interface
See WebUI
WebAuth **2**-372
external user group **2**-423
local user group **2**-420
pre-policy auth process **2**-212, **2**-400
with SSL (external user group) **2**-427
Websense **4**-113
WebTrends **3**-66, **3**-89
encryption **3**-89, **3**-97
messages **3**-90
WebUI **3**-3, **3**-30, **3**-31
conventions **1**-xxviii, **2**-xi, **3**-vii, **5**-vii, **6**-vii, **7**-v, **8**-vii
wildcard **5**-241
WinNuke attack **4**-69
WINS
definition **1**-A-XVII
L2TP settings **5**-276
Work zone **2**-61

X

XAuth
address assignments **2**-436, **2**-438
address timeout **2**-438
auth and address **2**-452
client authentication **2**-458
defined **2**-436
external auth server queries **2**-437
external user auth **2**-444
external user group auth **2**-447
IP address lifetime **2**-438– **2**-439
lifetime **2**-439
local user auth **2**-440
local user group auth **2**-442
query remote settings **2**-437
ScreenOS as client **2**-458
TCP/IP assignments **2**-437
user authentication **2**-436
virtual adapter **2**-436
VPN idletime **2**-439
VPN monitoring **5**-309
XAuth users **2**-436– **2**-458
point of authentication **2**-397
server support **2**-372
with L2TP **2**-467

Z

zip files, blocking **4**-168
zombie agent **4**-35, **4**-37
zones **2**-45– **2**-54
definition **1**-A-XVII
function **2**-54
global **2**-48, **2**-359
Layer 2 **2**-93
security **1**-A-XIII, **2**-48
shared **7**-15
tunnel **1**-A-XVI, **2**-49
VLAN **2**-54, **2**-93
vsys **7**-7

NetScreen Concepts & Examples

ScreenOS Reference Guide

Volume 2: Fundamentals

ScreenOS 5.0.0

P/N 093-0925-000

Rev. B

Copyright Notice

NetScreen Technologies, Inc.
Building #3
805 11th Avenue
Sunnyvale, CA 94089
www.netscreen.com

FCC Statement

The following information is for FCC compliance of Class A devices: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. The equipment generates, uses, and can radiate radio-frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case users will be required to correct the interference at their own expense.

The following information is for FCC compliance of Class B devices: The equipment described in this manual generates and may radiate radio-frequency energy. If it is not installed in accordance with NetScreen's installation instructions, it may cause interference with radio and television reception. This equipment has been tested and found to comply with the limits for a Class B digital device in accordance with the specifications in part 15 of the FCC rules. These specifications are designed to provide reasonable protection against such interference in a residential installation. However, there is no guarantee that interference will not occur in a particular installation.

If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Consult the dealer or an experienced radio/TV technician for help.
- Connect the equipment to an outlet on a circuit different from that to which the receiver is connected.

Caution: Changes or modifications to this product could void the user's warranty and authority to operate this device.

Disclaimer

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR NETSCREEN REPRESENTATIVE FOR A COPY.

Contents

Preface ix
Conventions x
CLI Conventions x
WebUI Conventions xi
Illustration Conventions xiii
Naming Conventions and Character Types xiv
NetScreen Documentation xv
Chapter 1 ScreenOS Architecture 1
Security Zones 2
Security Zone Interfaces 3
Physical Interfaces 3
Subinterfaces 4
Virtual Routers 5
Policies 6
VPNs 8
Virtual Systems 10
Packet Flow Sequence 11
Example (Part 1): Enterprise with Six Zones 14
Example (Part 2): Interfaces for Six Zones 16
Example (Part 3): Enterprise with Two Routing Domains 20
Example (Part 4): Policies for an Enterprise with Six Zones 22
Chapter 2 Routing Tables and Static Routing 29
Routing Essentials 30
Routing Methods 30
Static Routing 30
Dynamic Routing 30
Routing Tables 31
Routing with Static Routes 33
Virtual Routers on NetScreen Devices 35
When to Configure Static Routes 36
Configuring Static Routes 38
Example: Configuring Static Routes 39
Example: Static Route through a Tunnel Interface 43
Chapter 3 Zones 45
Security Zones 48
Global Zone 48
SCREEN Options 48
Tunnel Zones 49
Example: Binding a Tunnel Interface to a Tunnel Zone 50
Configuring Security Zones and Tunnel Zones 51
Creating a Zone 51
Modifying a Zone 52
Deleting a Zone 53
Function Zones 54
Null Zone 54
MGT Zone 54
HA Zone 54
Self Zone 54
VLAN Zone 54

Port Modes 55
Setting the Port Mode on NetScreen Appliances 59
Example: Setting Home-Work Port Mode 60
Home Zone/Work Zone 61
Example: Configuring Home and Work Zones 63

Chapter 4 Interfaces 65

Interface Types 66
Security Zone Interfaces 66
Physical 66
Subinterface 66
Aggregate Interfaces 67
Redundant Interfaces 67
Virtual Security Interfaces 68
Function Zone Interfaces 68
Management Interface 68
HA Interface 69
Tunnel Interfaces 69
Deleting Tunnel Interfaces 72
Example: Deleting a Tunnel Interface 72
Viewing Interfaces 74
Interface Table 74
Configuring Security Zone Interfaces 76
Binding an Interface to a Security Zone 76
Example: Binding an Interface 76
Defining an Address for a L3 Security Zone Interface 77
Public IP Addresses 77
Private IP Addresses 78
Example: Addressing an Interface 79
Unbinding an Interface from a Security Zone 80
Example: Unbinding an Interface 80
Modifying Interfaces 81
Example: Modifying Settings on an Interface 81
Creating Subinterfaces 82
Example: Creating a Subinterface in the Root System 82
Deleting Subinterfaces 83
Example: Deleting a Security Zone Interface 83
Secondary IP Addresses 84
Secondary IP Address Properties 84
Example: Creating a Secondary IP Address 85
Loopback Interfaces 86
Example: Creating a Loopback Interface 86
Using Loopback Interfaces 87
Example: Using the Loopback Interface to Manage a Device 87
Example: Enabling BGP on a Loopback Interface 88
Example: Configuring NSRP VSIs on a Loopback Interface 88
Example: Specifying a Loopback Interface as a Source Interface 89

Chapter 5 Interface Modes 91

Transparent Mode 92
Zone Settings 93
VLAN Zone 93
Predefined Layer 2 Zones 93
Traffic Forwarding 94
Unknown Unicast Options 95
Flood Method 96
ARP/Trace-Route Method 98
Example: VLAN1 Interface for Management 102
Example: Transparent Mode 105

NAT Mode 110
Inbound and Outbound NAT Traffic 112
Interface Settings 113
Example: NAT Mode 114
Route Mode 118
Interface Settings 119
Example: Route Mode 120
Chapter 6 Building Blocks for Policies 125
Addresses 126
Address Entries 127
Example: Adding Addresses 127
Example: Modifying Addresses 128
Example: Deleting Addresses 129
Address Groups 129
Example: Creating an Address Group 131
Example: Editing a Group Address Entry 132
Example: Removing an Address Group Member and a Group 133
Services 134
Predefined Services 134
Example: Setting a Predefined Service Timeout 136
Custom Services 136
Example: Adding a Custom Service 136
Example: Modifying a Custom Service 138
Example: Removing a Custom Service 138
ICMP Services 139
Example: Defining an ICMP Service 140
RSH ALG 140
H.323 Protocol for Voice-over-IP 141
Example: Gatekeeper in the Trust Zone (Transparent or Route Mode) 141
Example: Gatekeeper in the Trust Zone (NAT Mode) 143
Example: Gatekeeper in the Untrust Zone (Transparent or Route Mode) 148
Example: Gatekeeper in the Untrust Zone (NAT Mode) 151
SIP – Session Initiation Protocol 156
SIP Request Methods 157
Classes of SIP Responses 157
ALG – Application-Layer Gateway 159
SDP 160
Pinhole Creation 161
Session Inactivity Timeout 163
Example: Creating a Policy to Permit SIP 164
Example: Signaling and Media Inactivity Timeouts 166
Service Groups 167
Example: Creating a Service Group 168
Example: Modifying a Service Group 169
Example: Removing a Service Group 170
DIP Pools 171
Port Address Translation 172
Example: Creating a DIP Pool with PAT 172
Example: Modifying a DIP Pool 174
Sticky DIP Addresses 174
Extended Interface and DIP 175
Example: Using DIP in a Different Subnet 175
Loopback Interface and DIP 183
Example: DIP on a Loopback Interface 184
DIP Groups 189
Example: DIP Group 191
Schedules 193
Example: Recurring Schedule 193

Chapter 7 Policies 197
Basic Elements 199
Three Types of Policies 200
Interzone Policies 200
Intrazone Policies 201
Global Policies 201
Policy Set Lists 202
Policies Defined 203
Policies and Rules 203
Anatomy of a Policy 205
ID 206
Zones 206
Addresses 206
Services 206
Action 207
Application 207
Name 208
VPN Tunneling 208
L2TP Tunneling 209
Deep Inspection 209
Placement at the Top of the Policy List 209
Source Address Translation 210
Destination Address Translation 210
User Authentication 210
HA Session Backup 212
URL Filtering 213
Logging 213
Counting 213
Traffic Alarm Threshold 213
Schedules 214
Antivirus Scanning 214
Traffic Shaping 215
Policies Applied 216
Viewing Policies 216
Policy Icons 216
Creating Policies 217
Policy Location 218
Example: Interzone Policies for E-Mail Service 218
Example: Interzone Policy Set 223
Example: Intrazone Policies 231
Example: Global Policy 234
Entering a Policy Context 235
Multiple Items per Policy Component 236
Address Negation 237
Example: Destination Address Negation 237
Modifying and Disabling Policies 241
Policy Verification 242
Reordering Policies 243
Removing a Policy 244
Chapter 8 Address Translation 245
Introduction to Address Translation 246
Policy-Based Translation Options 253
Directional Nature of NAT-Src and NAT-Dst 257
Source Network Address Translation 259
NAT-Src from a DIP Pool with PAT Enabled 260
Example: NAT-Src with PAT Enabled 261
NAT-Src from a DIP Pool with PAT Disabled 264
Example: NAT-Src with PAT Disabled 264
NAT-Src from a DIP Pool with Address Shifting 267
Example: NAT-Src with Address Shifting 268
NAT-Src from the Egress Interface IP Address 273
Example: NAT-Src without DIP 273

Destination Network Address Translation 276
Packet Flow for Destination Translation 278
Routing for Destination Translation 282
Addresses Connected to the Same Interface 283
Addresses Connected to the Same Interface but Separated by a Router 284
Addresses Separated by an Interface 285
NAT-Dst: One-to-One Mapping 286
Example: One-to-One Destination Translation 287
Translating from One Address to Multiple Addresses 291
Example: One-to-Many Destination Translation 291
NAT-Dst: Many-to-One Mapping 295
Example: Many-to-One Destination Translation 295
NAT-Dst: Many-to-Many Mapping 300
Example: Many-to-Many Destination Translation 301
NAT-Dst with Port Mapping 305
Example: NAT-Dst with Port Mapping 305
NAT-Src and NAT-Dst in the Same Policy 310
Example: NAT-Src and NAT-Dst Combined 310
Mapped IP Addresses 331
MIP and the Global Zone 332
Example: Adding a MIP to an Untrust Zone Interface 333
Example: Reaching a MIP from Different Zones 336
Example: Adding a MIP to a Tunnel Interface 341
MIP-Same-as-Untrust 342
Example: MIP on the Untrust Interface 343
MIP and the Loopback Interface 346
Example: MIP for Two Tunnel Interfaces 347
Virtual IP Addresses 356
VIP and the Global Zone 359
Example: Configuring Virtual IP Servers 359
Example: Editing a VIP Configuration 362
Example: Removing a VIP Configuration 362
Example: VIP with Custom and Multiple-Port Services 363
Chapter 9 User Authentication 371
Authentication Servers 372
Local Database 374
Supported User Types and Features 374
Example: Setting the Local Database Timeout 375
External Auth Servers 376
Auth Server Object Properties 377
Auth Server Types 379
RADIUS 379
RADIUS Auth Server Object Properties 380
Supported User Types and Features 380
NetScreen Dictionary File 381
RADIUS Access-Challenge 382
SecurID 384
SecurID Auth Server Object Properties 385
Supported User Types and Features 385
LDAP 386
LDAP Auth Server Object Properties 387
Supported User Types and Features 387

Defining Auth Server Objects 388
Example: Defining an Auth Server Object for RADIUS 388
Example: Defining an Auth Server Object for SecurID 391
Example: Defining an Auth Server Object for LDAP 393
Defining Default Auth Servers 395
Example: Changing the Default Auth Servers 395
Authentication Types and Applications 397
Auth Users and User Groups 398
Referencing Auth Users in Policies 398
Referencing Auth User Groups in Policies 402
Example: Run-Time Authentication (Local User) 403
Example: Run-Time Authentication (Local User Group) 406
Example: Run-Time Authentication (External User) 409
Example: Run-Time Authentication (External User Group) 412
Example: Local Auth User in Multiple Groups 416
Example: WebAuth (Local User Group) 420
Example: WebAuth (External User Group) 423
Example: WebAuth + SSL (External User Group) 427
IKE Users and User Groups 431
Example: Defining IKE Users 432
Example: Creating an IKE User Group 434
Referencing IKE Users in Gateways 435
XAuth Users and User Groups 436
XAuth Users in IKE Negotiations 437
Example: XAuth Authentication (Local User) 440
Example: XAuth Authentication (Local User Group) 442
Example: XAuth Authentication (External User) 444
Example: XAuth Authentication (External User Group) 447
Example: XAuth Authentication and Address Assignments (Local User Group) 452
XAuth Client 458
Example: NetScreen Device as an XAuth Client 459
L2TP Users and User Groups 460
Example: Local and External L2TP Auth Servers 461
Admin Users 465
Multiple-Type Users 467
Group Expressions 468
Example: Group Expressions (AND) 470
Example: Group Expressions (OR) 472
Example: Group Expressions (NOT) 474
Banner Customization 476
Example: Customizing the WebAuth Success Message 476
Chapter 10 Traffic Shaping 477
Applying Traffic Shaping 478
Managing Bandwidth at the Policy Level 478
Example: Traffic Shaping 479
Setting Service Priorities 485
Example: Priority Queuing 486

Chapter 11 System Parameters 493
Domain Name System Support 495
DNS Lookup 496
DNS Status Table 497
Example: Defining DNS Server Addresses and Scheduling Lookups 498
Example: Setting a DNS Refresh Interval 499
DHCP 500
DHCP Server 502
Example: NetScreen Device as DHCP Server 502
DHCP Server in an NSRP Cluster 508
DHCP Server Detection 508
Example: Turning on DHCP Server Detection 509
Example: Turning off DHCP Server Detection 509
DHCP Relay Agent 510
Example: NetScreen Device as DHCP Relay Agent 511
DHCP Client 516
Example: NetScreen Device as DHCP Client 516
TCP/IP Settings Propagation 518
Example: Forwarding TCP/IP Settings 519
PPPoE 521
Example: Setting Up PPPoE 521
Example: Configuring PPPoE on Primary and Backup Untrust Interfaces 526
Downloading/Uploading Settings and Software 528
Saving and Importing Settings 528
Uploading and Downloading Software 530
Configuration Rollback 531
Last-Known-Good Configuration 531
Automatic and Manual Configuration Rollback 531
Loading a New Configuration File 533
Locking the Configuration File 534
Adding Comments to a Configuration File 535
License Keys 536
Example: Expanding User Capacity 537
Registration and Activation of Signature Services 538
Temporary Service 538
AV and DI Bundled with a New Device 538
AV Upgrade with DI 539
DI Upgrade Only 540
System Clock 541
Date and Time 541
Time Zone 541
NTP 542
Multiple NTP Servers 542
Maximum Time Adjustment 542
NTP and NSRP 543
Example: Configuring NTP Servers and a Maximum Time Adjustment Value 544
Secure NTP Servers 545
Index ..... IX-I

Preface

Volume 2, “Fundamentals”describes the ScreenOS architecture and its elements, including examples for configuring various elements. This volume describes the following:

- Security, tunnel, and function zones
- Routing basics, including route tables and how to configure static routes.
- Various interface types, such as physical interfaces, subinterfaces, virtual security interfaces (VSIs), redundant interfaces, aggregate interfaces, and VPN tunnel interfaces
- Source and destination Network Address Translation (NAT-src and NAT-dst), Mapped IP (MIP) addresses, virtual IP (VIP) addresses, dynamic IP (DIP) addresses
- Interface modes in which NetScreen interfaces can operate: Network Address Translation (NAT), Route, and Transparent
- Policies, which are used to control the traffic flow across an interface, and the elements that are used to create policies and virtual private networks, such as addresses, users, and services
- User authentication methods available to NetScreen devices and how to configure user accounts and user groups
- Traffic management concepts
- System parameters for the following functions:
  - Domain Name System (DNS) addressing
  - Dynamic Host Configuration Protocol (DHCP) for assigning or relaying TCP/IP settings
  - URL filtering
  - Uploading and downloading of configuration settings and software to and from a NetScreen device
  - License keys to expand the capabilities of a NetScreen device
  - System clock configuration

CONVENTIONS

This document contains several types of conventions, which are introduced in the following sections:

- "CLI Conventions"
- "WebUI Conventions" on page xi
- "Illustration Conventions" on page xiii
- "Naming Conventions and Character Types" on page xiv

CLI Conventions

The following conventions are used when presenting the syntax of a command line interface (CLI) command:

- Anything inside square brackets [ ] is optional.
- Anything inside braces { } is required.
- If there is more than one choice, each choice is separated by a pipe (|). For example,

  ```
  set interface { ethernet1 | ethernet2 | ethernet3 } manage
  ```

  means "set the management options for the ethernet1, ethernet2, or ethernet3 interface".
- Variables appear in *italic.* For example:

  ```
  set admin user name password
  ```

When a CLI command appears within the context of a sentence, it is in **bold** (except for variables, which are always in *italic*). For example: "Use the **get system** command to display the serial number of a NetScreen device."

Note: *When typing a keyword, you only have to type enough letters to identify the word uniquely. For example, typing* ***set adm u joe j12fmt54*** *is enough to enter the command* ***set admin user joe j12fmt54****. Although you can use this shortcut when entering commands, all the commands documented here are presented in their entirety.*

WebUI Conventions

Throughout this book, a chevron (>) is used to indicate navigation through the WebUI by clicking menu options and links. For example, the path to the address configuration dialog box is presented as **Objects > Addresses > List > New**. This navigational sequence is shown below.

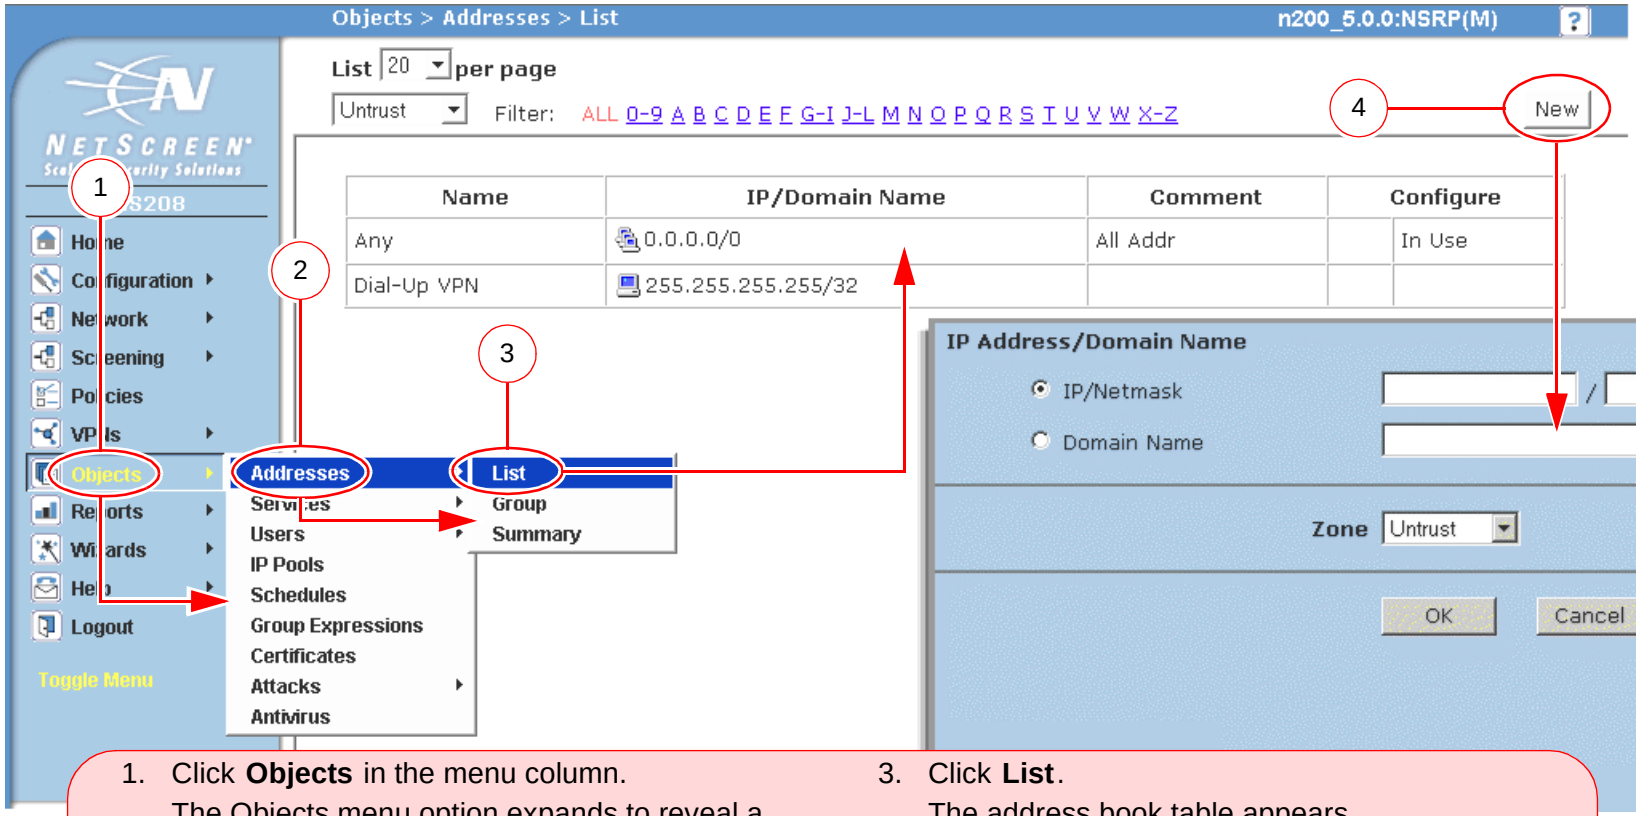

1. Click **Objects** in the menu column.
   The Objects menu option expands to reveal a subset of options for Objects.
2. (Applet menu) Hover the mouse over **Addresses**.
   (DHTML menu) Click **Addresses**.
   The Addresses option expands to reveal a subset of options for Addresses.
3. Click **List**.
   The address book table appears.
4. Click the **New** link.
   The new address configuration dialog box appears.

To perform a task with the WebUI, you must first navigate to the appropriate dialog box where you can then define objects and set parameters. The set of instructions for each task is divided into two parts: a navigational path and configuration details. For example, the following set of instructions includes the path to the address configuration dialog box and the settings for you to configure:

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: addr_1

IP Address/Domain Name:

IP/Netmask: (select), 10.2.2.5/32

Zone: Untrust

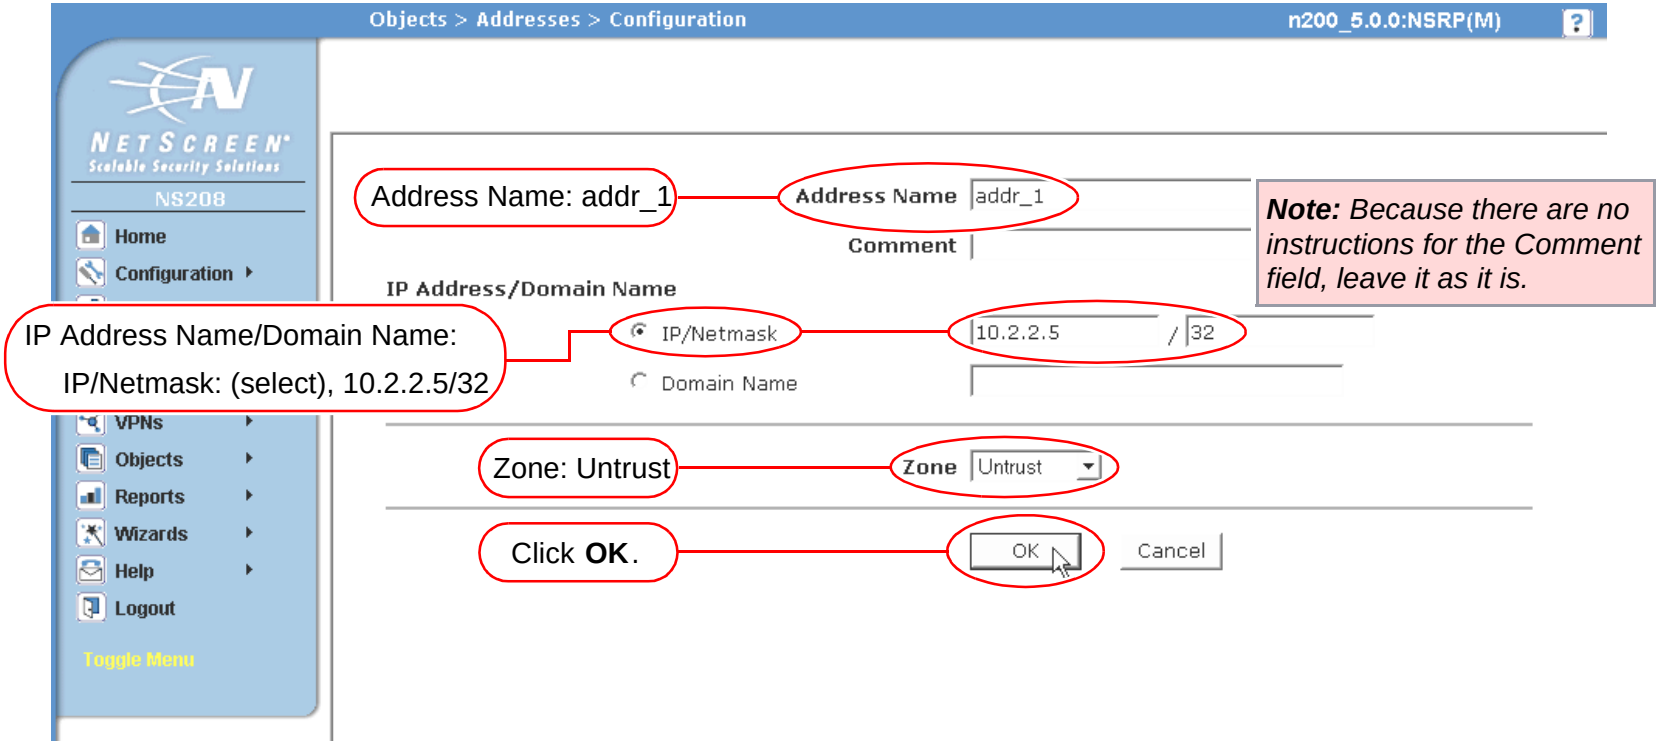

Illustration Conventions

The following graphics make up the basic set of images used in illustrations throughout this book:

Generic NetScreen Device

Virtual Routing Domain

Security Zone

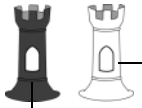

Security Zone Interfaces

White = Protected Zone Interface (example: Trust Zone)

Black = Outside Zone Interface (example: Untrust Zone)

Tunnel Interface

VPN Tunnel

Router Icon

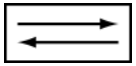

Switch Icon

Local Area Network (LAN) with a Single Subnet

(example: 10.1.1.0/24)

Internet

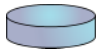

Dynamic IP (DIP) Pool

Desktop Computer

Laptop Computer

Generic Network Device

(examples: NAT server, Access Concentrator)

Server

Naming Conventions and Character Types

ScreenOS employs the following conventions regarding the names of objects—such as addresses, admin users, auth servers, IKE gateways, virtual systems, VPN tunnels, and zones—defined in ScreenOS configurations.

- If a name string includes one or more spaces, the entire string must be enclosed within double quotes ("); for example, **set address trust "local LAN" 10.1.1.0/24**.
- NetScreen trims any spaces leading or trailing text within a set of double quotes; for example, **" local LAN "** becomes **"local LAN"**.
- NetScreen treats multiple consecutive spaces as a single space.
- Name strings are case sensitive, although many CLI key words are case insensitive. For example, **"local LAN"** is different from **"local lan"**.

ScreenOS supports the following character types:

- Single-byte character sets (SBCS) and multiple-byte character sets (MBCS). Examples of SBCS are ASCII, European, and Hebrew. Examples of MBCS—also referred to as double-byte character sets (DBCS)—are Chinese, Korean, and Japanese.

  Note: *A console connection only supports SBCS. The WebUI supports both SBCS and MBCS, depending on the character sets that your Web browser supports.*

- ASCII characters from 32 (0x20 in hexidecimals) to 255 (0xff), except double quotes ("), which have special significance as an indicator of the beginning or end of a name string that includes spaces.

NETSCREEN DOCUMENTATION

To obtain technical documentation for any NetScreen product, visit www.netscreen.com/resources/manuals/.

To obtain the latest software version, visit www.netscreen.com/services/download_soft. Select a category of software product from the dropdown list, then follow the displayed instructions. (You must be a registered user to download NetScreen software.)

If you find any errors or omissions in the following content, please contact us at the e-mail address below:

techpubs@netscreen.com

1

ScreenOS Architecture

The NetScreen ScreenOS architecture offers great flexibility in designing the layout of your network security. On NetScreen devices with more than two interfaces, you can create numerous security zones and configure policies to regulate traffic between and within zones. You can bind one or more interfaces to each zone and enable a unique set of management and firewall attack screening options on a per-zone basis. Essentially, ScreenOS allows you to create the number of zones your network environment requires, assign the number of interfaces each zone requires, and design each interface to your specifications.

This chapter presents an overview of ScreenOS, covering the following key components:

- "Security Zones" on page 2
- "Security Zone Interfaces" on page 3
- "Virtual Routers" on page 5
- "Policies" on page 6
- "VPNs" on page 8
- "Virtual Systems" on page 10

Furthermore, to better understand the ScreenOS mechanism for processing traffic, you can see the flow sequence for an incoming packet in "Packet Flow Sequence" on page 11.

The chapter concludes with a four-part example that illustrates a basic configuration for a NetScreen device using ScreenOS:

- "Example (Part 1): Enterprise with Six Zones" on page 14
- "Example (Part 2): Interfaces for Six Zones" on page 16
- "Example (Part 3): Enterprise with Two Routing Domains" on page 20
- "Example (Part 4): Policies for an Enterprise with Six Zones" on page 22

SECURITY ZONES

A security zone is a collection of one or more network segments requiring the regulation of inbound and outbound traffic via policies (see "Policies" on page 6)[1]. Security zones are logical entities to which one or more interfaces are bound. With many types of NetScreen devices, you can define multiple security zones, the exact number of which you determine based on your network needs. In addition to user-defined zones, you can also use the predefined zones: Trust, Untrust, and DMZ (for Layer 3 operation), or V1-Trust, V1-Untrust, and V1-DMZ (for Layer 2 operation)[2]. If you want, you can continue using just the predefined zones. You can also ignore the predefined zones and use user-defined zones exclusively[3]. Optionally, you can use both kinds of zones—predefined and user-defined—side by side. This flexibility for zone configuration allows you to create a network design that best suits your specific needs.

A network configured with 5 security zones—3 default zones (Trust, Untrust, DMZ), and 2 user-defined zones (Finance, Eng)

Traffic (indicated by black lines) passes from one security zone to another only if a policy permits it.

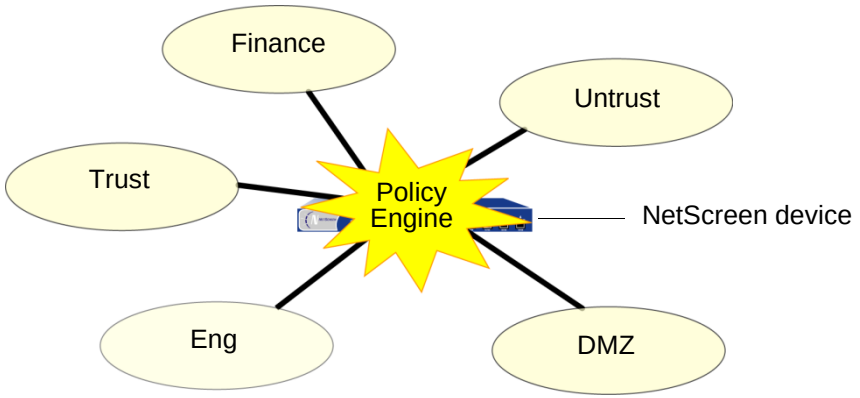

1. The one security zone that requires no network segment is the global zone. (For more information, see Global zone "Global Zone" on page 48.) Additionally, any zone without an interface bound to it nor any address book entries can also be said not to contain any network segments.
2. If you upgrade from an earlier version of ScreenOS, all your configurations for these zones remain intact.
3. You cannot delete a predefined security zone. You can, however, delete a user-defined zone. When you delete a security zone, you also automatically delete all addresses configured for that zone.

SECURITY ZONE INTERFACES

An interface for a security zone can be thought of as a doorway through which TCP/IP traffic can pass between that zone and any other zone.

Through the policies you define, you can permit traffic between zones to flow in one direction or in both[4]. With the routes that you define, you specify the interfaces that traffic from one zone to another must use. Because you can bind multiple interfaces to a zone, the routes you chart are important for directing traffic to the interfaces of your choice.

To permit traffic to flow from zone to zone, you bind an interface to the zone and—for an interface in Route or NAT mode (see Chapter 5, "Interface Modes")—assign an IP address to the interface. Two common interface types are physical interfaces and—for those devices with virtual system support—subinterfaces (that is, a layer 2 substantiation of a physical interface). For more information, see Chapter 4, "Interfaces".

Physical Interfaces

A physical interface relates to components that are physically present on the NetScreen device. The interface naming convention differs from device to device. On the NetScreen-500, for example, a physical interface is identified by the position of an interface module and an ethernet port on that module. For example, the interface *ethernet1/2* designates the interface module in the **first bay** *(ethernet**1**/2)* and the **second port** *(ethernet1/**2**)*.

Physical Interface Assignments

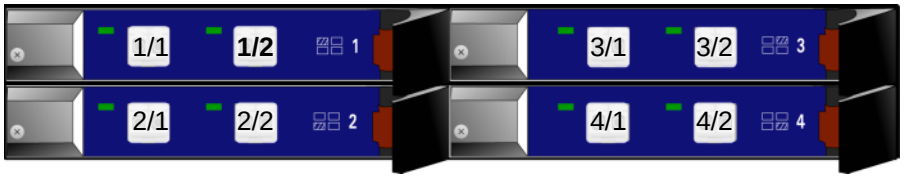

Note: *To see the naming convention for a specific NetScreen device, refer to the User's Guide for that device.*

4. For traffic to flow between interfaces bound to the same zone, no policy is required because both interfaces have security equivalency. ScreenOS requires policies for traffic between zones, not within a zone.

Subinterfaces

On devices that support virtual LANs (VLANs), you can logically divide a physical interface into several virtual subinterfaces, each of which borrows the bandwidth it needs from the physical interface from which it stems. A subinterface is an abstraction that functions identically to a physical interface and is distinguished by 802.1Q VLAN tagging[5]. The NetScreen device directs traffic to and from a zone with a subinterface via its IP address and VLAN tag. For convenience, administrators usually use the same number for a VLAN tag as the subinterface number. For example, the interface ethernet1/2 using VLAN tag 3 is named *ethernet1/2.3*. This refers to the interface module in the first bay, the second port on that module, and subinterface number **3** *(ethernet1/2.**3**)*.

Note that although a subinterface shares part of its identity with a physical interface, the zone to which you bind it is not dependent on the zone to which you bind the physical interface. You can bind the subinterface *ethernet1/2.3* to a different zone than that to which you bind the physical interface *ethernet1/2*, or to which you bind *ethernet1/2.2*. Similarly, there are no restrictions in terms of IP address assignments. The term *subinterface* does not imply that its address be in a subnet of the address space of the physical interface.

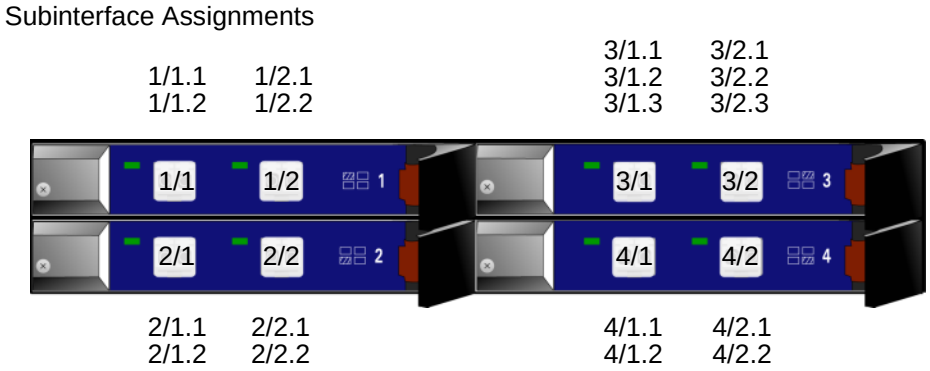

5. 802.1Q is an IEEE standard that defines the mechanisms for the implementation of virtual bridged LANs and the ethernet frame formats used to indicate VLAN membership via VLAN tagging.

VIRTUAL ROUTERS

A virtual router (VR) functions as a router. It has its own interfaces and its own routing table. In ScreenOS, a NetScreen device supports two predefined virtual routers. This allows the NetScreen device to maintain two separate routing tables and to conceal the routing information in one virtual router from the other. For example, the untrust-vr is typically used for communication with untrusted parties and does not contain any routing information for the protected zones. Routing information for the protected zones is maintained by the trust-vr. Thus, no internal network information can be gleaned by the surreptitious extraction of routes from the untrust-vr.

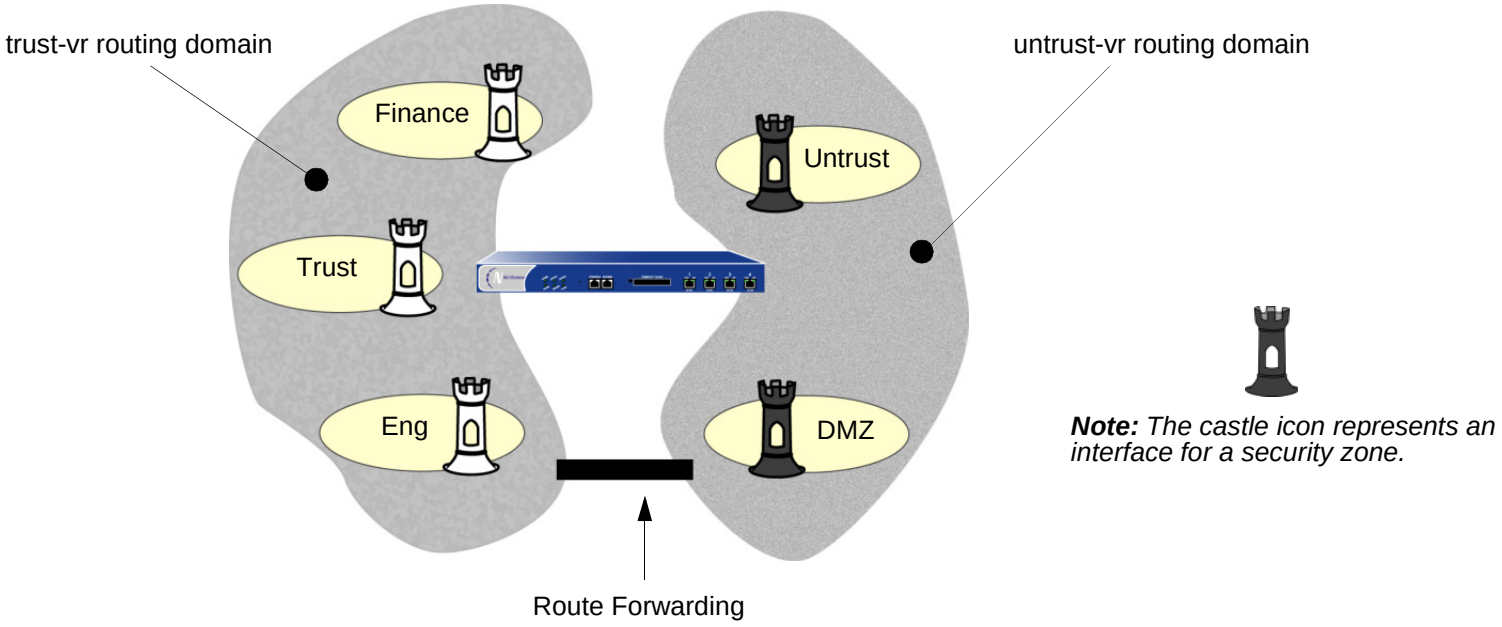

When there are two virtual routers on a NetScreen device, traffic is *not* automatically forwarded between zones that reside in different VRs, even if there are policies that permit the traffic. If you want traffic to pass between virtual routers, you need to either export routes between the VRs or configure a static route in one VR that defines the other VR as the next-hop. For more information about using two virtual routers, see Volume 6 “Dynamic Routing”.

POLICIES

NetScreen devices secure a network by inspecting, and then allowing or denying, all connection attempts that require passage from one security zone to another.

By default, a NetScreen device denies all traffic in all directions[6]. Through the creation of policies, you can then control the traffic flow from zone to zone by defining the kinds of traffic permitted to pass from specified sources to specified destinations at scheduled times. At the broadest level, you can allow all kinds of traffic from any source in one zone to any destination in all other zones without any scheduling restrictions. At the narrowest level, you can create a policy that allows only one kind of traffic between a specified host in one zone and another specified host in another zone during a scheduled period of time.

Broadly defined Internet Access: Any service from any point in the Trust zone to any point in the Untrust zone at any time

Narrowly defined Internet Access: SMTP service from a mail server in the Trust zone to a mail server in the Untrust zone from 5:00 AM to 7:00 PM

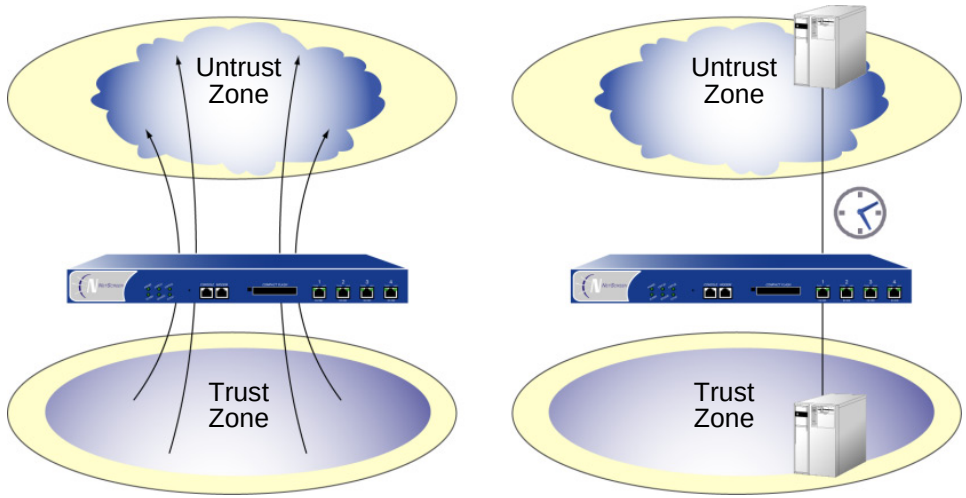

6. Some NetScreen devices ship with a default policy that allows all outbound traffic from the Trust to the Untrust zone but denies all inbound traffic from the Untrust zone to the Trust zone.

Every time a packet attempts to pass from one zone to another or between two interfaces bound to the same zone, the NetScreen device checks its policy set lists for a policy that permits such traffic (see "Policy Set Lists" on page 202). To allow traffic to pass from one security zone to another—for example, from zone A to zone B—you must configure a policy that permits zone A to send traffic to zone B. To allow traffic to flow the other way, you must configure another policy permitting traffic from zone B to zone A. For any traffic to pass from one zone to another, there must be a policy that permits it. Also, if intrazone blocking is enabled, there must be a policy to permit traffic to pass from one interface to another within that zone.

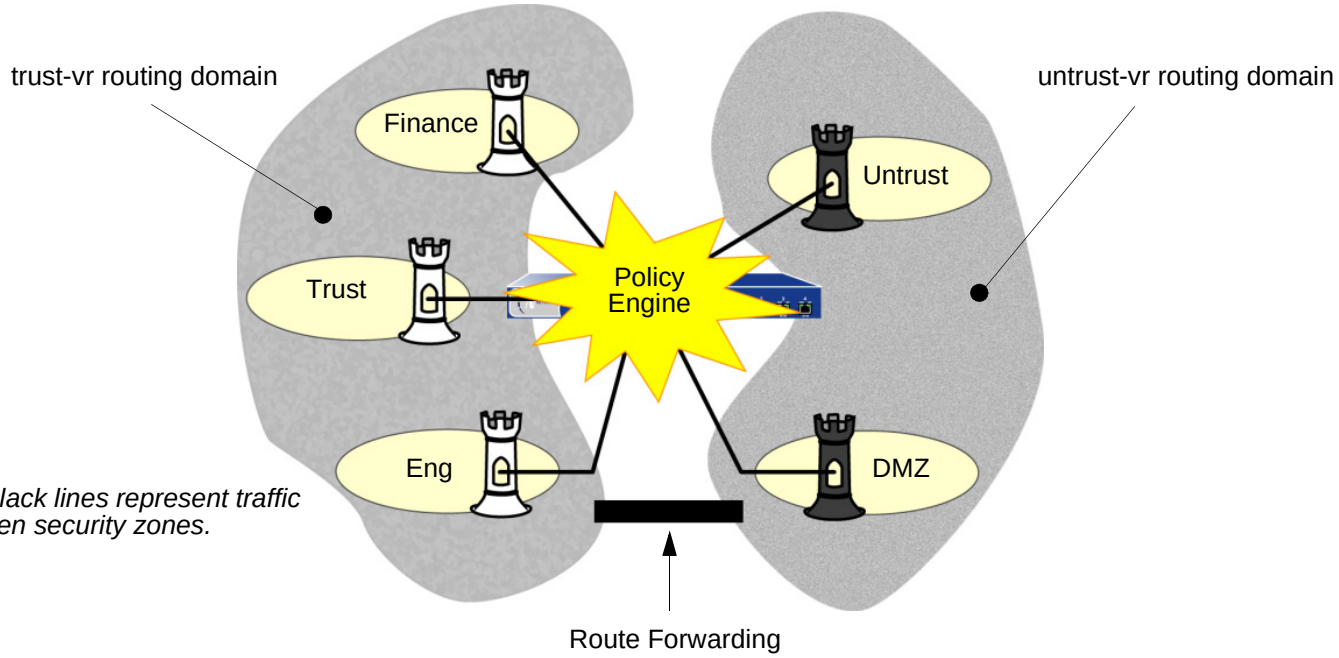

Note: *The black lines represent traffic between security zones.*

Note: *For information about policies, see Chapter 7, "Policies".*

VPNs

ScreenOS supports several virtual private network (VPN) configuration options. The two main types are as follows:

- **Route-based VPN** – A route lookup determines which traffic the NetScreen device encapsulates. Policies either permit or deny traffic to the destination specified in the route. If the policy permits the traffic and the route references a tunnel interface bound to a VPN tunnel, then the NetScreen device also encapsulates it. This configuration separates the application of policies from the application of VPN tunnels. Once configured, such tunnels exist as available resources for securing traffic en route between one security zone and another.
- **Policy-based VPN** – A policy lookup determines which traffic the NetScreen device encapsulates when the policy references a particular VPN tunnel and specifies "tunnel" as the action.

A route-based VPN is good choice for site-to-site VPN configurations because you can be apply multiple policies to traffic passing through a single VPN tunnel. A policy-based VPN is a good choice for dialup VPN configurations because the dialup client might not have an internal IP address to which you can set a route.

The following steps provide a sense of the main elements involved in a route-based VPN configuration:

1. While configuring the VPN tunnel (for example, *vpn-to-SF*, where *SF* is the destination or end entity), specify a physical interface or subinterface on the local device as the outgoing interface. (The IP address for this interface is what the remote peer must use when configuring its remote gateway.)
2. Create a tunnel interface (for example, *tunnel.1*), and bind it to a security zone[7].
3. Bind the tunnel interface *tunnel.1* to the VPN tunnel *vpn-to-SF*.
4. To direct traffic through this tunnel, set up a route stating that traffic to *SF* must use *tunnel.1*.

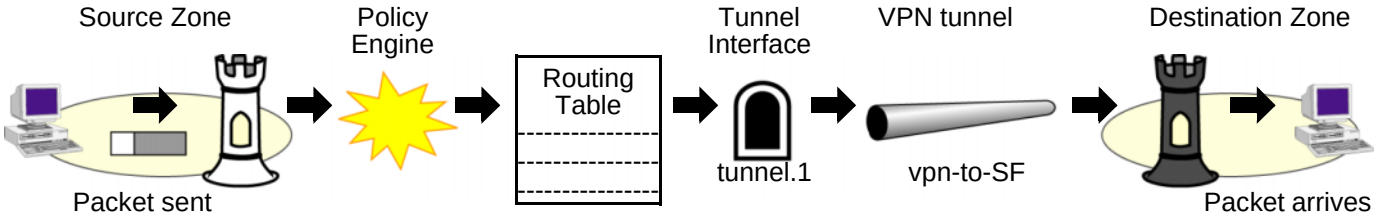

7. You do not have to bind the tunnel interface to the same zone for which VPN traffic is destined. Traffic to any zone can access a tunnel interface if a route points to that interface.

At this point, the tunnel is ready for traffic bound for *SF*. You can now create address book entries, such as “Trust LAN” (10.1.1.0/24) and “SF LAN” (10.2.2.0/24) and set up policies to permit or block different types of traffic from a specified source, such as “Trust LAN”, to a specified destination, such as “SF LAN”.

The local NetScreen device routes traffic from the Trust zone to “SF LAN” in the Untrust zone through the tunnel.1 interface. Because tunnel.1 is bound to the VPN tunnel “vpn-to-SF”, the NetScreen device encrypts the traffic and sends it through that tunnel to the remote peer.

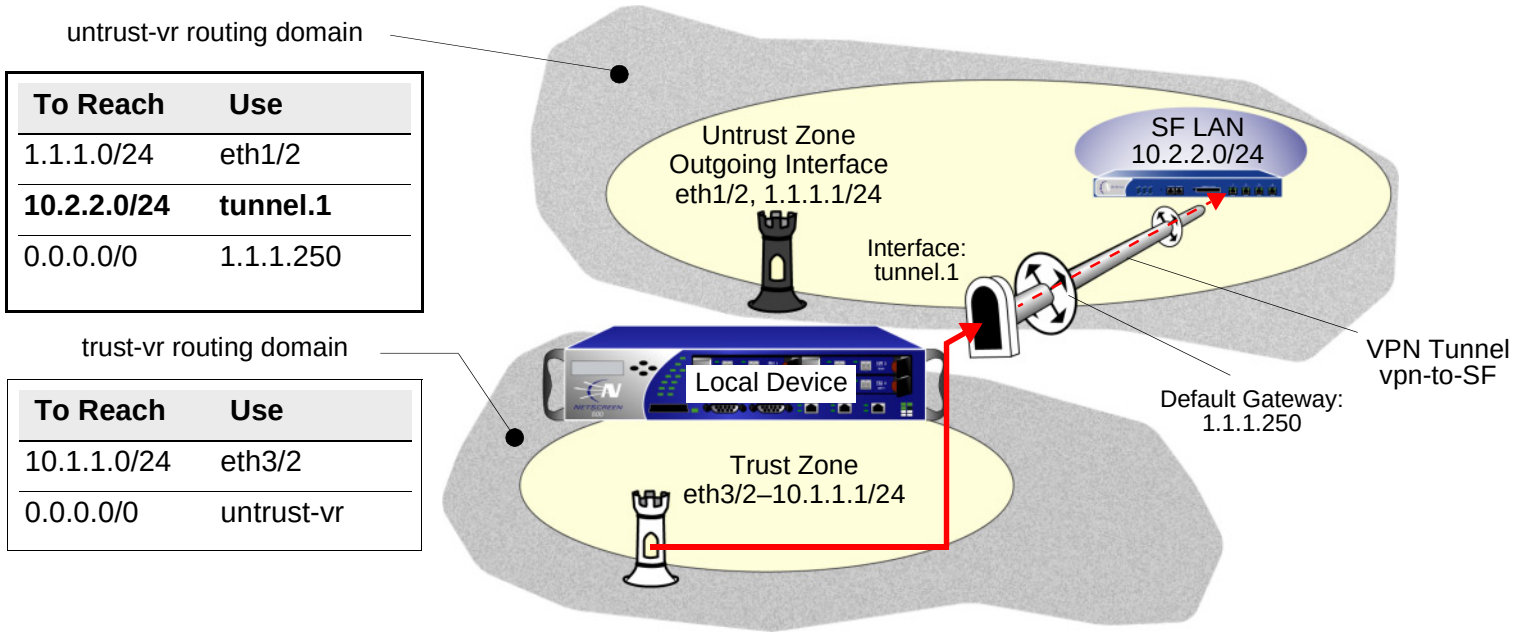

Note: *For detailed information about VPNs, see Volume 5, “VPNs”.*

VIRTUAL SYSTEMS

Some NetScreen devices support virtual systems (vsys). A virtual system is a subdivision of the main system that appears to the user to be a stand-alone entity. Virtual systems reside separately from each other and from the root system within the same NetScreen device. The application of ScreenOS to virtual systems involves the coordination of three main components: zones, interfaces, and virtual routers. The following illustration presents a conceptual overview of how ScreenOS integrates these components at both the root and vsys levels.

Note: *The castle icon represents a security zone interface.*

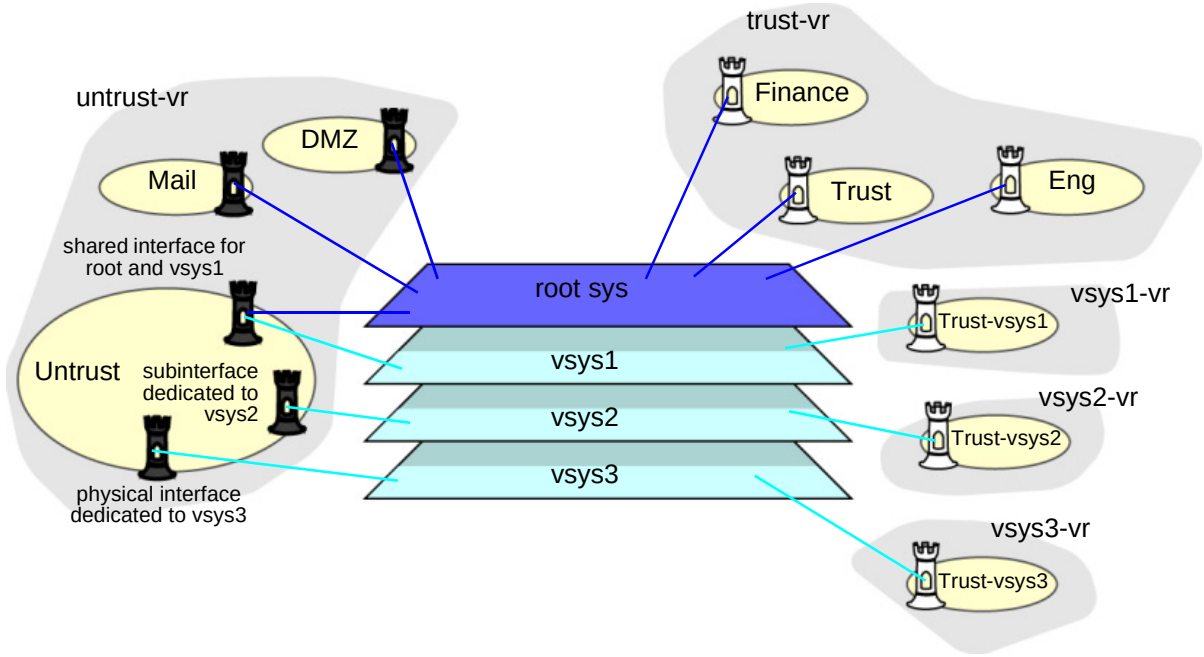

Note: *For further information on virtual systems and the application of zones, interfaces, and virtual routers within the context of virtual systems, see Volume 7, "Virtual Systems".*

PACKET FLOW SEQUENCE

In ScreenOS, the flow sequence of an incoming packet progresses as presented below.

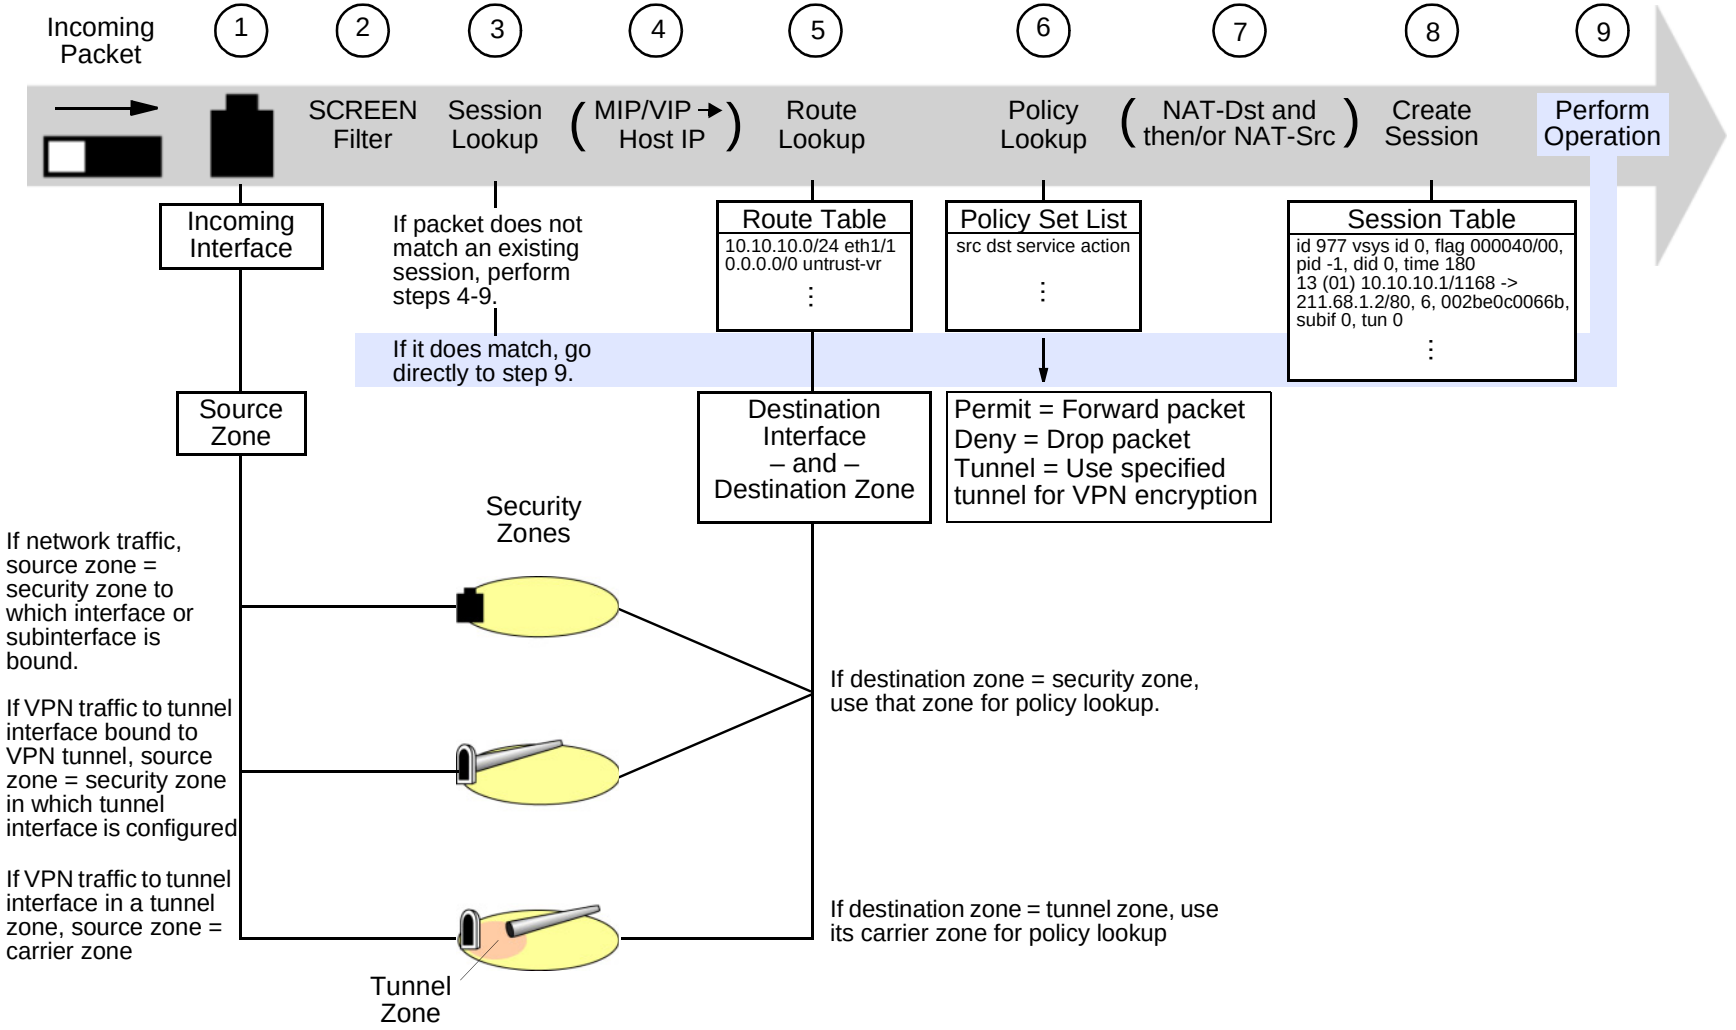

1. The interface module identifies the incoming interface and, consequently, the source zone to which the interface is bound.

   The source zone determination is based on the following criteria:

   – If the packet is not encapsulated, the source zone is the security zone to which the incoming interface or subinterface is bound.
   – If the packet is encapsulated and the tunnel interface is bound to a VPN tunnel, the source zone is the security zone in which the tunnel interface is configured.
   – If the packet is encapsulated and the tunnel interface is in a tunnel zone, the source zone is the corresponding carrier zone (a security zone that *carries* a tunnel zone) for that tunnel zone.

2. If you have enabled SCREEN options for the source zone, the NetScreen device activates the SCREEN module at this point. SCREEN checking can produce one of the following three results:

   – If a SCREEN mechanism detects anomalous behavior for which it is configured to block the packet, the NetScreen device drops the packet and makes an entry in the event log.
   – If a SCREEN mechanism detects anomalous behavior for which it is configured to record the event but not block the packet, the NetScreen device records the event in the SCREEN counters list for the ingress interface and proceeds to the next step.
   – If the SCREEN mechanisms detect no anomalous behavior, the NetScreen device proceeds to the next step.

3. The session module performs a session lookup, attempting to match the packet with an existing session.

   If the packet does not match an existing session, the NetScreen device performs First Packet Processing, a procedure involving the following steps 4 through 9.

   If the packet matches an existing session, the NetScreen device performs Fast Processing, using the information available from the existing session entry to process the packet. Fast Processing bypasses steps 4 through 8 because the information generated by those steps has already been obtained during the processing of the first packet in the session.

4. If a mapped IP (MIP) or virtual IP (VIP) address is used, the address-mapping module resolves the MIP or VIP so that the routing table can search for the actual host address.

5. The route table lookup finds the interface that leads to the destination address. In so doing, the interface module identifies the destination zone to which that interface is bound.

   The destination zone determination is based on the following criteria:

   – If the destination zone is a security zone, that zone is used for the policy lookup.
   – If the destination zone is a tunnel zone, the corresponding carrier zone is used for the policy lookup.

6. The policy engine searches the policy set lists for a policy between the addresses in the identified source and destination zones.

   The action configured in the policy determines what the NetScreen firewall does with the packet:

   – If the action is **permit**, the NetScreen device determines to forward the packet to its destination.
   – If the action is **deny**, the NetScreen device determines to drop the packet.
   – If the action is **tunnel**, the NetScreen device determines to forward the packet to the VPN module, which encapsulates the packet and transmits it using the specified VPN tunnel settings.

7. If destination address translation (NAT-dst) is specified in the policy, the NAT module translates the original destination address in the IP packet header to a different address.

   If source address translation is specified (either interface-based NAT or policy-based NAT-src), the NAT module translates the source address in the IP packet header before forwarding it either to its destination or to the VPN module.

   (If both NAT-dst and NAT-src are specified in the same policy, the NetScreen device first performs NAT-dst and then NAT-src.)

8. The session module creates a new entry in the session table containing the results of steps 1 through 7.

   The NetScreen device then uses the information maintained in the session entry when processing subsequent packets of the same session.

9. The NetScreen device performs the operation specified in the session.

   Some typical operations are source address translation, VPN tunnel selection and encryption, decryption, and packet forwarding.

Example (Part 1): Enterprise with Six Zones

This is the first of a four-part example, the purpose of which is to illustrate some of the concepts covered in the previous sections. For this second part, in which the interfaces for each zone are set, see "Example (Part 2): Interfaces for Six Zones" on page 16. Here you configure the following six zones for an enterprise:

- Finance
- Trust
- Eng
- Mail
- Untrust
- DMZ

The Trust, Untrust, and DMZ zones are preconfigured. You must define the Finance, Eng, and Mail zones. By default, a user-defined zone is placed in the trust-vr routing domain. Thus, you do not have to specify a virtual router for the Finance and Eng zones. However, in addition to configuring the Mail zone, you must also specify that it be in the untrust-vr routing domain. You must also shift virtual router bindings for the Untrust and DMZ zones from the trust-vr to the untrust-vr[8].

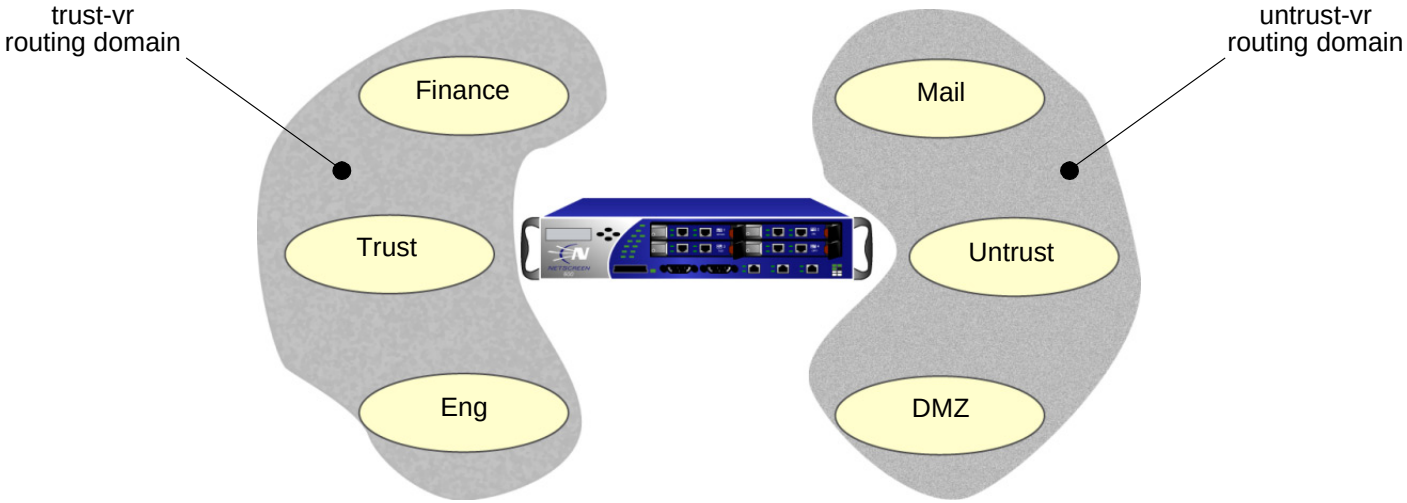

8. For more information on virtual routers and their routing domains, see Chapter 2, "Routing Tables and Static Routing".

WebUI

Network > Zones > New: Enter the following, and then click **OK**:

Zone Name: Finance

Virtual Router Name: trust-vr

Zone Type: Layer 3: (select)

Network > Zones > New: Enter the following, and then click **OK**:

Zone Name: Eng

Virtual Router Name: trust-vr

Zone Type: Layer 3: (select)

Network > Zones > New: Enter the following, and then click **OK**:

Zone Name: Mail

Virtual Router Name: untrust-vr

Zone Type: Layer 3: (select)

Network > Zones > Edit (for Untrust): Select **untrust-vr** in the Virtual Router Name drop-down list, and then click **OK**.

Network > Zones > Edit (for DMZ): Select **untrust-vr** in the Virtual Router Name drop-down list, and then click **OK**.

CLI

```
set zone name finance
set zone name eng
set zone name mail
set zone mail vrouter untrust-vr
set zone untrust vrouter untrust-vr
set zone dmz vrouter untrust-vr
save
```

Example (Part 2): Interfaces for Six Zones

This is the second part of an ongoing example. For the first part, in which zones are configured, see “Example (Part 1): Enterprise with Six Zones” on page 14. For the next part, in which virtual routers are configured, see “Example (Part 3): Enterprise with Two Routing Domains” on page 20. This part of the example demonstrates how to bind interfaces to zones and configure them with an IP address and various management options.

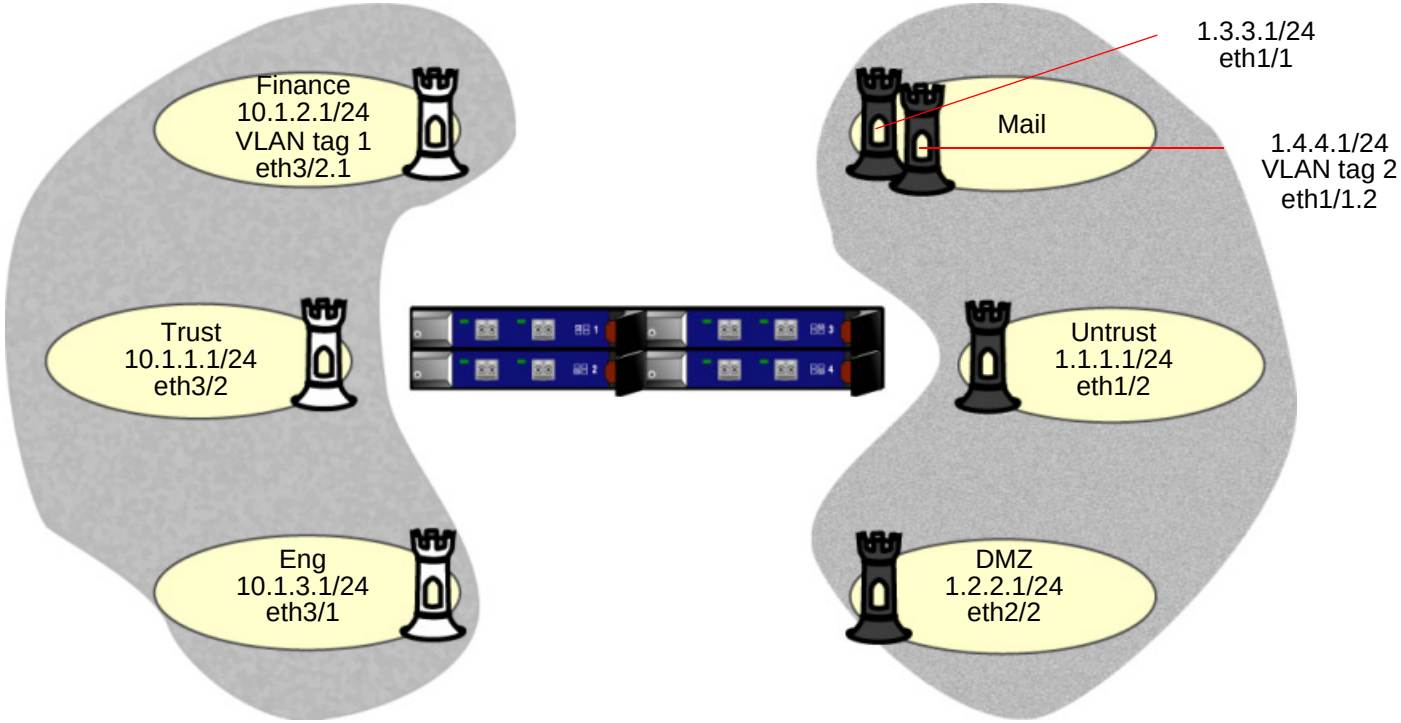

WebUI

1. **Interface ethernet3/2**

   Network > Interfaces > Edit (for ethernet3/2): Enter the following, and then click **OK**:

   Zone Name: Trust

   Static IP: (select this option when present)

   IP Address/Netmask: 10.1.1.1/24

   Manageable: (select)

   Management Services: WebUI, Telnet, SNMP, SSH (select)

   Other Services: Ping (select)

2. **Interface ethernet3/2.1**

   Network > Interfaces > Sub-IF New: Enter the following, and then click **OK**:

   Interface Name: ethernet3/2.1

   Zone Name: Finance

   Static IP: (select this option when present)

   IP Address/Netmask: 10.1.2.1/24

   VLAN Tag: 1

   Other Services: Ping (select)

3. **Interface ethernet3/1**

   Network > Interfaces > Edit (for ethernet3/1): Enter the following, and then click **OK**:

   Zone Name: Eng

   Static IP: (select this option when present)

   IP Address/Netmask: 10.1.3.1/24

   Other Services: Ping (select)

4. Interface ethernet1/1

Network > Interfaces > Edit (for ethernet1/1): Enter the following, and then click **OK**:

Zone Name: Mail

Static IP: (select this option when present)

IP Address/Netmask: 1.3.3.1/24

5. Interface ethernet1/1.2

Network > Interfaces > Sub-IF New: Enter the following, and then click **OK**:

Interface Name: ethernet1/1.2

Zone Name: Mail

Static IP: (select this option when present)

IP Address/Netmask: 1.4.4.1/24

VLAN Tag: 2

6. Interface ethernet1/2

Network > Interfaces > Edit (for ethernet1/2): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

Manageable: (select)

Management Services: SNMP (select)

7. Interface ethernet2/2

Network > Interfaces > Edit (for ethernet2/2): Enter the following, and then click **OK**:

Zone Name: DMZ

Static IP: (select)

IP Address/Netmask: 1.2.2.1/24

CLI

1. Interface ethernet3/2

```
set interface ethernet3/2 zone trust
set interface ethernet3/2 ip 10.1.1.1/24
set interface ethernet3/2 manage ping
set interface ethernet3/2 manage webui
set interface ethernet3/2 manage telnet
set interface ethernet3/2 manage snmp
set interface ethernet3/2 manage ssh
```

2. Interface ethernet3/2.1

```
set interface ethernet3/2.1 tag 1 zone finance
set interface ethernet3/2.1 ip 10.1.2.1/24
set interface ethernet3/2.1 manage ping
```

3. Interface ethernet3/1

```
set interface ethernet3/1 zone eng
set interface ethernet3/1 ip 10.1.3.1/24
set interface ethernet3/1 manage ping
```

4. Interface ethernet1/1

```
set interface ethernet1/1 zone mail
set interface ethernet1/1 ip 1.3.3.1/24
```

5. Interface ethernet1/1.2

```
set interface ethernet1/1.2 tag 2 zone mail
set interface ethernet1/1.2 ip 1.4.4.1 /24
```

6. Interface ethernet1/2

```
set interface ethernet1/2 zone untrust
set interface ethernet1/2 ip 1.1.1.1/24
set interface ethernet1/2 manage snmp
```

7. Interface ethernet2/2

```
set interface ethernet2/2 zone dmz
set interface ethernet2/2 ip 1.2.2.1/24
save
```

Example (Part 3): Enterprise with Two Routing Domains

This is the third part of an ongoing example. For the previous part, in which interfaces for the various security zones are defined, see "Example (Part 2): Interfaces for Six Zones" on page 16. For the next part, in which the polices are set, see "Example (Part 4): Policies for an Enterprise with Six Zones" on page 22. In this example, you only have to configure a route for the default gateway to the Internet. The other routes are automatically created by the NetScreen device when you create the interface IP addresses.

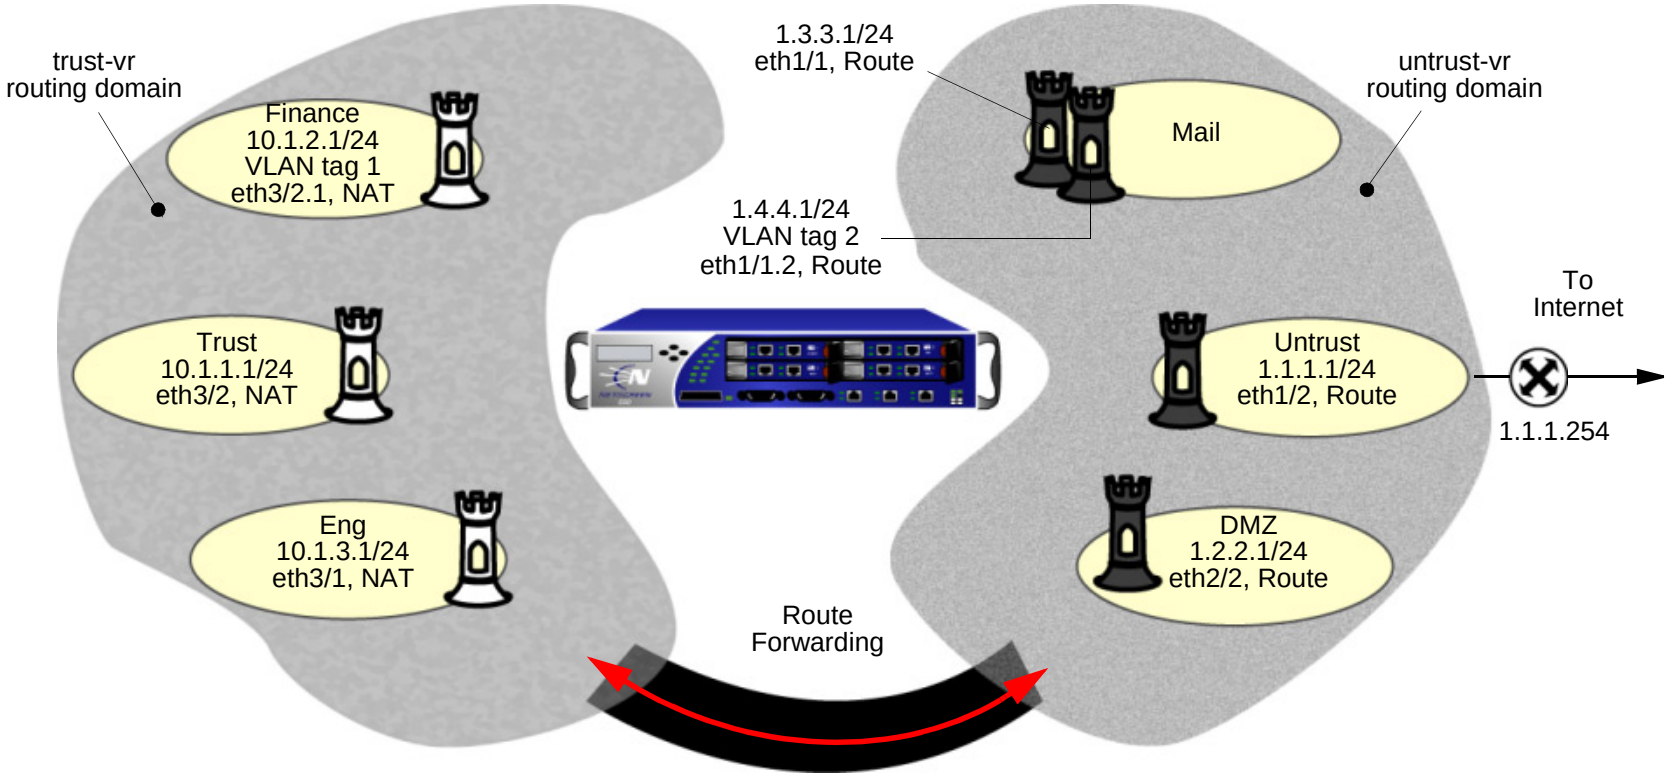

WebUI

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Next Hop Virtual Router Name: (select); untrust-vr

Network > Routing > Routing Entries > untrust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet1/2

Gateway IP Address: 1.1.1.254

CLI

```
set vrouter trust-vr route 0.0.0.0/0 vrouter untrust-vr
set vrouter untrust-vr route 0.0.0.0/0 interface eth1/2 gateway 1.1.1.254
save
```

The NetScreen device automatically creates the following routes (in black):

trust-vr

To Reach:	Use Interface:	Use Gateway/Vrouter:
0.0.0.0/0	n/a	untrust-vr
10.1.3.0/24	eth3/1	0.0.0.0
10.1.1.0/24	eth3/2	0.0.0.0
10.1.2.0/24	eth3/2.1	0.0.0.0

untrust-vr

To Reach:	Use Interface:	Use Gateway/Vrouter:
1.2.2.0/24	eth2/2	0.0.0.0
1.1.1.0/24	eth1/2	0.0.0.0
1.4.4.0/24	eth1/1.2	0.0.0.0
1.3.3.0/24	eth1/1	0.0.0.0
0.0.0.0/0	eth1/2	1.1.1.254

Note: These are the only user-configured entries.

Example (Part 4): Policies for an Enterprise with Six Zones

This is the last part of an ongoing example. The previous part is "Example (Part 3): Enterprise with Two Routing Domains" on page 20. This part of the example demonstrates how to configure new policies.

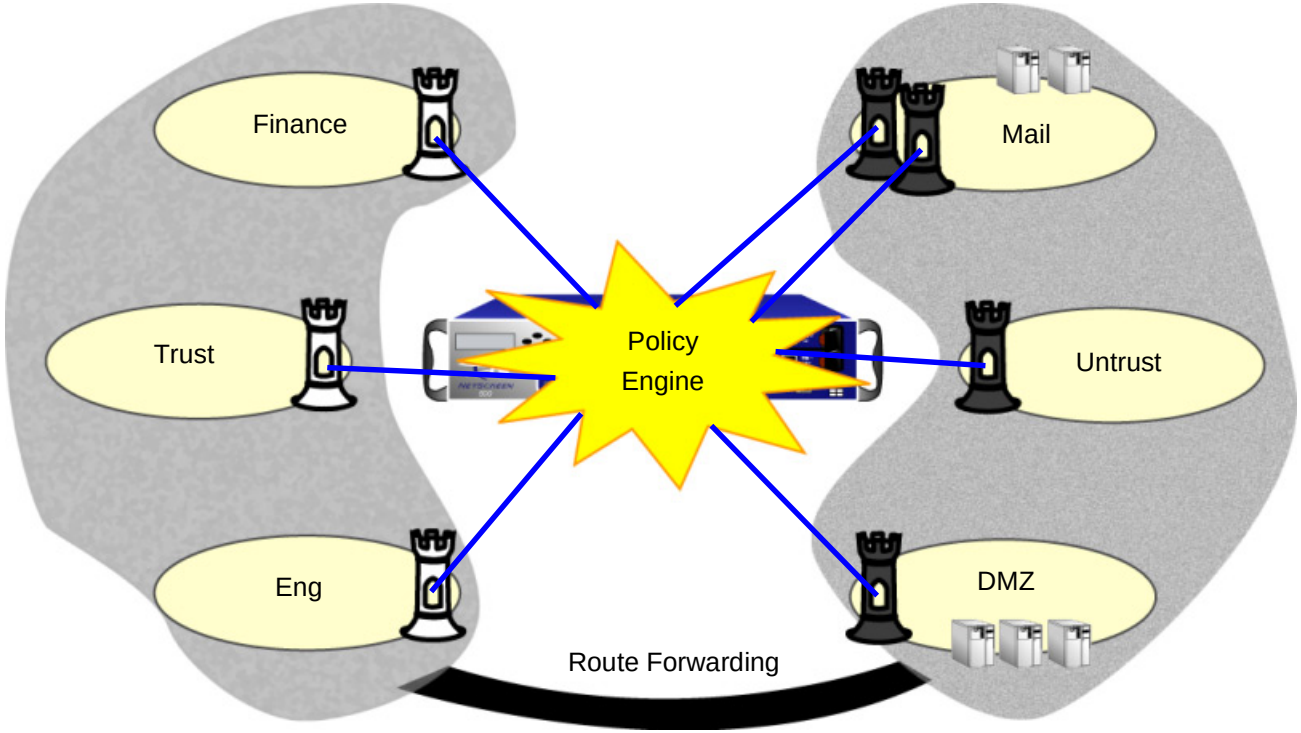

For the purpose of this example, before you begin configuring new policies, you need to create new service groups.

Note: *When you create a zone, the NetScreen device automatically creates the address* ***Any*** *for all hosts within that zone. This example makes use of the address* ***Any*** *for the hosts.*

WebUI

1. Service Groups

Objects > Services > Group > New: Enter the following, and then click **OK**:

Group Name: Mail-Pop3

Select **Mail** and use the **<<** button to move that service from the Available Members column to the Group Members column.

Select **Pop3** and use the **<<** button to move that service from the Available Members column to the Group Members column.

Object > Services > Groups > New: Enter the following, and then click **OK**:

Group Name: HTTP-FTPGet

Select **HTTP** and use the **<<** button to move that service from the Available Members column to the Group Members column.

Select **FTP-Get** and use the **<<** button to move that service from the Available Members column to the Group Members column.

2. Policies

Policies > (From: Finance, To: Mail) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Any

Service: Mail-Pop3

Action: Permit

Policies > (From: Trust, To: Mail) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Any

Service: Mail-Pop3

Action: Permit

Policies > (From: Eng, To: Mail) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Any

Service: Mail-Pop3

Action: Permit

Policies > (From: Untrust, To: Mail) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Any

Service: Mail

Action: Permit

Policies > (From: Finance, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Any

Service: HTTP-FTPGet

Action: Permit

Policies > (From: Finance, To: DMZ) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Any

Service: HTTP-FTPGet

Action: Permit

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Any

Service: HTTP-FTPGet

Action: Permit

Policies > (From: Trust, To: DMZ) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Any

Service: HTTP-FTPGet

Action: Permit

Policies > (From: Eng, To: DMZ) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Any

Service: HTTP-FTPGet

Action: Permit

Policies > (From: Eng, To: DMZ) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Any

Service: FTP-Put

Action: Permit

Policies > (From: Untrust, To: DMZ) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Any

Service: HTTP-FTPGet

Action: Permit

CLI

1. Service Groups

```
set group service mail-pop3 add mail
set group service mail-pop3 add pop3
set group service http-ftpget add http
set group service http-ftpget add ftp-get
```

2. Policies

```
set policy from finance to mail any any mail-pop3 permit
set policy from trust to mail any any mail-pop3 permit
set policy from eng to mail any any mail-pop3 permit
set policy from untrust to mail any any mail permit
set policy from finance to untrust any any http-ftpget permit
set policy from finance to dmz any any http-ftpget permit
set policy from trust to untrust any any http-ftpget permit
set policy from trust to dmz any any http-ftpget permit
set policy from eng to untrust any any http-ftpget permit
set policy from eng to dmz any any http-ftpget permit
set policy from eng to dmz any any ftp-put permit
set policy from untrust to dmz any any http-ftpget permit
save
```

2

Routing Tables and Static Routing

In order for a NetScreen device to forward packets from one network to another, ScreenOS maintains a *routing table* that contains entries for all the network addresses it knows about. The routing table usually contains one or more *static routes*, which are manually entered configurations that define a path to a specific destination.

This chapter describes the ScreenOS routing table, the basic routing process on the NetScreen device, and how to configure static routes on NetScreen devices. It contains the following sections:

- "Routing Essentials" on page 30
  - "Routing Methods" on page 30
  - "Routing Tables" on page 31
  - "Routing with Static Routes" on page 33
- "Virtual Routers on NetScreen Devices" on page 35
- "When to Configure Static Routes" on page 36
- "Configuring Static Routes" on page 38

Note: *For information about configuring dynamic routing on NetScreen devices, including dynamic routing protocols, see Volume 6, "Dynamic Routing".*

ROUTING ESSENTIALS

Routing is the process of forwarding packets from one network to another toward a final destination. A router is a point where one network meets another network. NetScreen security devices provide integrated routing functions that enable ScreenOS to effectively forward protected traffic to its destination.

Routing Methods

There are two types of routing you can configure on NetScreen devices: static and dynamic. When a network uses static routing, an administrator must manually configure routes and maintain the routing tables on the routers. For networks that have many connections to other networks or where inter-network connections change often, you should use dynamic routing protocols to automatically update routing tables. Dynamic routing protocols enable routers to automatically update their routing tables when network topology changes occur locally, or when neighboring routers announce changes in distant networks.

Static Routing

Static routes are mappings of IP network addresses to next-hop[1] destinations that you define on a layer 3 forwarding device, such as a router. These mappings do not change unless you alter them. For networks that have few connections to other networks or where inter-network connections are relatively unchanging, it is usually more efficient to define static routes than to set up dynamic routing. ScreenOS retains static routes until you explicitly remove them. However, you can override static routes with dynamic routing information if necessary.

Dynamic Routing

Dynamic routing involves routers exchanging information about the reachability of networks and subnetworks and adjusting routing tables by analyzing incoming routing update messages. These messages populate the network, directing routers to recalculate routes and change their routing tables accordingly. For more information about dynamic routing protocols and configuring dynamic routing on NetScreen devices, see Volume 6**,** “Dynamic Routing”.

1. A next-hop destination is a router.

Routing Tables

Typically, routers are attached to multiple networks and are responsible for directing traffic across these networks. Each router maintains a routing table, which is a list of known networks and directions on how to reach them. While processing an incoming packet on a NetScreen device, ScreenOS performs a routing table lookup to find the appropriate interface that leads to the destination address. See Chapter 1, “ScreenOS Architecture” for more information on the packet flow sequence in ScreenOS.

Each entry in a routing table — called a *route entry* or simply a *route* — is identified by the destination network to which traffic can be forwarded. The destination network, in the form of an IP address and netmask, can be an IP network, subnetwork, supernet, or a host. ScreenOS routing table entries entries can originate from the following sources:

- Directly-connected networks (the destination network is the IP address that you assign to an interface in Route mode)[2]
- Dynamic routing protocols, such as OSPF, BGP, or RIP
- Routes that are imported from other routers or virtual routers
- Statically-configured routes

2. When you set an IP address for an interface in Route mode, the routing table automatically creates a connected route to the adjacent subnet for traffic traversing the interface.

The following is an example of a ScreenOS routing table:

```
C - Connected, S - Static, A - Auto-Exported, I - Imported
    iB - IBGP, eB - EBGP, R - RIP, O - OSPF, E1 - OSPF external type 1
    E2 - OSPF external type 2

Total 8 entries

     ID  IP-Prefix       Interface     Gateway       P     Pref    Mtr   Vsys
-------------------------------------------------------------------------------
*     9 0.0.0.0/0          eth3        10.31.1.1     eB     40     100   root
*    11 192.168.1.100/32   eth2        10.3.3.100    iB     250    0     root
*    10 1.1.0.0/16         eth3        10.31.1.1     eB     40     100   root
*     4 10.1.1.1/32        eth3        10.2.2.250    S      20     1     root
*     1 192.168.1.1/32     eth1        0.0.0.0       C      0      0     root
*     5 2.2.0.0/16         eth3        10.2.2.250    S      20     1     root
*     2 10.3.3.0/24        eth2        0.0.0.0       C      0      0     root
*     3 10.2.2.0/24        eth3        0.0.0.0       C      0      0     root
```

Default route

Destination Network | Interface to Forward Data | Next Hop | Protocol | Preference | Metric | Vsys

For each destination network, the routing table contains the following information:

- The interface on the NetScreen device on which traffic for the destination network is forwarded.
- The next-hop, which can be either another virtual router on the NetScreen device or a gateway IP address (usually a router address).
- The protocol from which the route is derived.
- The *preference* is used to select the route to use when there are multiple routes to the same destination network. This value is determined by the protocol or the origin of the route. The lower the preference value of a route, the more likely the route is to be selected as the active route.

  You can modify the preference value for each protocol or route origin on a per-virtual router basis. See the “Virtual Routers” chapter in Volume 6 for more information.

- The *metric* can also be used to select the route to use when there are multiple routes for the same destination network with the same preference value. The metric value for connected routes is always 0. The default metric value for static routes is 1, but you can specify a different value when defining a static route.
- The virtual system (vsys) to which this route belongs. For more information about virtual routers and vsys, see the “Virtual Routers” chapter in Volume 6.

Most routing tables include a *default route* (network address 0.0.0.0/0), which is a catch-all entry for packets that are destined for networks other than those defined in the routing table.

Routing with Static Routes

When a host sends packets to another host that resides on a different network, each packet header contains the address of the destination host. When a router receives a packet, it compares the destination address to all addresses contained in its routing table. The router selects the most specific[3] route in the routing table to the destination address and, from the selected route entry, determines the next-hop to forward the packet.

The following illustration represents a network using static routing. For the purpose of this representation, host 1 in network A wants to reach host 2 in network C and thus creates a packet that contains the following in the header:

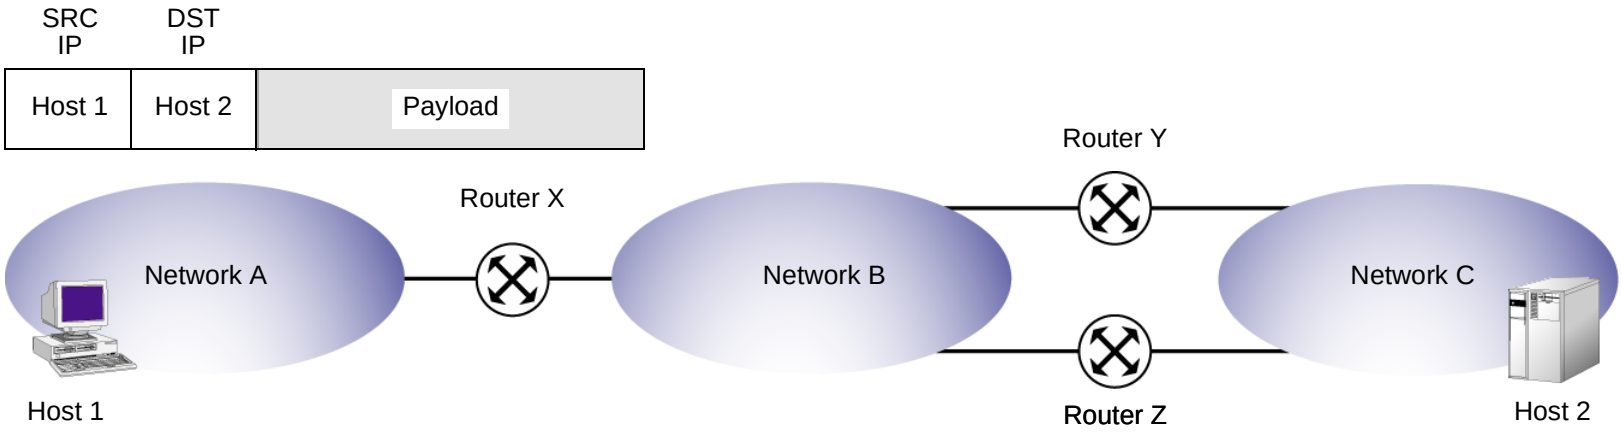

3. The most specific route is determined by first performing a bit-wise logical AND of the destination address and network mask for each entry in the routing table. For example, a bit-wise logical AND of the IP address 10.1.1.1 with the subnet mask 255.255.255.0 is 10.1.1.0. The route that has the highest number of bits set to 1 in the subnet mask is the most specific route (also called the “longest matching route”).

The following is a representation of the routing table on each router.

Routing Tables

Router X

Network	Gateway
Net A	Connected
Net B	Connected
Net C	Router Y

Router Y

Network	Gateway
Net A	Router X
Net B	Connected
Net C	Connected

Router Z

Network	Gateway
Net A	Router X
Net B	Connected
Net C	Connected

In the example above, router X has a static route configured for network C with the gateway (next-hop) as router Y. When router X receives the packet destined for host 2 in network C, it compares the destination address in the packet with its routing table and finds that the last route entry in the table is the most specific route to the destination address. The last route entry specifies to send traffic destined for network C to router Y for delivery. Router Y receives the packet and because it knows that network C is directly connected, it sends the packet through the interface connected to that network.

Note that if router Y fails or the link between router Y and network C is unavailable, the packet cannot reach host 2. While there is another route for network C through router Z, that route has not been statically configured on router X, so router X does not know about the alternate route.

VIRTUAL ROUTERS ON NETSCREEN DEVICES

ScreenOS can divide its routing component into two or more virtual routers. A virtual router supports static routing and dynamic routing protocols, which you can enable simultaneously in one virtual router. There are two predefined virtual routers on NetScreen devices:

- trust-vr, which by default contains all the predefined security zones and any user-defined zones
- untrust-vr, which by default does not contain any security zones

Some NetScreen devices allow you to create additional custom virtual routers. By separating routing information into two (or more) virtual routers, you can control the information in a given routing domain that is visible to other routing domains. For example, you can keep the routing information for all the security zones inside a corporate network on the predefined virtual router trust-vr, and the routing information for all the zones outside the corporate network on the other predefined virtual router untrust-vr. Because the information in the routing table of one virtual router is not visible to the other, you can keep internal network routing information separate from untrusted sources outside the company. This also means that traffic from zones in one virtual router are *not* automatically forwarded to zones in another virtual router even if there are policies that permit the traffic. If you want traffic to pass between virtual routers, you need to either export routes between the VRs or configure a static route in one VR that defines the other VR as the next-hop.

This chapter does not include information about creating custom virtual routers, using two or more virtual routers, or exporting routes between VRs. For more information about virtual routers, see “Virtual Routers” in Volume 6.

When to Configure Static Routes

The routing table provides information that helps a virtual router direct traffic to different interfaces and subnets. You probably need to define static routes even if you are using dynamic routing in the NetScreen device. You need to define static routes for conditions such as the following:

- If a network is not directly connected to the NetScreen device but is accessible through a router from an interface within a VR, you need to define a static route for the network with the IP address of the router. For example, the Untrust zone interface can be on a subnet with two routers that each connect to different Internet connections and you must define which router to use for forwarding traffic to specific ISPs.
- You need to define a static route to add a default route (0.0.0.0/0) into a virtual router's routing table. For example, if you are using two virtual routers on the same NetScreen device, the trust-vr routing table could contain a default route that specifies the untrust-vr as the next hop. This allows traffic for destinations that are not in the trust-vr routing table to be routed to the untrust-vr. You can also define a default route in the untrust-vr to route traffic for destinations not found in the untrust-vr routing table to a specific router IP address.
- If you are using two virtual routers on the same NetScreen device, and inbound traffic arrives on an untrust-vr interface that is destined for a network connected to a trust-vr interface, you need to define a static entry in the untrust-vr routing table for the destination network with the trust-vr as the next hop. (Note that if routing table entries in the trust-vr are exported to the untrust-vr, you do not need to define this static route.)
- When the device is in transparent mode, you must define static routes that direct management traffic originating from the device itself (as opposed to user traffic traversing the firewall) to remote destinations. For example, you need to define static routes directing syslog, SNMP, and WebTrends messages to a remote administrator's address. You must also define routes that direct authentication requests to the RADIUS, SecurID, and LDAP servers, and URL checks to the Websense server.

Note: *When the NetScreen device is in Transparent mode, you must define a static route for management traffic from the device even if the destination is on the same subnet as the device. This route is necessary to specify the interface through which to send traffic.*

- For outbound VPN traffic where there is more than one outgoing interface to the destination, you need to set a route for directing the outbound traffic through the desired interface to the external router.
- If the operational mode of an interface for a security zone in the trust-vr routing domain is NAT, and if you configured a MIP or VIP on that interface to receive incoming traffic from a source in the untrust-vr routing domain, then you must create a route to the MIP or VIP in the untrust-vr that points to the trust-vr as the gateway.

CONFIGURING STATIC ROUTES

To configure a static route, you need to define the following:

- The virtual router in which you are adding the route.
- The IP address and netmask of the destination network.
- The next hop for the route, which can be either another virtual router on the NetScreen device or a gateway (router) IP address.
- If you specify another virtual router, make sure that an entry for the destination network exists in the routing table of that virtual router.
- The interface through which the routed traffic is forwarded. The interface can be any ScreenOS-supported interface, such a physical interface (for example, ethernet1/2), or a tunnel interface.
- Optionally, you can specify a route metric and/or a route tag. The route metric is used to select the active route when there are multiple routes to the same destination network, all with the same preference value. The default metric for static routes is 1. The route tag is a value that can be used as a filter when redistributing routes. For example, you can import only routes that contain specified tag values.

Example: Configuring Static Routes

In the following example, a NetScreen device operating with its Trust zone interface in NAT mode protects a multilevel network. There is both local and remote management (via NetScreen-Security Manager). The NetScreen device sends SNMP traps and syslog reports to the local administrator (located on a network in the Trust zone) and it sends NetScreen-Security Manager reports to the remote administrator (located on a network in the Untrust zone). The device uses a SecurID server in the DMZ zone to authenticate users and a Websense server in the Trust zone to perform URL filtering.

The trust-vr and untrust-vr routing tables must contain routes for the following destinations (the numbers below correspond to the graphic shown on page 40):

untrust-vr

1. Default gateway to the Internet (this is the default route for the VR)
2. Remote administrator in the 3.3.3.0/24 subnet
3. The 2.2.40.0/24 subnet in the DMZ zone
4. The 2.20.0.0/16 subnet in the DMZ zone

trust-vr

5. untrust-vr for all addresses not found in the trust-vr routing table (this is the default route for the VR)
6. The 10.10.0.0/16 subnet in the Trust zone
7. The 10.20.0.0/16 subnet in the Trust zone
8. The 10.30.1.0/24 subnet in the Trust zone

> ***Note:*** *The following example assumes that you have already bound ethernet1 to the Trust zone, ethernet2 to the DMZ zone, and ethernet3 to the Untrust zone. The interface IP addresses are 10.1.1.1/24, 2.2.10.1/24, and 2.2.2.1/24 respectively.*

Static Route Configuration

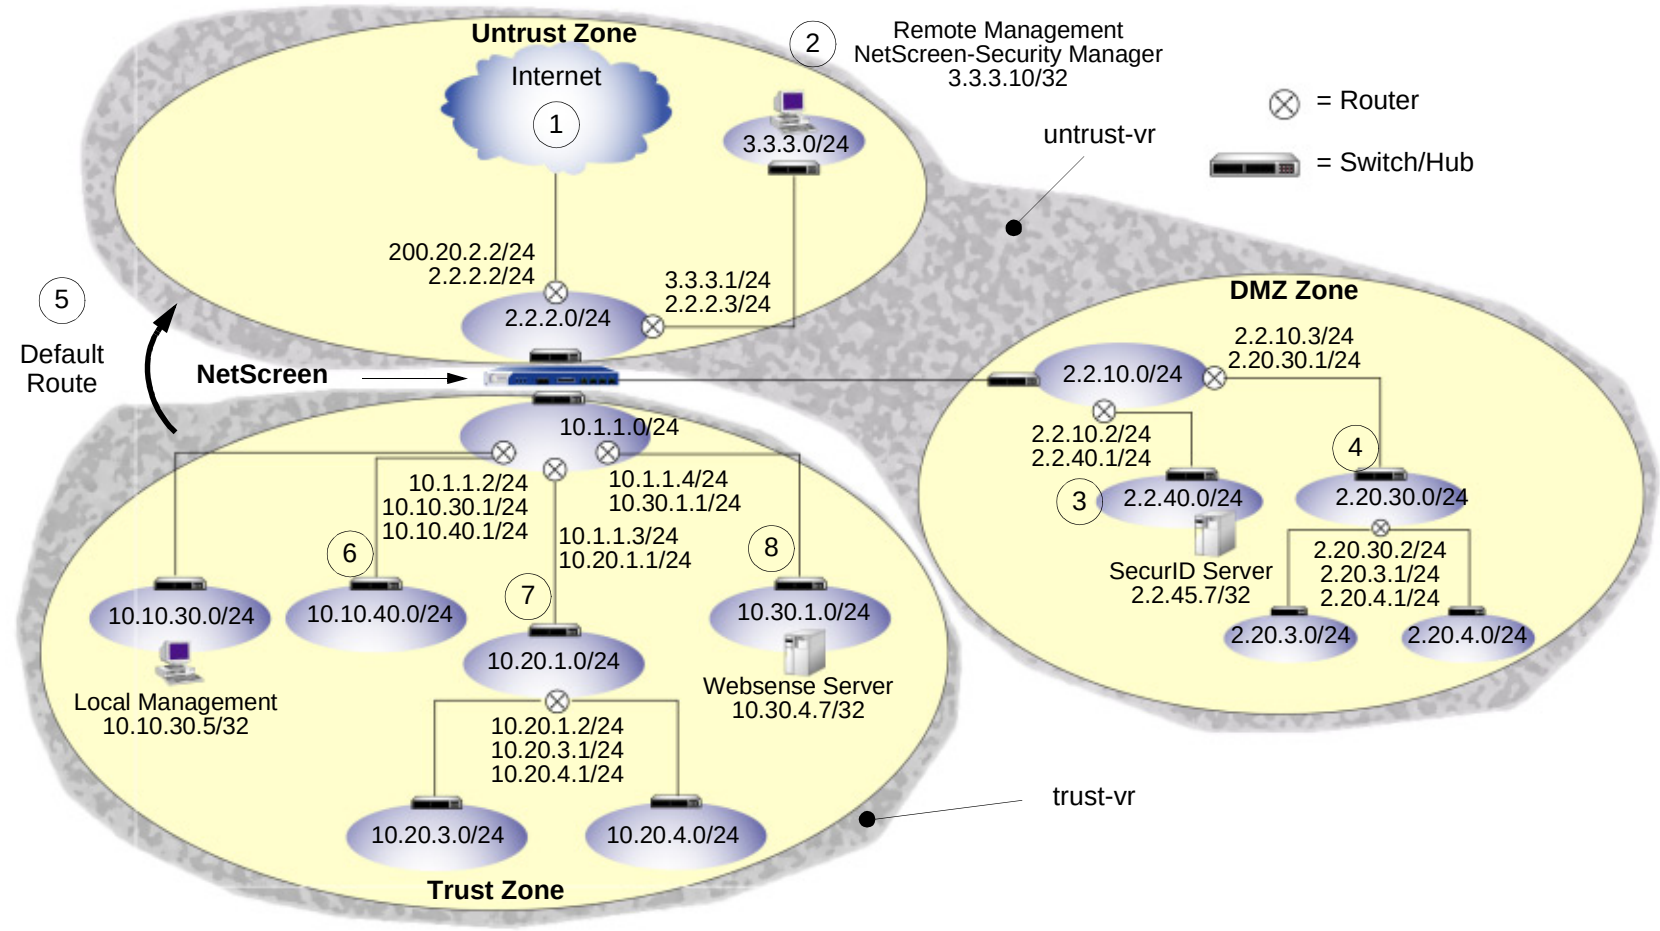

WebUI

1. untrust-vr

Network > Routing > Routing Entries > untrust-vr New: Enter the following to create the untrusted default gateway, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 2.2.2.2

Network > Routing > Routing Entries > untrust-vr New: Enter the following to direct system reports generated by the NetScreen device to remote management, and then click **OK**:

Network Address/Netmask: 3.3.3.0/24

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 2.2.2.3

Network > Routing > Routing Entries > untrust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 2.2.40.0/24

Gateway: (select)

Interface: ethernet2

Gateway IP Address: 2.2.10.2

Network > Routing > Routing Entries > untrust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 2.20.0.0/16

Gateway: (select)

Interface: ethernet2

Gateway IP Address: 2.2.10.3

2. trust-vr

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Next Hop Virtual Router Name: (select); untrust-vr

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 10.10.0.0/16

Gateway: (select)

Interface: ethernet1

Gateway IP Address: 10.1.1.2

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 10.20.0.0/16

Gateway: (select)

Interface: ethernet1

Gateway IP Address: 10.1.1.3

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 10.30.1.0/32

Gateway: (select)

Interface: ethernet1

Gateway IP Address: 10.1.1.4

Note: *To remove an entry, click* ***Remove****. A System Message appears prompting you to confirm the removal. Click* ***OK*** *to proceed, or* ***Cancel*** *to cancel the action.*

CLI

1. untrust-vr

```
set vrouter untrust-vr route 0.0.0.0/0 interface ethernet3 gateway 2.2.2.2
set vrouter untrust-vr route 3.3.3.0/24 interface ethernet3 gateway 2.2.2.3
set vrouter untrust-vr route 2.2.40.0/24 interface ethernet2 gateway 2.2.10.2
set vrouter untrust-vr route 2.20.0.0/16 interface ethernet2 gateway 2.2.10.3
```

2. trust-vr

```
set vrouter trust-vr route 0.0.0.0/0 vrouter untrust-vr
set vrouter trust-vr route 10.10.0.0/16 interface ethernet1 gateway 10.1.1.2
set vrouter trust-vr route 10.20.0.0/16 interface ethernet1 gateway 10.1.1.3
set vrouter trust-vr route 10.30.1.0/24 interface ethernet1 gateway 10.1.1.4
save
```

Example: Static Route through a Tunnel Interface

In this example, a trusted host resides in a different subnet than the trusted interface. An FTP server receives inbound traffic through a VPN tunnel. You need to set a route to direct traffic exiting the tunnel interface to the internal router leading to the subnet where the server resides.

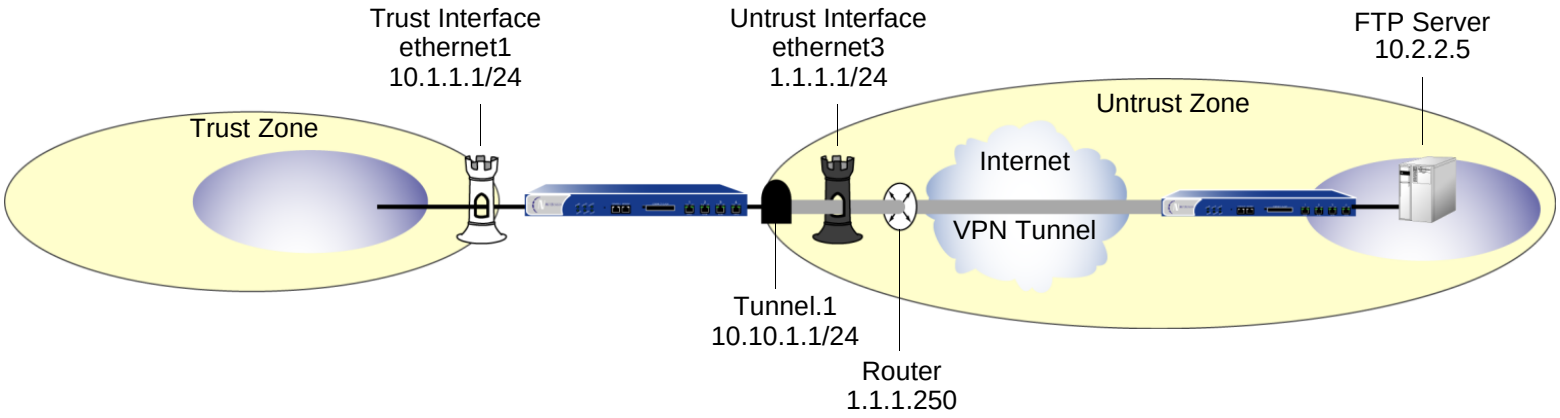

WebUI

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 10.2.2.5/32

Gateway: (select)

Interface: tunnel.1

Gateway IP Address: 0.0.0.0

Note: *For* ***tunnel.1*** *to appear in the Interface drop-down list, you must first create the tunnel.1 interface.*

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 1.1.1.250

CLI

```
set vrouter trust-vr route 10.2.2.5/32 interface tunnel.1
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
save
```

3

Zones

A zone can be a segment of network space to which security measures are applied (a security zone), a logical segment to which a VPN tunnel interface is bound (a tunnel zone), or either a physical or logical entity that performs a specific function (a function zone). This chapter examines each type of zone, with particular emphasis given to the security zone, and is organized into the following sections:

- "Security Zones" on page 48
  - "Global Zone" on page 48
  - "SCREEN Options" on page 48
- "Tunnel Zones" on page 49
- "Configuring Security Zones and Tunnel Zones" on page 51
  - "Creating a Zone" on page 51
  - "Modifying a Zone" on page 52
  - "Deleting a Zone" on page 53
- "Function Zones" on page 54
  - "Null Zone" on page 54
  - "MGT Zone" on page 54
  - "HA Zone" on page 54
  - "Self Zone" on page 54
  - "VLAN Zone" on page 54
- "Port Modes" on page 55
  - "Setting the Port Mode on NetScreen Appliances" on page 59
  - "Home Zone/Work Zone" on page 61

When you first boot up a NetScreen device, you can see a number of preconfigured zones. In the WebUI, click **Network > Zones** in the menu column on the left. In the CLI, use the **get zone** command.

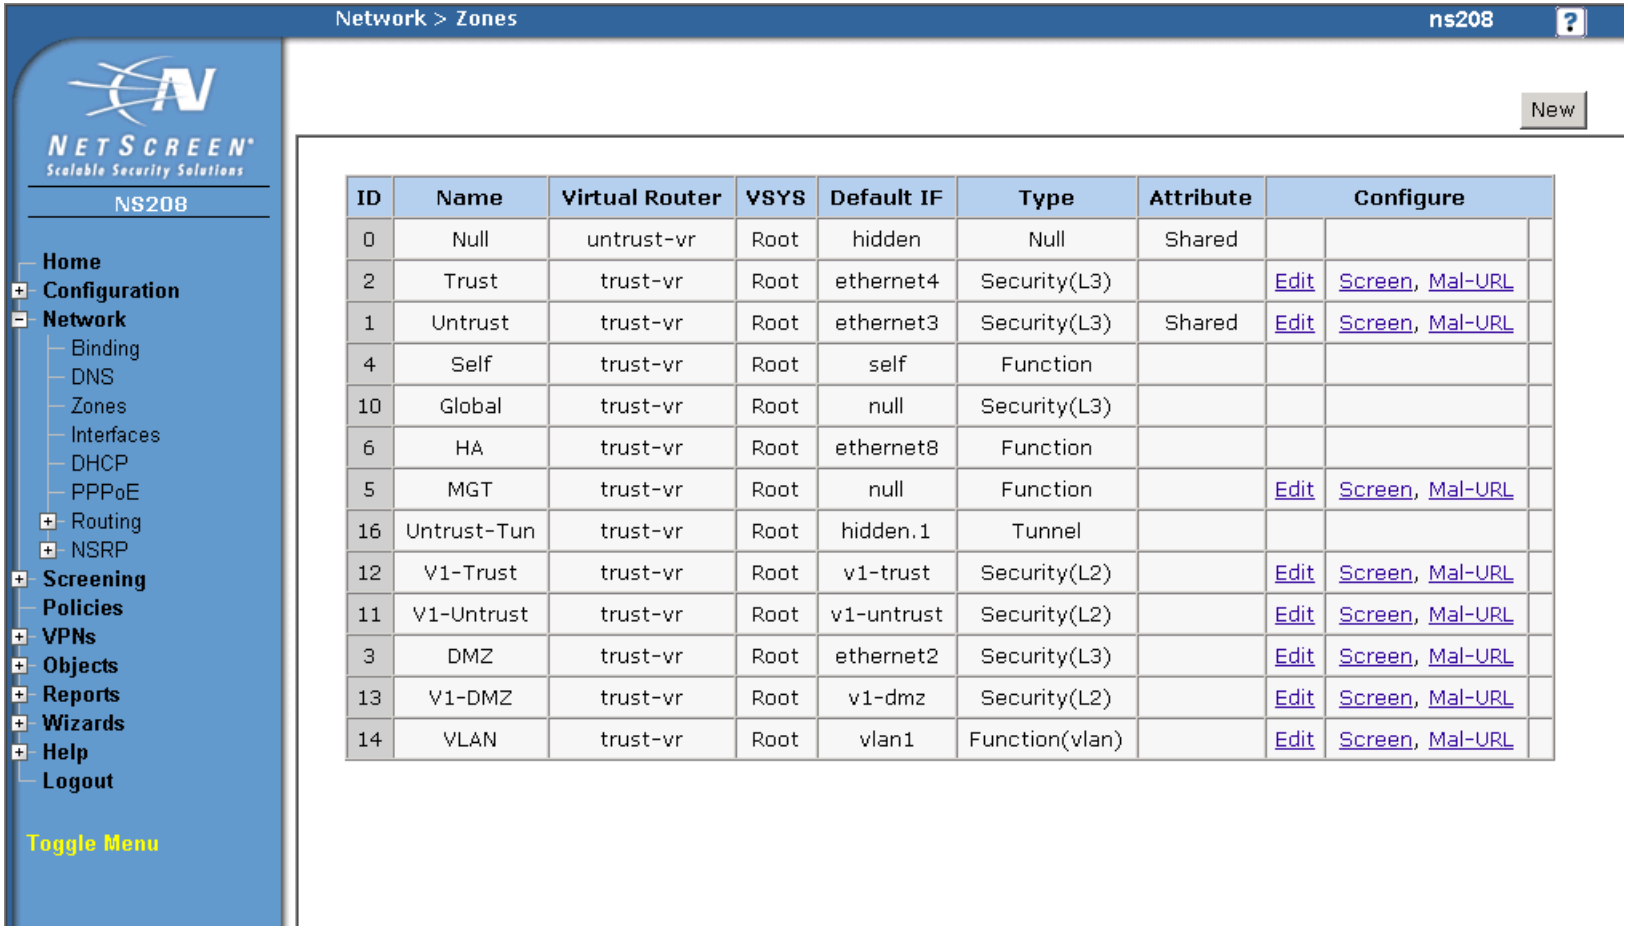

ID	Name	Virtual Router	VSYS	Default IF	Type	Attribute	Configure	
0	Null	untrust-vr	Root	hidden	Null	Shared		
2	Trust	trust-vr	Root	ethernet4	Security(L3)		Edit	Screen, Mal-URL
1	Untrust	trust-vr	Root	ethernet3	Security(L3)	Shared	Edit	Screen, Mal-URL
4	Self	trust-vr	Root	self	Function			
10	Global	trust-vr	Root	null	Security(L3)			
6	HA	trust-vr	Root	ethernet8	Function			
5	MGT	trust-vr	Root	null	Function		Edit	Screen, Mal-URL
16	Untrust-Tun	trust-vr	Root	hidden.1	Tunnel			
12	V1-Trust	trust-vr	Root	v1-trust	Security(L2)		Edit	Screen, Mal-URL
11	V1-Untrust	trust-vr	Root	v1-untrust	Security(L2)		Edit	Screen, Mal-URL
3	DMZ	trust-vr	Root	ethernet2	Security(L3)		Edit	Screen, Mal-URL
13	V1-DMZ	trust-vr	Root	v1-dmz	Security(L2)		Edit	Screen, Mal-URL
14	VLAN	trust-vr	Root	vlan1	Function(vlan)		Edit	Screen, Mal-URL

The output of the **get zone** command:

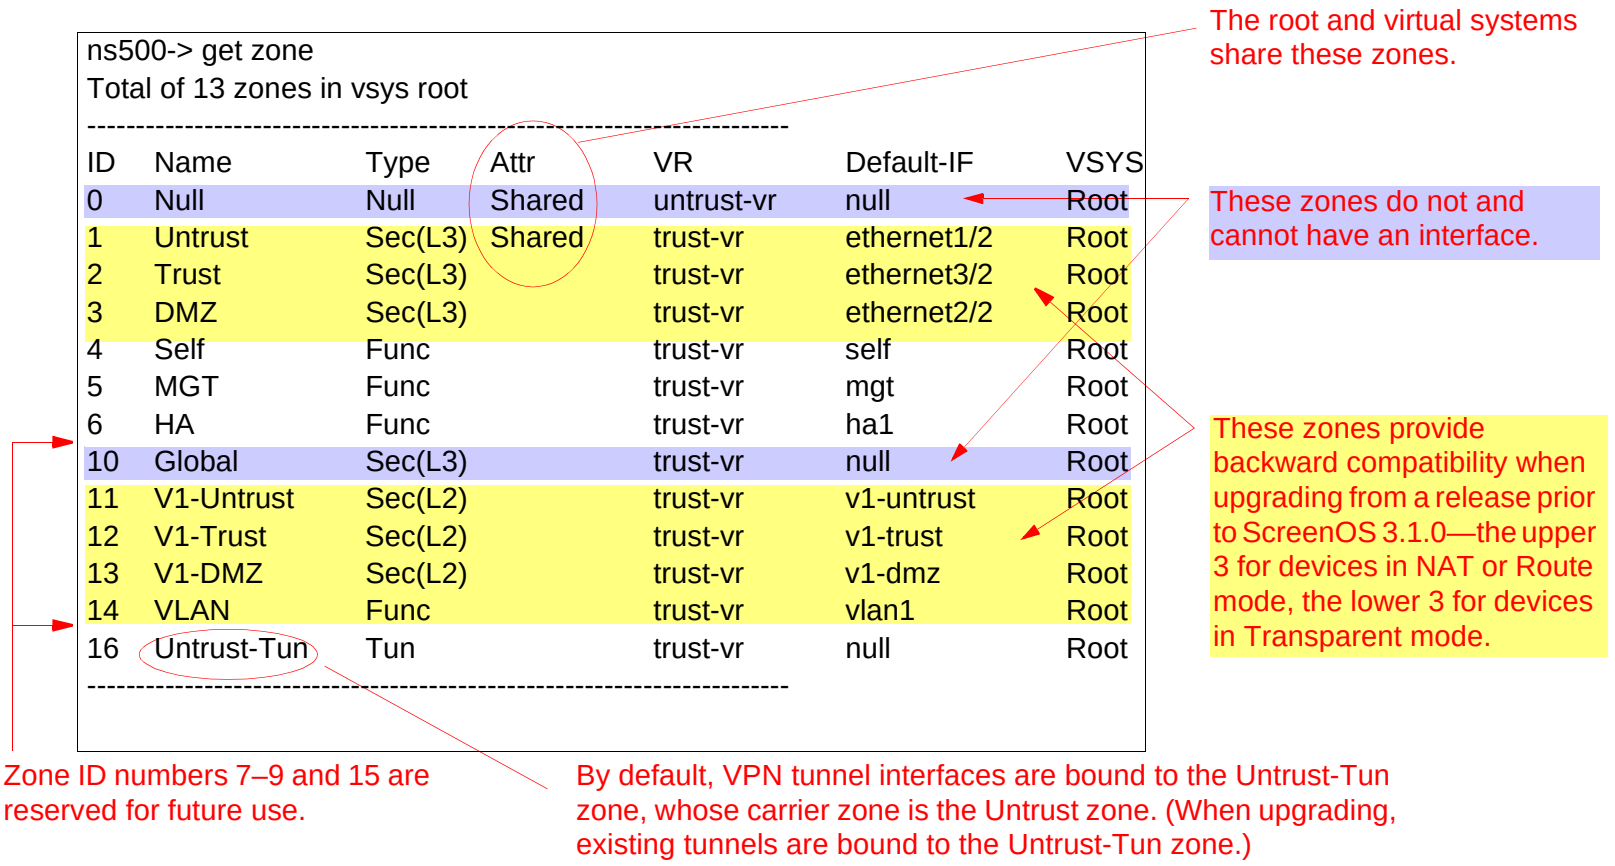

The preconfigured zones shown above can be grouped into three different types:

Security Zones: Untrust, Trust, DMZ, Global, V1-Untrust, V1-Trust, V1-DMZ

Tunnel Zone: Untrust-Tun

Function Zones: Null, Self, MGT, HA, VLAN

SECURITY ZONES

On a single NetScreen device, you can configure multiple security zones, sectioning the network into segments to which you can apply various security options to satisfy the needs of each segment. At a minimum, you must define two security zones, basically to protect one area of the network from the other. On some NetScreen platforms, you can define many security zones, bringing finer granularity to your network security design— and without deploying multiple security appliances to do so.

Global Zone

You can identify a security zone because it has an address book and can be referenced in policies. The Global zone satisfies these criteria. However, it does not have one element that all other security zones have—an interface. The Global zone serves as a storage area for mapped IP (MIP) and virtual IP (VIP) addresses. The predefined Global zone address “Any” applies to all MIPs, VIPs, and other user-defined addresses set in the Global zone. Because traffic going to these addresses is mapped to other addresses, the Global zone does not require an interface for traffic to flow through it.

The Global zone also contains addresses for use in global policies. For information about global policies, see “Global Policies” on page 201.

Note: *Any policy that uses the Global zone as its destination cannot support NAT or traffic shaping.*

SCREEN Options

A NetScreen firewall secures a network by inspecting, and then allowing or denying, all connection attempts that require passage from one security zone to another. For every security zone, and the MGT zone, you can enable a set of predefined SCREEN options that detect and block various kinds of traffic that the NetScreen determines as potentially harmful. For more information about the many SCREEN options available, see Volume 4, “Attack Detection and Defense Mechanisms”.

TUNNEL ZONES

A tunnel zone is a logical segment that hosts one or more tunnel interfaces. A tunnel zone is conceptually affiliated with a security zone in a “child-parent” relationship. The security zone acting as the “parent”, which you can also conceive of as a carrier zone, provides the firewall protection to the encapsulated traffic. The tunnel zone provides packet encapsulation/decapsulation, and—by supporting tunnel interfaces with IP addresses and netmasks that can host mapped IP (MIP) addresses and dynamic IP (DIP) pools—can also provide policy-based NAT services.

The NetScreen device uses the routing information for the carrier zone to direct traffic to the tunnel endpoint. The default tunnel zone is Untrust-Tun, and it is associated with the Untrust zone. You can create other tunnel zones and bind them to other security zones, with a maximum of one tunnel zone per carrier zone per virtual system[1].

By default, a tunnel zone is in the trust-vr routing domain, but you can also move a tunnel zone into another routing domain.

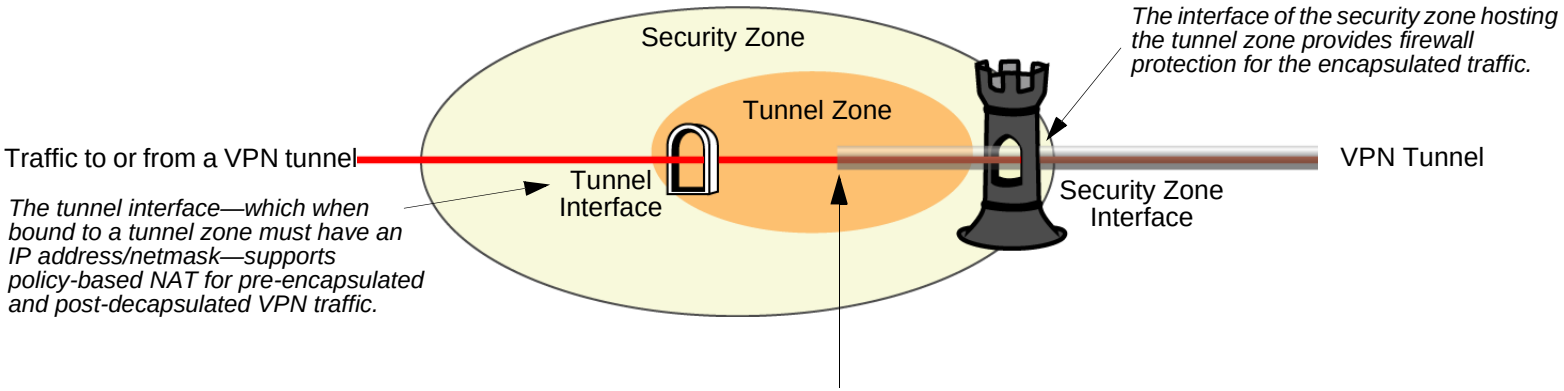

Outbound traffic enters the tunnel zone via the tunnel interface, is encapsulated, and exits via the security zone interface. Inbound traffic enters via the security zone interface, is decapsulated in the tunnel zone, and exits via the tunnel interface.

When upgrading from a version of ScreenOS earlier than 3.1.0, existing tunnel interfaces are bound by default to the preconfigured Untrust-Tun tunnel zone, which is a “child” of the preconfigured Untrust security zone. You can bind multiple tunnel zones to the same security zone; however, you cannot bind a tunnel zone to another tunnel zone.

1. The root system and all virtual systems can share the Untrust zone. However, each system has its own separate Untrust-Tun zone.

Example: Binding a Tunnel Interface to a Tunnel Zone

In this example, you create a tunnel interface and name it tunnel.3. You bind it to the Untrust-Tun zone, and assign it IP address 3.3.3.3/24. You then define a mapped IP (MIP) address on tunnel.3, translating 3.3.3.5 to 10.1.1.5, which is the address of a server in the Trust zone. Both the Untrust zone, which is the carrier zone for the Untrust-Tun zone, and the Trust zone are in the trust-vr routing domain.

WebUI

1. Tunnel Interface

   Network > Interfaces > New Tunnel IF: Enter the following, and then click **OK**:

   Tunnel Interface Name: tunnel.3

   Zone (VR): Untrust-Tun (trust-vr)

   Fixed IP: (select)

   IP Address / Netmask 3.3.3.3/24

2. MIP

   Network > Interfaces > Edit (for tunnel.3) > MIP > New: Enter the following, and then click **OK**:

   Mapped IP: 3.3.3.5

   Netmask: 255.255.255.255

   Host IP Address: 10.1.1.5

   Host Virtual Router Name: trust-vr

CLI

1. Tunnel Interface

```
set interface tunnel.3 zone Untrust-Tun
set interface tunnel.3 ip 3.3.3.3/24
```

2. MIP

```
set interface tunnel.3 mip 3.3.3.5 host 10.1.1.5
save
```

CONFIGURING SECURITY ZONES AND TUNNEL ZONES

The creation, modification and deletion of Layer 3 or Layer 2 security zones and tunnel zones are quite similar.

Note: *You cannot delete predefined security zones or the predefined tunnel zone, although you can edit them.*

Creating a Zone

To create a Layer 3 or Layer 2 security zone, or a tunnel zone, use either the WebUI or CLI:

WebUI

Network > Zones > New: Enter the following, and then click **OK**:

Zone Name: Type a name for the zone[2].

Virtual Router Name: Select the virtual router in whose routing domain you want to place the zone.

Zone Type: Select **Layer 3** to create a zone to which you can bind interfaces in NAT or Route mode. Select **Layer 2** to create a zone to which you can bind interfaces in Transparent mode. Select **Tunnel Out Zone** when creating a tunnel zone and binding it to a carrier zone, and then select a specific carrier zone from the drop-down list.

Block Intra-Zone Traffic: Select this option to block traffic between hosts within the same security zone. By default, intra-zone blocking is disabled.

CLI

```
set zone name zone [ l2 vlan_id_num[3] | tunnel sec_zone ]
set zone zone block
set zone zone vrouter name_str
```

2. The name of a Layer 2 security zone must begin with “L2-”; for example, “L2-Corp” or “L2-XNet”.
3. When creating a Layer 2 security zone, the VLAN ID number must be 1 (for VLAN1).

Modifying a Zone

To modify the name of a security zone or tunnel zone, or to change the carrier zone for a tunnel zone, you must first delete the zone[4], and then create it again with the changes. You can change the intra-zone blocking option and the virtual router[5] on an existing zone.

WebUI

1. Modifying the Zone Name

Network > Zones: Click **Remove** (for the security zone or tunnel zone whose name you want to change, or for the tunnel zone whose carrier zone you want to change).

When the prompt appears, asking for confirmation of the removal, click **Yes**.

Network > Zones > New: Enter the zone settings with your changes, and then click **OK**.

2. Changing the Intra-Zone Blocking Option or Virtual Router

Network > Zones > Edit (for the zone that you want to modify): Enter the following, and then click **OK**:

Virtual Router Name: From the drop-down list, select the virtual router into whose routing domain you want to move the zone.

Block Intra-Zone Traffic: To enable, select the check box. To disable, clear it.

CLI

1. Modifying the Zone Name

```
unset zone zone
set zone name zone [ l2 vlan_id_num | tunnel sec_zone ]
```

2. Changing the Intra-Zone Blocking Option or Virtual Router

```
{ set | unset } zone zone block
set zone zone vrouter name_str
```

4. Before you can remove a zone, you must first unbind all interfaces bound to it.
5. You must first remove any interfaces bound to a zone before changing its virtual router.

Deleting a Zone

To delete a security zone or tunnel zone, do either of the following[6]:

WebUI

Network > Zones: Click **Remove** (for the zone you want to delete).

When the prompt appears, asking for confirmation of the removal, click **Yes**.

CLI

```
unset zone zone
```

6. Before you can remove a zone, you must first unbind all interfaces bound to it. To unbind an interface from a zone, see "Binding an Interface to a Security Zone" on page 76.

FUNCTION ZONES

The five function zones are Null, MGT, HA, Self, and VLAN. Each zone exists for a single purpose, as explained below.

Null Zone

This zone serves as temporary storage for any interfaces that are not bound to any other zone.

MGT Zone

This zone hosts the out-of-band management interface, MGT. You can set firewall options on this zone to protect the management interface from different types of attacks. For more information on firewall options, see Volume 4, “Attack Detection and Defense Mechanisms”.

HA Zone

This zone hosts the high availability interfaces, HA1 and HA2. Although you can set interfaces for the HA zone, the zone itself is not configurable.

Self Zone

This zone hosts the interface for remote management connections. When you connect to the NetScreen device via HTTP, SCS, or Telnet, you connect to the Self zone.

VLAN Zone

This zone hosts the VLAN1 interface, which you use to manage the device and terminate VPN traffic when the device is in Transparent mode. You can also set firewall options on this zone to protect the VLAN1 interface from various attacks.

PORT MODES

You can select a *port mode* for some NetScreen appliances. The port mode automatically sets different port, interface, and zone bindings[7] for the device. For example, on the NetScreen-5XT, you can configure one of the following four port modes:

Warning: Changing the port mode removes any existing configurations on the NetScreen device, and requires a system reset.

- Trust-Untrust mode is the default port mode. This mode provides the following port, interface, and zone bindings:
  - Binds the Untrusted Ethernet port to the Untrust interface, which is bound to the Untrust security zone
  - Binds the Modem port to the serial interface, which you can bind as a backup interface to the Untrust security zone
  - Binds the Ethernet ports 1 through 4 to the Trust interface, which is bound to the Trust security zone

Trust-Untrust

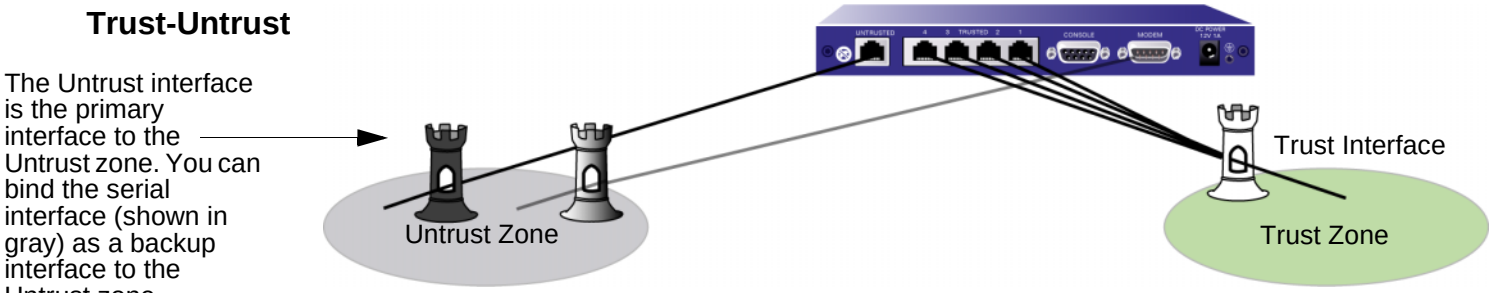

***Note:** The Initial Configuration Wizard only runs in Trust-Untrust port mode.*

7. In the port mode context, *port* refers to a physical interface on the back of the NetScreen device. The ports are referenced by their labels: Untrusted, 1-4, Console, or Modem. The term *interface* refers to a logical interface that can be configured through the WebUI or CLI. Each port can be bound to only one interface, but multiple ports can be bound to an interface.

- Home-Work mode binds interfaces to the Untrust security zone and to new Home and Work security zones. The Work and Home zones allow you to segregate users and resources in each zone. In this mode, default policies allow traffic flow and connections from the Work zone to the Home zone, but do not allow traffic from the Home zone to the Untrust zone. By default, there are no restrictions for traffic from the Home zone to the Untrust zone. This mode provides the following port, interface, and zone bindings:
  - Binds the Ethernet ports 1 and 2 to the ethernet1 interface, which is bound to the Work security zone
  - Binds the Ethernet ports 3 and 4 to the ethernet2 interface, which is bound to the Home security zone
  - Binds the Untrusted Ethernet port to the ethernet3 interface, which is bound to the Untrust security zone
  - Binds the Modem port to the serial interface, which you can bind as a backup interface to the Untrust security zone

Home-Work

The ethernet3 interface is the primary interface to the Untrust zone. You can bind the serial interface (shown in gray) as a backup interface to the Untrust zone.

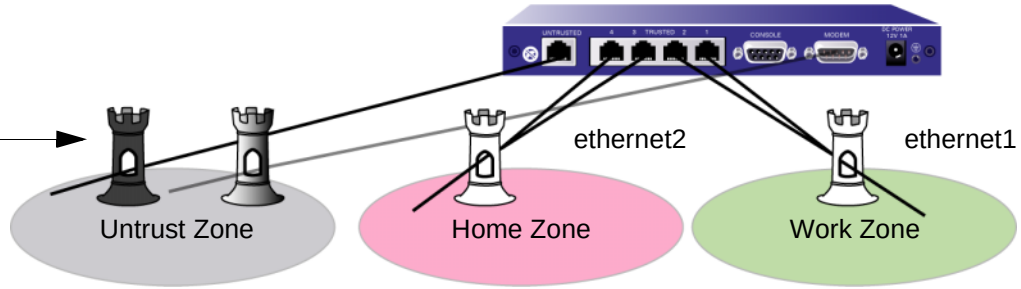

See “Home Zone/Work Zone” on page 61 for more information about configuring and using Home-Work mode.

- Dual Untrust mode binds two interfaces, a primary and a backup, to the Untrust security zone. The primary interface is used to pass traffic to and from the Untrust zone, while the backup interface is used only when there is a failure on the primary interface. This mode provides the following port, interface, and zone bindings:
  - Binds the Untrusted Ethernet port to the ethernet3 interface, which is bound to the Untrust security zone
  - Binds Ethernet port 4 to the ethernet2 interface, which is bound as a backup interface to the Untrust security zone (the ethernet3 interface is the primary interface to the Untrust security zone)
  - Binds the Ethernet ports 1, 2, and 3 to the ethernet1 interface, which is bound to the Trust security zone

Dual Untrust

The eth3 interface is the primary interface to the Untrust zone. The eth2 interface (shown in gray) is a backup interface to the Untrust zone.

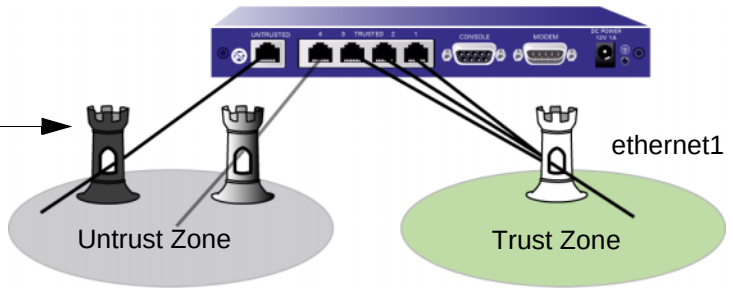

Note: *The serial interface is not available in Dual Untrust port mode.*

See Volume 8, “High Availability” for more information about configuring and using Dual Untrust mode.

- Combined mode allows both primary and backup interfaces to the Internet and the segregation of users and resources in Work and Home zones.

  Note: *For the NetScreen-5XT, the Combined port mode is supported only on the NetScreen-5XT Elite (unrestricted users) platform.*

  This mode provides the following port, interface, and zone bindings:

  – Binds the Untrusted Ethernet port to the ethernet4 interface, which is bound to the Untrust zone
  – Binds Ethernet port 4 to the ethernet3 interface, which is bound as a backup interface to the Untrust zone (the ethernet4 interface is the primary interface to the Untrust security zone)
  – Binds the Ethernet ports 3 and 2 to the ethernet2 interface, which is bound to the Home zone
  – Binds Ethernet port 1 to the ethernet1 interface, which is bound to the Work zone

Combined

The ethernet4 interface is the primary interface to the Untrust zone. The ethernet3 interface (shown in gray) is the backup interface to the Untrust zone.

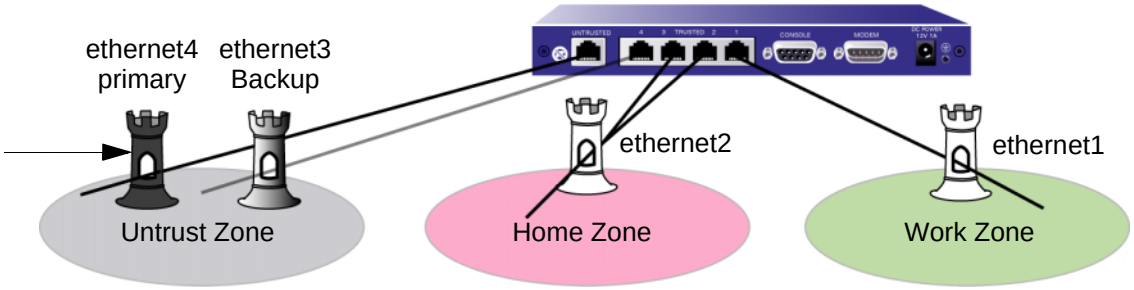

Note: *The serial interface is not available in Combined port mode.*

See Volume 8, “High Availability” and “Home Zone/Work Zone” on page 61 for more information about configuring and using the Combined mode.

Setting the Port Mode on NetScreen Appliances

The following table summarizes the port, interface, and zone bindings provided by the ScreenOS port modes:

Port*	Trust-Untrust Mode†		Home-Work Mode		Dual Untrust Mode		Combined Mode	
	Interface	Zone	Interface	Zone	Interface	Zone	Interface	Zone
Untrusted	Untrust	Untrust	ethernet3	Untrust	ethernet3	Untrust	ethernet4	Untrust
1	Trust	Trust	ethernet1	Work	ethernet1	Trust	ethernet1	Work
2	Trust	Trust	ethernet1	Work	ethernet1	Trust	ethernet2	Home
3	Trust	Trust	ethernet2	Home	ethernet1	Trust	ethernet2	Home
4	Trust	Trust	ethernet2	Home	ethernet2	Untrust	ethernet3	Untrust
Modem	serial	Null	serial	Null	N/A	N/A	N/A	N/A

* As labeled on the NetScreen appliance chassis.
† Default port mode

You change the port mode setting on the NetScreen device through either the WebUI or the CLI. Before setting the port mode, note the following:

- Changing the port mode *removes* any existing configurations on the NetScreen device and requires a system reset.
- Issuing the **unset all** CLI command does not affect the port mode setting on the NetScreen device. For example, if you want to change the port mode setting from the Combined mode back to the default Trust-Untrust mode, issuing the **unset all** command removes the existing configuration but does *not* set the device to the Trust-Untrust mode.

Example: Setting Home-Work Port Mode

In this example, you set the port mode on the NetScreen-5XT to the Home-Work mode.

Note: *Changing the port mode removes any existing configurations on the NetScreen device and requires a system reset.*

WebUI

Configuration > Port Mode > Port Mode: Select Home-Work from the drop-down list, and then click **Apply**.

At the following prompt, click **OK**:

Operational mode change will erase current configuration and reset the device, continue?

CLI

```
exec port-mode home-work
```

At the following prompt, enter **y** (for yes):

Change port mode from <trust-untrust> to <home-work> will erase system configuration and reboot box

Are you sure y/[n] ?

To see the current port mode setting on the NetScreen device:

WebUI

Configuration > Port Mode

CLI

```
get system
```

Home Zone/Work Zone

Security conflicts can arise as both employee telecommuting and home networks become commonplace. The home network used by both telecommuters and family members can become a dangerous back door to a corporate network, carrying threats such as worms and allowing access to corporate resources, such as servers and networks, by non-employees.

The Home-Work and Combined port modes[8] bind ScreenOS interfaces to special Work and Home zones. This allows segregation of business and home users and resources, while allowing users in both Home and Work zones access to the Untrust zone.

Home-Work

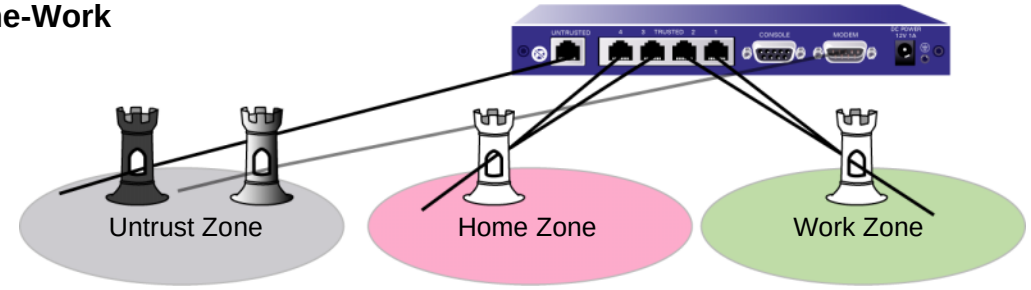

Combined

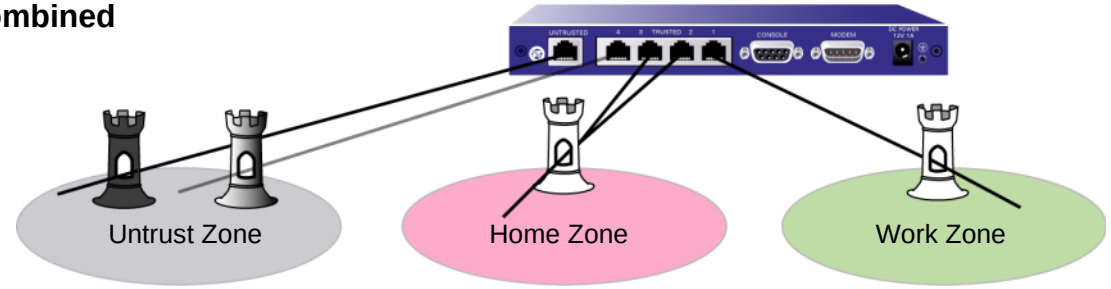

8. You can set port modes only on certain NetScreen appliances. See "Port Modes" on page 55.

The Home-Work port mode also binds the Modem port to a serial interface, which you can bind as a backup interface to the Untrust security zone. For more information about using the serial interface as a backup interface to the Untrust security zone, see Volume 8, “High Availability”.

The Combined port mode also binds the Trusted4 Ethernet port as a backup interface (ethernet3) to the Untrust security zone. The backup interface is used only when there is a failure on the primary interface to the Untrust zone. For more information about using the ethernet3 interface as a backup interface to the Untrust security zone, see Volume 8, “High Availability”.

By default, the NetScreen-5XT acts as a Dynamic Host Configuration Protocol (DHCP) server, allocating dynamic IP addresses to DHCP clients in the Work zone. (For more information about the DHCP server, see “DHCP Server” on page 502.)

You can configure the NetScreen device using a Telnet connection or the WebUI from the Work zone only. You cannot configure the NetScreen device from the Home zone. You cannot use any management services, including ping, on the Home zone interface. The default IP address of the Work zone interface, ethernet1, is 192.168.1.1/24.

The default policies in the Home-Work and Combined port modes provide the following traffic control between zones:

- Allow all traffic from the Work zone to the Untrust zone
- Allow all traffic from the Home zone to the Untrust zone
- Allow all traffic from the Work zone to the Home zone
- Block all traffic from the Home zone to the Work zone (you cannot remove this policy)

You can create new policies for traffic from the Work zone to the Untrust zone, from the Home zone to the Untrust zone, and from the Work zone to the Home zone. You can also remove the default policies that allow all traffic from the Work zone to the Untrust zone, from the Home zone to the Untrust zone, and from the Work zone to the Home zone. Note, however, that you cannot create a policy to allow traffic from the Home zone to the Work zone.

Example: Configuring Home and Work Zones

In this example, you first set a NetScreen-5XT appliance in Home-Work port mode. You then configure a policy to allow only FTP traffic from the Home zone to the Untrust zone and remove the default policy that allows all traffic from the Home zone to the Untrust zone. In this example, the default policy, which allows traffic from any source address to any destination address for any service, has an ID of 2.

Warning: *Changing the port mode removes any existing configurations on the NetScreen device and requires a system reset.*

WebUI

Configuration > Port Mode > Port Mode: Select Home-Work from the drop-down list, and then click **Apply**.

At the following prompt, click **OK**:

Operational mode change will erase current configuration and reset the device, continue?

Policies > (From: Home, To: Untrust) > New: Enter the following, and then click **OK**.

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Any

Service: FTP

Action: Permit

Policies: In the “From Home to Untrust” policy list, click **Remove** in the Configure column for the policy with ID 2.

CLI

```
exec port-mode home-work

At the following prompt, enter y (for yes):
    Change port mode from <trust-untrust> to <home-work> will erase system
         configuration and reboot box
    Are you sure y/[n] ?

set policy from home to untrust any any ftp permit
unset policy 2
save
```

4

Interfaces

Physical interfaces and subinterfaces, like doorways, allow traffic to enter and exit a security zone. To allow network traffic to flow in and out of a security zone, you must bind an interface to that zone and, if it is a Layer 3 zone, assign it an IP address. Then, you must configure policies to allow traffic to pass from interface to interface between zones. You can assign multiple interfaces to a zone, but you cannot assign a single interface to multiple zones.

This chapter contains the following sections:

- "Interface Types" on page 66
  - "Security Zone Interfaces" on page 66
  - "Function Zone Interfaces" on page 68
  - "Tunnel Interfaces" on page 69
- "Viewing Interfaces" on page 74
- "Configuring Security Zone Interfaces" on page 76
  - "Binding an Interface to a Security Zone" on page 76
  - "Defining an Address for a L3 Security Zone Interface" on page 77
  - "Unbinding an Interface from a Security Zone" on page 80
  - "Modifying Interfaces" on page 81
  - "Creating Subinterfaces" on page 82
  - "Deleting Subinterfaces" on page 83
- "Secondary IP Addresses" on page 84
  - "Secondary IP Address Properties" on page 84
- "Loopback Interfaces" on page 86

INTERFACE TYPES

This section describes security zone, function zone, and tunnel interfaces. For information on how to view a table of all these interfaces, see “Viewing Interfaces” on page 74.

Security Zone Interfaces

The purpose of physical interfaces and subinterfaces is to provide an opening through which network traffic can pass between zones.

Physical

Each port on your NetScreen device represents a physical interface, and the name of the interface is predefined. The name of a physical interface is composed of the media type, slot number (for some NetScreen devices), and port number, for example, *ethernet3/2* or *ethernet2* (see also “Security Zone Interfaces” on page 3). You can bind a physical interface to any security zone where it acts as a doorway through which traffic enters and exits the zone. Without an interface, no traffic can access the zone or leave it.

On NetScreen devices that support changes to interface-to-zone bindings, three of the physical ethernet interfaces are pre-bound to specific Layer 2 security zones—V1-Trust, V1-Untrust, and V1-DMZ. Which interface is bound to which zone is specific to each platform. (For more information on security zones, see “Security Zones” on page 2.)

Subinterface

A subinterface, like a physical interface, acts as a doorway through which traffic enters and exits a security zone. You can logically divide a physical interface into several virtual subinterfaces. Each virtual subinterface borrows the bandwidth it needs from the physical interface from which it stems, thus its name is an extension of the physical interface name, for example, *ethernet3/2.1* or *ethernet2.1.* (See also “Security Zone Interfaces” on page 3.)

You can bind a subinterface to any zone. You can bind a subinterface to the same zone as its physical interface, or you can bind it to a different zone. (For more information, see “Binding an Interface to a Security Zone” on page 76 and “Defining Subinterfaces and VLAN Tags” on page **7** -23.)

Aggregate Interfaces

The NetScreen-5000 series supports aggregate interfaces. An aggregate interface is the accumulation of two or more physical interfaces, each of which shares the traffic load directed to the IP address of the aggregate interface equally among themselves. By using an aggregate interface, you can increase the amount of bandwidth available to a single IP address. Also, if one member of an aggregate interface fails, the other member or members can continue processing traffic—although with less bandwidth than previously available.

Note: *For more information about aggregate interfaces, see* “Interface Redundancy” on page **8** *-93*.

Redundant Interfaces

You can bind two physical interfaces together to create one redundant interface, which you can then bind to a security zone. One of the two physical interfaces acts as the primary interface and handles all the traffic directed to the redundant interface. The other physical interface acts as the secondary interface and stands by in case the active interface experiences a failure. If that occurs, traffic to the redundant interface fails over to the secondary interface, which becomes the new primary interface. The use of redundant interfaces provides a first line of redundancy before escalating a failover to the device level.

Note: *For more information about redundant interfaces, see the* “Interface Redundancy” chapter in Volume 8, “High Availability”*.*

Virtual Security Interfaces

Virtual security interfaces (VSIs) are the virtual interfaces that two NetScreen devices forming a virtual security device (VSD) share when operating in high availability (HA) mode. Network and VPN traffic use the IP address and virtual MAC address of a VSI. The VSD then maps the traffic to the physical interface, subinterface, or redundant interface to which you have previously bound the VSI. When two NetScreen devices are operating in HA mode, you must bind security zone interfaces that you want to provide uninterrupted service in the event of a device failover to one or more virtual security devices (VSDs). When you bind an interface to a VSD, the result is a virtual security interface (VSI).

Note: *For more information on VSIs and how they function with VSDs in an HA cluster, see Volume 8, "High Availability".*

Function Zone Interfaces

Function zone interfaces, such as Management and HA, serve a special purpose.

Management Interface

On some NetScreen devices, you can manage the device through a separate physical interface—the Management (MGT) interface—moving administrative traffic outside the regular network user traffic. Separating administrative traffic from network user traffic greatly increases security and assures constant management bandwidth.

Note: *For information on configuring the device for administration, see "Administration" on page* ***3****-1.*

HA Interface

The HA interface is a physical port used exclusively for HA functions. With NetScreen devices that have dedicated High Availability (HA) interfaces, you can link two devices together to form a redundant group, or cluster. In a redundant group, one unit acts as the master, performing the network firewall, VPN, and traffic-shaping functions, while the other unit acts as a backup, basically waiting to take over the firewall functions should the master unit fail. This is an active/passive configuration. You can also set up both members of the cluster to be master and backup for each other. This is an active/active configuration. Both configurations are explained fully in Volume 8, “High Availability”.

Virtual HA Interface

On NetScreen devices without a dedicated HA interface, a Virtual High Availability (HA) interface provides the same functionality. Because there is no separate physical port exclusively used for HA traffic, the Virtual HA interface must be bound to one of the physical ethernet ports. You use the same procedure for binding a network interface to the HA zone as you do for binding a network interface to a security zone (see “Binding an Interface to a Security Zone” on page 76).

Note: *For more information about HA interfaces, see “Dual HA Interfaces” on page* ***8*** *-38.*

Tunnel Interfaces

A tunnel interface acts as a doorway to a VPN tunnel. Traffic enters and exits a VPN tunnel via a tunnel interface.

When you bind a tunnel interface to a VPN tunnel, you can reference that tunnel interface in a route to a specific destination and then reference that destination in one or more policies. With this approach, you can finely control the flow of traffic through the tunnel. It also provides dynamic routing support for VPN traffic. When there is no tunnel interface bound to a VPN tunnel, you must specify the tunnel in the policy itself and choose **tunnel** as the action. Because the action **tunnel** implies permission, you cannot specifically deny traffic from a VPN tunnel.

You can perform policy-based NAT on outgoing or incoming traffic using a pool of dynamic IP (DIP) addresses in the same subnet as the tunnel interface. A typical reason for using policy-based NAT on a tunnel interface is to avoid IP address conflicts between the two sites on either end of the VPN tunnel.

You must bind a route-based VPN tunnel to a tunnel interface so that the NetScreen device can route traffic to and from it. You can bind a route-based VPN tunnel to a tunnel interface that is either numbered (with IP address/netmask) or unnumbered (without IP address/netmask). If the tunnel interface is unnumbered, you must specify an interface from which the tunnel interface borrows an IP address. The NetScreen device only uses the borrowed IP address as a source address when the NetScreen device itself initiates traffic—such as OSPF messages—through the tunnel. The tunnel interface can borrow the IP address from an interface in the same security zone or from an interface in a different one as long as both zones are in the same routing domain.

You can achieve very secure control of VPN traffic routing by binding all the unnumbered tunnel interfaces to one zone, which is in its own virtual routing domain, and borrowing the IP address from a loopback interface bound to the same zone. For example, you can bind all the unnumbered tunnel interfaces to a user-defined zone named “VPN” and configure them to borrow an IP address from the loopback.1 interface, also bound to the VPN zone. The VPN zone is in a user-defined routing domain named “vpn-vr”. You put all destination addresses to which the tunnels lead in the VPN zone. Your routes to these addresses point to the tunnel interfaces, and your policies control VPN traffic between other zones and the VPN zone.

```
set vrouter name vpn-vr
set zone name vpn vrouter vpn-vr
set interface loopback.1 zone vpn
set interface loopback.1 ip 172.16.1.1/24
set interface tunnel.1 zone vpn
set interface tunnel.1 ip unnumbered loopback.1

Configure addresses for src-1 and dst-1.
Configure a VPN tunnel and bind it to tunnel.1.

set vrouter trust-vr route 10.2.2.5/32 vrouter vpn-vr
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3
     gateway 1.1.1.250
set vrouter vpn-vr route 10.2.2.5 interface tunnel.1

set policy from trust to vpn scr-1 dst-1 any permit
```

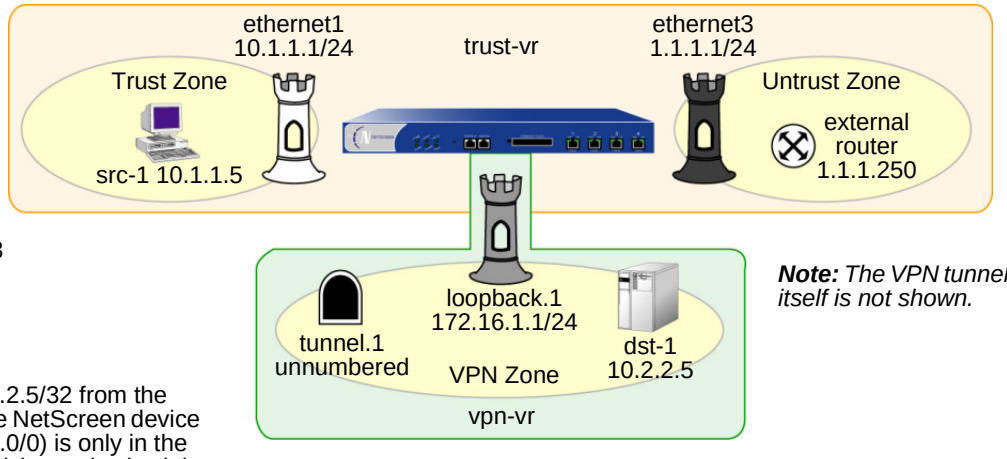

Note: *The VPN tunnel itself is not shown.*

The NetScreen device sends traffic destined for 10.2.2.5/32 from the trust-vr to the vpn-vr. If tunnel.1 becomes disabled, the NetScreen device drops the packet. Because the default route (to 0.0.0.0/0) is only in the trust-vr, the Netscreen device does not attempt to send the packet in plain text out ethernet3.

Putting all the tunnel interfaces in such a zone is very secure because there is no chance for the failure of a VPN, which causes the route to the associated tunnel interface to become inactive, to redirect traffic intended for tunneling to use a non-tunneled route—such as the default route. (For two suggestions about how to avoid such a problem, see “Security Consideration for a Route-Based VPN Design” on page **5**-323.)

You can also bind a tunnel interface to a tunnel zone. When you do, it must have an IP address. The purpose of binding a tunnel interface to a tunnel zone is to make NAT services available for policy-based VPN tunnels[1].

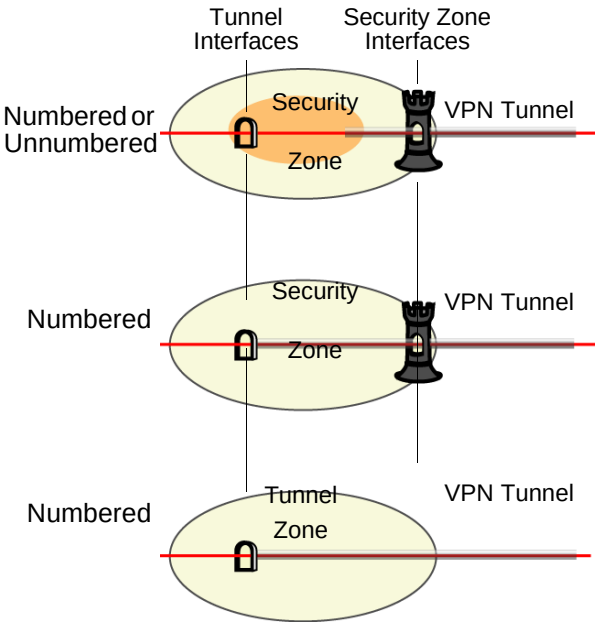

When a tunnel interface is in a security zone, you must bind a VPN tunnel to the tunnel interface. Doing so allows you to create a route-based VPN configuration.

The tunnel interface can be numbered or unnumbered. If it is unnumbered, the tunnel interface borrows the IP address from the default interface of the security zone in which you created it. Note: Only a tunnel interface with an IP address and netmask can support policy-based NAT.

When a numbered tunnel interface is in a security zone and is the only interface in that zone, you do not need to create a security zone interface. In this case, the security zone supports VPN traffic via the tunnel interface, but no other kind of traffic.

When a tunnel interface is bound to a tunnel zone, the tunnel interface must have an IP address and netmask. This allows you to define DIP pools and MIP addresses on that interface. If you bind a VPN tunnel to a tunnel zone, you cannot also bind it to a tunnel interface. In such cases, you must create a policy-based VPN configuration.

Conceptually, you can view VPN tunnels as pipes that you have laid. They extend from the local device to remote gateways, and the tunnel interfaces are the openings to these pipes. The pipes are always there, available for use whenever the routing engine directs traffic to one of their interfaces.

Generally, assign an IP address to a tunnel interface if you want the interface to support one or more dynamic IP (DIP) pools for source address translation (NAT-src) and mapped IP (MIP) addresses for destination address translation (NAT-dst). For more information about VPNs and address translation, see “VPN Sites with Overlapping Addresses” on page **5**-168. You can create a tunnel interface with an IP address and netmask in either a security zone or a tunnel zone.

1. Network address translation (NAT) services include dynamic IP (DIP) pools and mapped IP (MIP) addresses defined in the same subnet as an interface.

If the tunnel interface does not need to support address translation, and your configuration does not require the tunnel interface to be bound to a tunnel zone, you can specify the interface as unnumbered. You must bind an unnumbered tunnel interface to a security zone; you cannot bind it to a tunnel zone. You must also specify an interface with an IP address that is in the same virtual routing domain as the security zone to which the unnumbered interface is bound. The unnumbered tunnel interface borrows the IP address from that interface.

Note: *For examples showing how to bind a tunnel interface to a tunnel, see the route-based VPN examples in "Site-to-Site VPNs" on page* ***5****-69 and "Dialup VPNs" on page* ***5****-199.*

Deleting Tunnel Interfaces

You cannot immediately delete a tunnel interface that hosts mapped IP addresses (MIPs) or Dynamic IP (DIP) address pools. Before you delete a tunnel interface hosting any of these features, you must first delete any policies that reference them. Then you must delete the MIPs and DIP pools on the tunnel interface. Also, if a route-based VPN configuration references a tunnel interface, you must first delete the VPN configuration before you can delete the tunnel interface.

Example: Deleting a Tunnel Interface

In this example, tunnel interface tunnel.2 is linked to DIP pool 8. DIP pool 8 is referenced in a policy (ID 10) for VPN traffic from the Trust zone to the Untrust zone through a VPN tunnel named vpn1. To remove the tunnel interface, you must first delete the policy (or remove the reference to DIP pool 8 from the policy), and then the DIP pool. Then, you must unbind tunnel.2 from vpn1. After removing all the configurations that depend on the tunnel interface, you can then delete it.

WebUI

1. Deleting Policy 10, Which References DIP Pool 8

   Policies (From: Trust, To: Untrust): Click **Remove** for Policy ID 10.

2. Deleting DIP Pool 8, Which Is Linked to Tunnel.2

   Network > Interfaces > Edit (for tunnel.2) > DIP: Click **Remove** for DIP ID 8.

3. Unbinding tunnel.2 from vpn1

   VPNs > AutoKey IKE > Edit (for vpn1) > Advanced: Select **None** in the Bind to: Tunnel Interface drop-down list, click **Return**, and then click **OK**.

4. Deleting Tunnel.2

   Network > Interfaces: Click **Remove** for tunnel.2.

CLI

1. Deleting Policy 10, Which References DIP Pool 8

   ```
   unset policy 10
   ```

2. Deleting DIP Pool 8, Which Is Linked to Tunnel.2

   ```
   unset interface tunnel.2 dip 8
   ```

3. Unbinding tunnel.2 from vpn1

   ```
   unset vpn vpn1 bind interface
   ```

4. Deleting Tunnel.2

   ```
   unset interface tunnel.2
   save
   ```

VIEWING INTERFACES

You can view a table that lists all interfaces on your NetScreen device. Because they are predefined, physical interfaces are listed regardless of whether or not you configure them. Subinterfaces and tunnel interfaces are only listed once you create and configure them.

To view the interface table in the WebUI, click **Network > Interfaces**. You can specify the types of interfaces to display from the List Interfaces drop-down list.

To view the interface table in the CLI, use the **get interface** command.

Interface Table

The interface table displays the following information on each interface:

- **Name:** This field identifies the name of the interface.
- **IP/Netmask:** This field identifies the IP address and netmask address of the interface.
- **Zone:** This field identifies the zone to which the interface is bound.
- **Type:** This field indicates if the interface type: Layer 2, Layer 3, tunnel, redundant, aggregate, VSI.
- **Link:** This field identifies whether the interface is active (Up) or inactive (Down).
- **Configure:** This field allows you modify or remove interfaces.

WebUI Interface Table

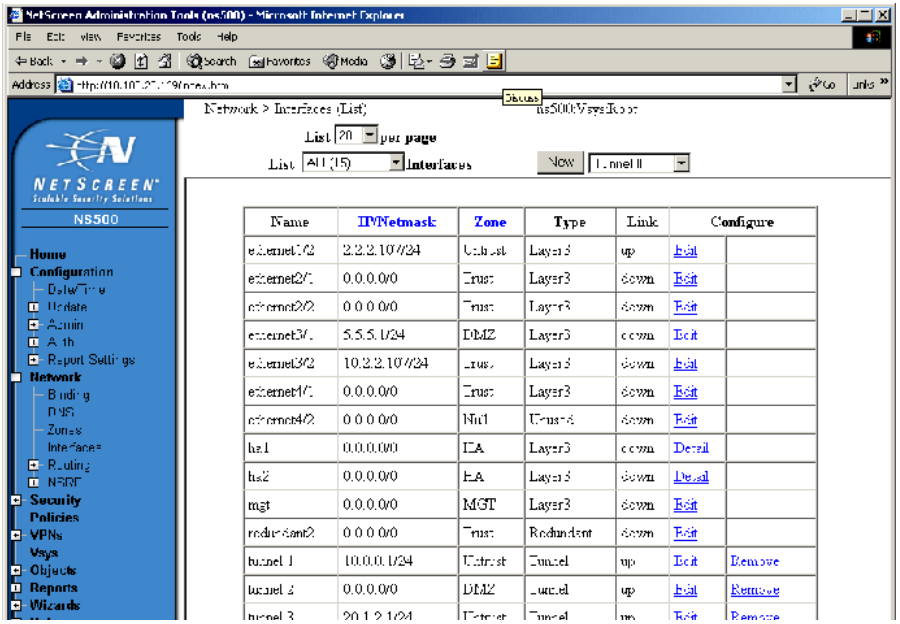

CLI Interface Table

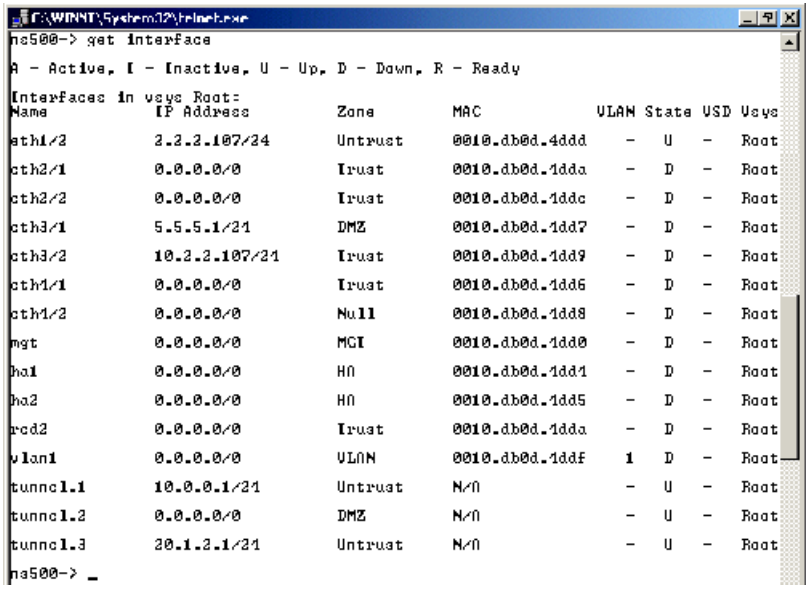

CONFIGURING SECURITY ZONE INTERFACES

This section describes how to configure the following aspects of security zone interfaces:

- Binding and unbinding an interface to a security zone
- Assigning an address to a Layer 3 (L3) security zone interface
- Modifying physical interfaces and subinterfaces
- Creating subinterfaces
- Deleting subinterfaces

> ***Note:*** *For information on setting traffic bandwidth for an interface, see Chapter 10, "Traffic Shaping". For more information on the management and other services options available per interface, see "Controlling Administrative Traffic" on page* ***3****-29.*

Binding an Interface to a Security Zone

You can bind any physical interface to either a L2 or L3 security zone. You can bind a subinterface only to a L3 security zone because a subinterface requires an IP address. You can only assign an IP address to an interface after you have bound it to a L3 security zone.

Example: Binding an Interface

In this example, you bind ethernet5 to the Trust zone.

WebUI

Network > Interfaces > Edit (for ethernet5): Select **Trust** from the Zone Name drop-down list, and then click **OK**.

CLI

```
set interface ethernet5 zone trust
save
```

Defining an Address for a L3 Security Zone Interface

When defining a Layer 3 (L3) security zone interface or subinterface, you must assign it an IP address and netmask. If you bind the interface to a zone in the trust-vr, you can also specify the interface mode as NAT or Route. (If the zone to which you bind the interface is in the untrust-vr, the interface is always in Route mode.)

Note: *For examples of NAT and Route mode configurations, see* *Chapter 5, "Interface Modes" on page 91**.*

The two basic types of IP addresses to be considered when making interface address assignments are as follows:

- Public addresses, which Internet service providers (ISPs) supply for use on a public network like the Internet and which must be unique
- Private addresses, which a local network administrator assigns for use on a private network and which other administrators can assign for use on other private networks too

Note: *When you add an IP address to an interface, the NetScreen device checks via an ARP request to make sure that the IP address does not already exist on the local network. (The physical link must be up at the time.) If the IP address already exists, a warning is displayed.*

Public IP Addresses

If an interface connects to a public network, it must have a public IP address. Also, if a L3 security zone in the untrust-vr connects to a public network and the interfaces of zones in the trust-vr are in Route mode, then all the addresses in the zones in the trust-vr—for interfaces and for hosts—must also be public addresses. Public IP addresses fall into three classes, A, B, and C[2], as shown below:

Address Class	Address Range	Excluded Address Range
A	0.0.0.0 – 127.255.255.255	10.0.0.0 – 10.255.255.255, 127.0.0.0 – 127.255.255.255
B	128.0.0.0 – 191.255.255.255	172.16.0.0 – 172.31.255.255
C	192.0.0.0 – 223.255.255.255	192.168.0.0 – 192.168.255.255

2. There are also D and E class addresses, which are reserved for special purposes.

An IP address is composed of four octets, each octet being 8 bits long. In a class A address, the first 8 bits indicate the network ID, and the final 24 bits indicate the host ID (nnn.hhh.hhh.hhh). In a class B address, the first 16 bits indicate the network ID, and the final 16 bits indicate the host ID (nnn.nnn.hhh.hhh). In a class C address, the first 24 bits indicate the network ID, and the the final 8 bits indicate the host ID (nnn.nnn.nnn.hhh).

Through the application of subnet masks (or netmasks), you can further divide networks. A netmask essentially masks part of the host ID so that the masked part becomes a subnet of the network ID. For example, the 24-bit mask[3] in the address 10.2.3.4/24 indicates that the first 8 bits (that is, the first octet—010) identify the network portion of this private class A address, the next 16 bits (that is, the second and third octets—002.003) identify the subnetwork portion of the address, and the last 8 bits (the last octet—004) identify the host portion of the address. Using subnets to narrow large network address spaces into smaller subdivisions greatly increases the efficient delivery of IP datagrams.

Private IP Addresses

If an interface connects to a private network, a local network administrator can assign it any address, although it is conventional to use an address from the range of addresses reserved for private use—10.0.0.0/8, 172.16.0.0 – 172.31.255.255, 192.168.0.0/16— as defined in RFC 1918, “Address Allocation for Private Internets”.

If a L3 security zone in the untrust-vr connects to a public network and the interfaces bound to zones in the trust-vr are in NAT mode, then all the addresses in the zones in the trust-vr—for interfaces and for hosts—can be private addresses.

3. The dotted-decimal equivalent of a 24-bit mask is 255.255.255.0.

Example: Addressing an Interface

In this example, you assign ethernet5 the IP address 210.1.1.1/24 and give it the Manage IP address 210.1.1.5. (Note that the Manage IP address must be in the same subnet as the security zone interface IP address.) Finally, you set the interface in NAT mode, which translates all internal IP addresses to the default interfaces[4] bound to the other security zones.

WebUI

Network > Interfaces > Edit (for ethernet5): Enter the following, and then click **OK**:

IP Address/Netmask: 210.1.1.1/24

Manage IP: 210.1.1.5

CLI

```
set interface ethernet5 ip 210.1.1.1/24
set interface ethernet5 manage-ip 210.1.1.5
save
```

4. The default interface in a security zone is the first interface bound to the zone. To learn which interface is the default interface for a zone, see the Default IF column on the Network > Zones page in the WebUI, or the Default-If column in the output from the **get zone** command in the CLI.

Unbinding an Interface from a Security Zone

If an interface is unnumbered, you can unbind it from one security zone and bind it to another. If an interface is numbered, you must first set its IP address and netmask to 0.0.0.0. Then, you can unbind it from one security zone and bind it to another one, and (optionally) reassign it an IP address/netmask.

Example: Unbinding an Interface

In this example, ethernet3 has the IP address 210.1.1.1/24 and is bound to the Untrust zone. You set its IP address and netmask to 0.0.0.0/0 and bind it to the Null zone.

WebUI

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Null

IP Address/Netmask: 0.0.0.0/0

CLI

```
set interface ethernet3 ip 0.0.0.0/0
set interface ethernet3 zone null
save
```

Modifying Interfaces

After you have configured a physical interface, a subinterface, a redundant interface, an aggregate interface, or a Virtual Security Interface (VSI), you can later change any of the following settings should the need arise:

- IP address and netmask
- Manage IP address
- (L3 zone interfaces) Management and network services
- (Subinterface) Subinterface ID number and VLAN tag number
- (Interfaces bound to L3 security zones in the trust-vr) Interface mode—NAT or Route
- (Physical interface) Traffic bandwidth settings (see Chapter 10, “Traffic Shaping” on page 477)
- (Physical, redundant, and aggregate interfaces) Maximum Transmission Unit (MTU) size
- (L3 interfaces) Block traffic from coming in and going out the same interface, including traffic between a primary and secondary subnet or between secondary subnets (this is done with the CLI **set interface** command with the **route-deny** option)

For physical interfaces on some NetScreen devices, you can force the physical state of the link to be down or up. By forcing the physical state of the link to be down, you can simulate a disconnect of the cable from the interface port. (This is done with the CLI **set interface** command with the **phy link-down** option.)

Example: Modifying Settings on an Interface

In this example, you make some modifications to ethernet1, an interface bound to the Trust zone. You change the Manage IP address from 10.1.1.2 to 10.1.1.12. To enforce tighter security of administrative traffic, you also change the management services options, enabling SCS and SSL and disabling Telnet and WebUI.

WebUI

Network > Interfaces > Edit (for ethernet1): Make the following modifications, and then click **OK**:

Manage IP: 10.1.1.12

Management Services: (select) SSH, SSL; (clear) Telnet, WebUI

CLI

```
set interface ethernet1 manage-ip 10.1.1.12
set interface ethernet1 manage ssh
set interface ethernet1 manage ssl
unset interface ethernet1 manage telnet
unset interface ethernet1 manage web
save
```

Creating Subinterfaces

You can create a subinterface on any physical interface[5] in the root system or virtual system. A subinterface makes use of VLAN tagging to distinguish traffic bound for it from traffic bound for other interfaces. Note that although a subinterface stems from a physical interface, from which it borrows the bandwidth it needs, you can bind a subinterface to any zone, not necessarily that to which its “parent” interface is bound. Additionally, the IP address of a subinterface must be in a different subnet from the IP addresses of all other physical interfaces and subinterfaces.

Example: Creating a Subinterface in the Root System

In this example, you create a subinterface for the Trust zone in the root system. You configure the subinterface on ethernet1, which is bound to the Trust zone. You bind the subinterface to a user-defined zone named “accounting”, which is in the trust-vr. You assign it subinterface ID 3, IP address 10.2.1.1/24, and VLAN tag ID 3. The interface mode is NAT.

WebUI

Network > Interfaces > New Sub-IF: Enter the following, and then click **OK**:

Interface Name: ethernet1.3

Zone Name: accounting

IP Address / Netmask: 10.2.1.1/24

VLAN Tag: 3

5. You can also configure subinterfaces on redundant interfaces and VSIs. For an example that includes the configuration of a subinterface on a redundant interface, see “Virtual System Failover” on page **8**-144.

CLI

```
set interface ethernet1.3 zone accounting
set interface ethernet1.3 ip 10.2.1.1/24 tag 3
save
```

Deleting Subinterfaces

You cannot immediately delete a subinterface that hosts mapped IP addresses (MIPs), virtual IP addresses (VIPs), or Dynamic IP (DIP) address pools. Before you delete a subinterface hosting any of these features, you must first delete any policies or IKE gateways that reference them. Then you must delete the MIPs, VIPs, and DIP pools on the subinterface.

Example: Deleting a Security Zone Interface

In this example, you delete the subinterface ethernet1:1.

WebUI

Network > Interfaces: Click **Remove** for ethernet1:1.

A system message prompts you to confirm the removal.

Click **Yes** to delete the subinterface.

CLI

```
unset interface ethernet1:1
save
```

SECONDARY IP ADDRESSES

Each NetScreen interface has a single, unique *primary* IP address. However, some situations demand that an interface have multiple IP addresses. For example, an organization might have additional IP address assignments and might not wish to add a router to accommodate them. In addition, an organization might have more network devices than its subnet can handle, as when there are more than 254 hosts connected to a LAN. To solve such problems, you can add *secondary* IP addresses to an interface in the Trust, DMZ, or user-defined zone.

Note: *You cannot set multiple secondary IP addresses for interfaces in the Untrust zone.*

Secondary IP Address Properties

Secondary addresses have certain properties that affect how you can implement such addresses. These properties are as follows:

- There can be no subnet address overlap between any two secondary IP addresses. In addition, there can be no subnet address overlap between a secondary IP and any existing subnet on the NetScreen device.
- When you manage a NetScreen device through a secondary IP address, the address always has the same management properties as the primary IP address. Consequently, you cannot specify a separate management configuration for the secondary IP address.
- You cannot configure a gateway for a secondary IP address.
- Whenever you create a new secondary IP address, the NetScreen device automatically creates a corresponding routing table entry. When you delete a secondary IP address, the device automatically deletes its routing table entry.

Enabling or disabling routing between two secondary IP addresses causes no change in the routing table. For example, if you disable routing between two such addresses, the NetScreen device drops any packets directed from one interface to the other, but no change occurs in the routing table.

Example: Creating a Secondary IP Address

In this example, you set up a secondary IP address—192.168.2.1/24—for ethernet1, an interface that has IP address 10.1.1.1/24 and is bound to the Trust zone.

WebUI

Network > Interfaces > Edit (for ethernet1) > Secondary IP: Enter the following, and then click **Add**:

IP Address/Netmask: 192.168.2.1/24

CLI

```
set interface ethernet1 ip 192.168.2.1/24 secondary
save
```

LOOPBACK INTERFACES

A loopback interface is a logical interface that emulates a physical interface on the NetScreen device. However, unlike a physical interface, a loopback interface is always in the up state as long as the device on which it resides is up. Loopback interfaces are named loopback.*id_num*, where *id_num* is a number greater than or equal to 1[6] and denotes a unique loopback interface on the device. Like a physical interface, you must assign an IP address to a loopback interface and bind it to a security zone.

After defining a loopback interface, you can then define other interfaces as members of its group. Traffic can reach a loopback interface if it arrives through one of the interfaces in its group. Any interface type can be a member of a loopback interface group—physical interface, subinterface, tunnel interface, redundant interface, or VSI.

Example: Creating a Loopback Interface

In the following example, you create the loopback interface loopback.1, bind it to the Untrust zone, and assign the IP address 1.1.1.27/24 to it.

WebUI

Network > Interfaces > New Loopback IF: Enter the following, and then click **OK**:

Interface Name: loopback.1

Zone: Untrust (select)

IP Address / Netmask: 1.1.1.27./24

CLI

```
set interface loopback.1 zone untrust
set interface loopback.1 ip 1.1.1.27
save
```

Note: *The loopback interface is not directly accessible from networks or hosts that reside in other zones. You must define a policy to permit traffic to and from the interface.*

6. The maximum *id_num* value you can specify is platform-specific.

Using Loopback Interfaces

You can use a loopback interface in many of the same ways as a physical interface. This section shows examples of the ways you can configure loopback interfaces.

Note: *You cannot bind a loopback interface to a HA zone, nor can you configure a loopback interface for layer 2 operation or as a redundant/aggregate interface. You cannot configure the following features on loopback interfaces: NTP, DNS, VIP, secondary IP, track IP, or Webauth.*

You can define a MIP on a loopback interface. This allows the MIP to be accessed by a group of interfaces; this capability is unique to loopback interfaces. For information about using the loopback interface with MIPs, see "MIP and the Loopback Interface" on page 346.

You can manage the NetScreen device using either the IP address of a loopback interface or the manage IP address that you assign to a loopback interface.

Example: Using the Loopback Interface to Manage a Device

In the following example, you configure the previously-defined loopback.1 interface as a management interface for the device.

WebUI

Network > Interfaces > loopback.1 > Edit: Select all the management options, and then click **OK**.

CLI

```
set interface loopback.1 manage
save
```

Example: Enabling BGP on a Loopback Interface

The loopback interface can support the BGP dynamic routing protocol on the NetScreen device. In the following example, you enable BGP on the loopback.1 interface.

Note: *To enable BGP on the loopback interface, you must first create a BGP instance for the virtual router in which you plan to bind the interface. For information about configuring BGP on NetScreen devices, See Volume 6, "Dynamic Routing".*

WebUI

Network > Interfaces > loopback.1 > Edit: Select **Protocol BGP**, and then click **OK**.

CLI

```
set interface loopback.1 protocol bgp
save
```

Example: Configuring NSRP VSIs on a Loopback Interface

You can configure Virtual Security Interfaces (VSIs) for NSRP on a loopback interface. The physical state of the VSI on the loopback interface is always up. The interface can be active or not, depending upon the state of the VSD group to which the interface belongs.

WebUI

Network > Interfaces > New VSI IF: Enter the following, and then click **OK**:

Interface Name: VSI Base: loopback.1

VSD Group: 1

IP Address / Netmask: 1.1.1.1/24

CLI

```
set interface loopback.1:1 ip 1.1.1.1/24
save
```

Example: Specifying a Loopback Interface as a Source Interface

You can use a loopback interface as a source interface for certain traffic that originates from the NetScreen device. (When you define a source interface for an application, the specified source interface address is used instead of the outbound interface address to communicate with an external device.) In the following example, you specify that the NetScreen device uses the previously-defined loopback.1 interface for sending syslog packets.

WebUI

Configuration > Report Settings > Syslog: Enter the following, and then click **Apply**:

Enable Syslog Messages: (select)

Source interface: loopback.1 (select)

Syslog Servers:

No.: 1 (select)

IP/Hostname: 10.1.1.1

Traffic Log: (select)

Event Log: (select)

CLI

```
set syslog config 10.1.1.1 log all
set syslog src-interface loopback.1
set syslog enable
save
```

5

Interface Modes

Interfaces can operate in three different modes: Network Address Translation (NAT), Route, and Transparent. If an interface bound to a Layer 3 zone has an IP address, you can define the operational mode for that interface as either NAT[1] or Route. An interface bound to a Layer 2 zone (such as the predefined v1-trust, v1-untrust, and v1-dmz zones, or a user-defined Layer 2 zone) must be in Transparent mode. You select an operational mode when you configure an interface.

This chapter contains the following sections:

- "Transparent Mode" on page 92
  - "Zone Settings" on page 93
  - "Traffic Forwarding" on page 94
  - "Unknown Unicast Options" on page 95
- "NAT Mode" on page 110
  - "Inbound and Outbound NAT Traffic" on page 112
  - "Interface Settings" on page 113
- "Route Mode" on page 118
  - "Interface Settings" on page 119

1. Although you can define the operational mode for an interface bound to any Layer 3 zone as NAT, the NetScreen device only performs NAT on traffic passing through that interface en route to the Untrust zone. NetScreen does not perform NAT on traffic destined for any zone other than the Untrust zone. Also, note that NetScreen allows you to set an Untrust zone interface in NAT mode, but doing so activates no NAT operations.

TRANSPARENT MODE

When an interface is in Transparent mode, the NetScreen device filters packets traversing the firewall without modifying any of the source or destination information in the IP packet header. All interfaces behave as though they are part of the same network, with the NetScreen device acting much like a Layer 2 switch or bridge. In Transparent mode, the IP addresses of interfaces are set at 0.0.0.0, making the presence of the NetScreen device invisible, or "transparent," to users.

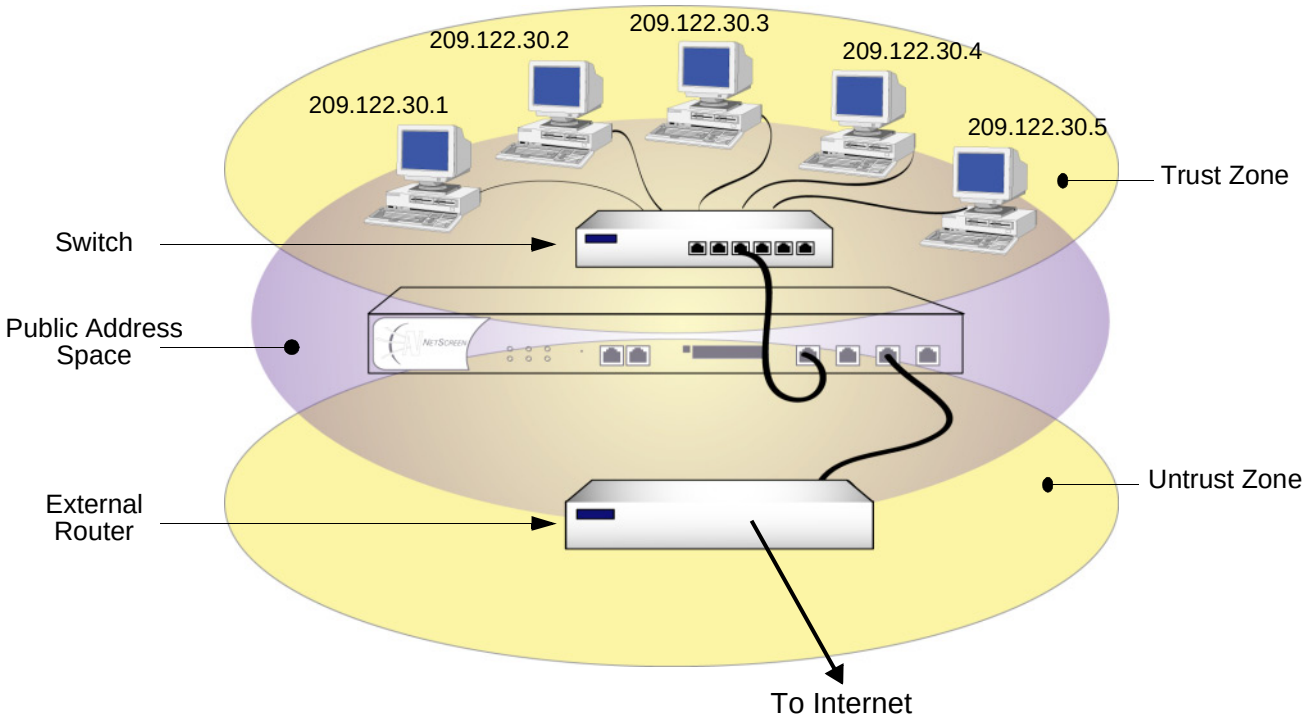

Transparent mode is a convenient means for protecting Web servers, or any other kind of server that mainly receives traffic from untrusted sources. Using Transparent mode offers the following benefits:

- No need to reconfigure the IP settings of routers or protected servers
- No need to create Mapped or Virtual IP addresses for incoming traffic to reach protected servers

Zone Settings

By default, ScreenOS creates one function zone, the VLAN zone, and three L2 security zones: V1-Trust, V1-Untrust, and V1-DMZ.

VLAN Zone

The VLAN zone hosts the VLAN1 interface, which has the same configuration and management abilities as a physical interface. When the NetScreen device is in Transparent mode, you use the VLAN1 interface for managing the device and terminating VPN traffic. You can configure the VLAN1 interface to permit hosts in the L2 security zones to manage the device. To do that, you must set the VLAN1 interface IP address in the same subnet as the hosts in the L2 security zones.

For management traffic, the VLAN1 Manage IP takes precedence over the VLAN1 interface IP. You can set the VLAN1 Manage IP for management traffic and dedicate the VLAN1 interface IP solely for VPN tunnel termination.

Predefined Layer 2 Zones

ScreenOS provides three L2 security zones by default: V1-Trust, V1-Untrust, and V1-DMZ. These three zones share the same L2 domain. When you configure an interface in one of the zones, it gets added to the L2 domain shared by all interfaces in all the L2 zones. All hosts in the L2 zones must be on the same subnet to communicate.

As stated in the previous section, when the device is in transparent mode, you use the VLAN1 interface to manage the device. For management traffic to reach the VLAN1 interface, you must enable the management options on the VLAN1 interface and on the zone(s) through which the management traffic passes. By default, all management options are enabled in the V1-Trust zone. To enable hosts in other zones to manage the device, you must set those options on the zones to which they belong.

Note: *To see which physical interfaces are prebound to the L2 zones for each NetScreen platform, refer to the installer's guide for that platform.*

Traffic Forwarding

A NetScreen device operating at Layer 2 (L2) does not permit any traffic between zones unless there is a policy configured on the device. For more information on how to set policies, see "Policies" on page 197. After you configure a policy on the NetScreen device, it does the following:

- Allows or denies the traffic specified in the policy
- Allows ARP and L2 non-IP multicast and broadcast traffic. The NetScreen device can then receive and pass L2 broadcast traffic for the spanning tree protocol.
- Continues to block all non-IP and non-ARP unicast traffic, and IPSec traffic

You can change the forwarding behavior of the device as follows:

- To block all L2 non-IP and non-ARP traffic, including multicast and broadcast traffic, enter the **unset interface vlan1 bypass-non-ip-all** command.
- To allow all L2 non-IP traffic to pass through the device, enter the **set interface vlan1 bypass-non-ip** command.
- To revert to the default behavior of the device, which is to block all non-IP and non-ARP unicast traffic, enter the **unset interface vlan1-bypass-non-ip** command.
  - Note that the **unset interface vlan1 bypass-non-ip-all** command always overwrites the **unset interface vlan1 bypass-non-ip** command when both commands are in the configuration file. Therefore, if you had previously entered the **unset interface vlan1 bypass-non-ip-all** command, and you now want the device to revert to its default behavior of blocking only the non-IP and non-ARP unicast traffic, you should first enter the **set interface vlan1 bypass-non-ip** command to allow all non-IP traffic to pass through the device. Then you must enter the **unset interface vlan1-bypass-non-ip** command to block only the non-IP, non-ARP unicast traffic.
- To allow a NetScreen device to pass IPSec traffic without attempting to terminate it, use the **set interface vlan1 bypass-others-ipsec** command. The NetScreen device then allows the IPSec traffic to pass through to other VPN termination points.

Note: *A NetScreen device with interfaces in Transparent mode requires routes for two purposes: to direct self-initiated traffic, such as SNMP traps, and to forward VPN traffic after encapsulating or decapsulating it.*

Unknown Unicast Options

When a host or any kind of network device does not know the MAC address associated with the IP address of another device, it uses the Address Resolution Protocol (ARP) to obtain it. The requestor broadcasts an ARP query (arp-q) to all the other devices on the same subnet. The arp-q requests the device at the specified destination IP address to send back an ARP reply (arp-r), which provides the requestor with the MAC address of the replier. When all the other devices on the subnet receive the arp-q, they check the destination IP address and, because it is not their IP address, drop the packet. Only the device with the specified IP address returns an arp-r. After a device matches an IP address with a MAC address, it stores the information in its ARP cache.

As ARP traffic passes through a NetScreen device in Transparent mode, the device notes the source MAC address in each packet and learns which interface leads to that MAC address. In fact, the NetScreen device learns which interface leads to which MAC address by noting the source MAC addresses in all packets it receives. It then stores this information in its forwarding table.

Note: *A NetScreen device in Transparent mode does not permit any traffic between zones unless there is a policy configured on the device. For more information on how the device forwards traffic when it is in Transparent mode, see* *"Traffic Forwarding" on page 94**.*

The situation can arise when a device sends a unicast packet with a destination MAC address, which it has in its ARP cache, but which the NetScreen device does not have in its forwarding table. For example, the NetScreen device clears its forwarding table every time it reboots. (You can also clear the forwarding table with the CLI command **clear arp**.) When a NetScreen device in Transparent mode receives a unicast packet for which it has no entry in its forwarding table, it can follow one of two courses:

- After doing a policy lookup to determine the zones to which traffic from the source address is permitted, flood the initial packet out the interfaces bound to those zones, and then continue using whichever interface receives a reply. This is the Flood option, which is enabled by default.
- Drop the initial packet, flood ARP queries (and, optionally, trace-route packets, which are ICMP echo requests with the time-to-live value set to 1) out all interfaces (except the interface at which the packet arrived), and then send subsequent packets through whichever interface receives an ARP (or trace-route) reply from the router or host whose MAC address matches the destination MAC address in the initial packet. The trace-route option allows the NetScreen device to discover the destination MAC address when the destination IP address is in a nonadjacent subnet.

> ***Note:*** *Of the two methods—flood and ARP/trace-route—ARP/trace-route is more secure because the NetScreen device floods ARP queries and trace-route packets—not the initial packet—out all interfaces.*

Flood Method

The flood method forwards packets in the same manner as most Layer 2 switches. A switch maintains a forwarding table that contains MAC addresses and associated ports for each Layer 2 domain. The table also contains the corresponding interface through which the switch can forward traffic to each device. Every time a packet arrives with a new source MAC address in its frame header, the switch adds the MAC address to its forwarding table. It also tracks the interface at which the packet arrived. If the destination MAC address is unknown to the switch, the switch duplicates the packet and floods it out all interfaces (other than the interface at which the packet arrived). It learns the previously unknown MAC address and its corresponding interface when a reply with that MAC address arrives at one of its interfaces.

When you enable the flood method and the NetScreen device receives an ethernet frame with a destination MAC address that is not listed in the NetScreen device MAC table, it floods the packet out all interfaces.

Flood Method

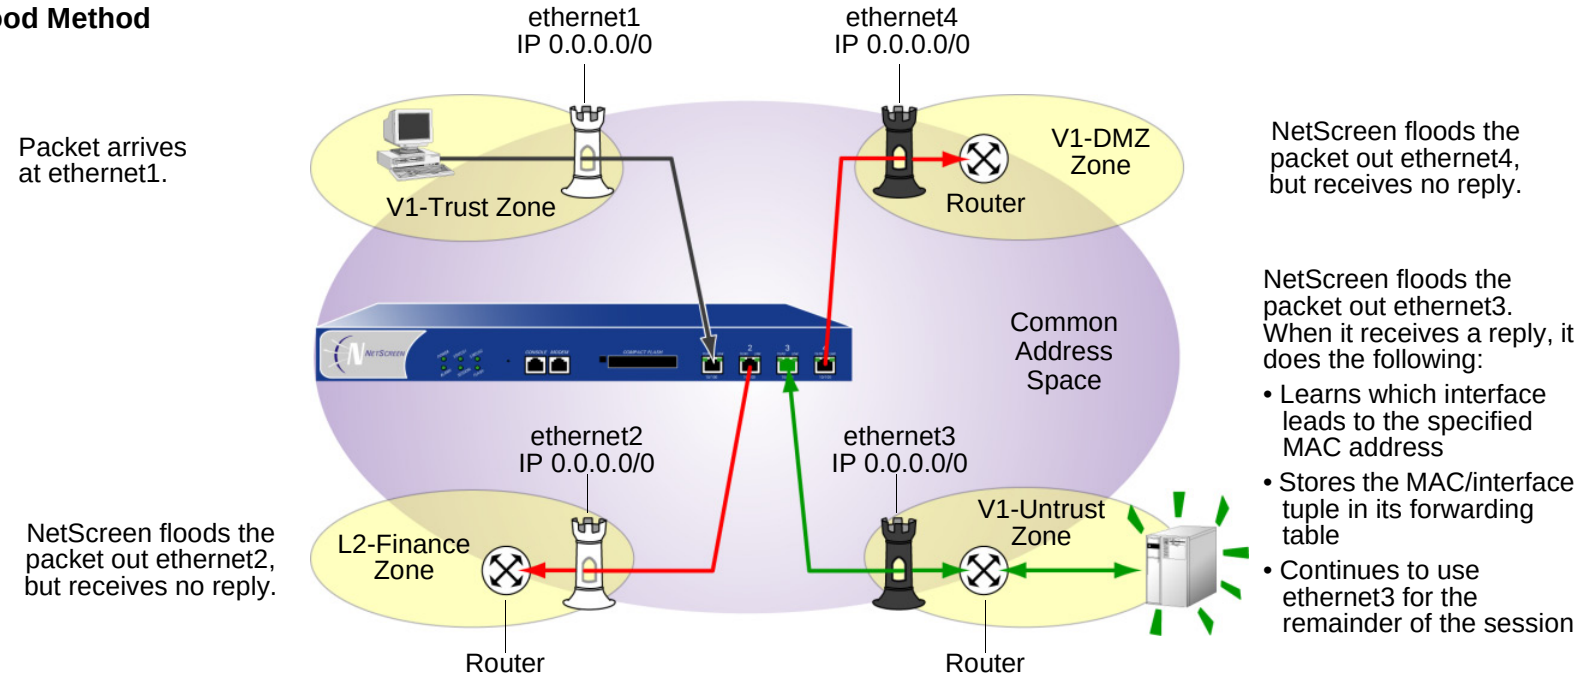

To enable the flood method for handling unknown unicast packets, do either of the following:

WebUI

Network > Interface > Edit (for VLAN1): For the broadcast options, select **Flood**, and then click **OK**.

CLI

```
set interface vlan1 broadcast flood
save
```

ARP/Trace-Route Method

When you enable the ARP method with the trace-route option[2] and the NetScreen device receives an ethernet frame with a destination MAC address that is not listed in its MAC table, the NetScreen device performs the following series of actions:

1. The NetScreen device notes the destination MAC address in the initial packet (and, if it is not already there, adds the source MAC address and its corresponding interface to its forwarding table).
2. The NetScreen device drops the initial packet.
3. The NetScreen device generates two packets—ARP query (arp-q) and a trace-route (an ICMP echo request, or PING) with a time-to-live (TTL) field of 1—and floods those packets out all interfaces except the interface at which the initial packet arrived. For the arp-q packets and ICMP echo requests, the NetScreen device uses the source and destination IP addresses from the initial packet. For arp-q packets, the NetScreen device replaces the source MAC address from the initial packet with the MAC address for VLAN1, and it replaces the destination MAC address from the initial packet with ffff.ffff.ffff. For the trace-route option, the NetScreen device uses the source and destination MAC addresses from the initial packet in the ICMP echo requests that it broadcasts.

   If the destination IP address belongs to a device in the same subnet as the ingress IP address[3], the host returns an ARP reply (arp-r) with its MAC address, thus indicating the interface through which the NetScreen device must forward traffic destined for that address. (See “ARP Method” on page 100.)

   If the destination IP address belongs to a device in a subnet beyond that of the ingress IP address, the trace-route returns the IP and MAC addresses of the router leading to the destination[4], and more significantly, indicates the interface through which the NetScreen device must forward traffic destined for that MAC address. (See “Trace-Route” on page 101.)

2. When you enable the ARP method, the trace-route option is enabled by default. You can also enable the ARP method without the trace-route option. However, this method only allows the NetScreen device to discover the destination MAC address for a unicast packet if the destination IP address is in the same subnet as the ingress IP address. (For more information about the ingress IP address, see the next footnote.)
3. The ingress IP address refers to the IP address of the last device to send the packet to the NetScreen device. This device might be the source that sent the packet or a router forwarding the packet.
4. Actually, the trace-route returns the IP and MAC addresses of all the routers in the subnet. The NetScreen device then matches the destination MAC address from the initial packet with the source MAC address on the arp-r packets to determine which router to target, and consequently, which interface to use to reach that target.

4. Combining the destination MAC address gleaned from the initial packet with the interface leading to that MAC address, the NetScreen device adds a new entry to its forwarding table.
5. The NetScreen device forwards all subsequent packets it receives out the correct interface to the destination.

To enable the ARP/trace-route method for handling unknown unicast packets, do either of the following:

WebUI

Network > Interface > Edit (for VLAN1): For the broadcast options, select **ARP**, and then click **OK**.

CLI

```
set interface vlan1 broadcast arp
save
```

Note: *The trace-route option is enabled by default. If you want to use ARP without the trace-route option, enter the following command:* ***unset interface vlan1 broadcast arp trace-route****. This command unsets the trace-route option but does not unset ARP as the method for handling unknown unicast packets.*

The following illustration shows how the ARP method can locate the destination MAC when the destination IP address is in an adjacent subnet.

ARP Method

Note: *Only the relevant elements of the packet header and the last four digits in the MAC addresses are shown below.*

If the following packet

Ethernet Frame			IP Datagram	
dst	src	type	src	dst
11bb	11aa	0800	210.1.1.5	210.1.1.75

arrives at ethernet1 and the forwarding table does not have an entry for MAC address 00bb.11bb.11bb, the NetScreen device floods the following arp-q packet out eth2, eth3, and eth4.

Ethernet Frame			ARP Message	
dst	src	type	src	dst
ffff	39ce	0806	210.1.1.5	210.1.1.75

When the NetScreen device receives the following arp-r at eth2,

Ethernet Frame			ARP Message	
dst	src	type	src	dst
39ce	11bb	0806	210.1.1.75	210.1.1.5

it can now associate the MAC address with the interface leading to it.

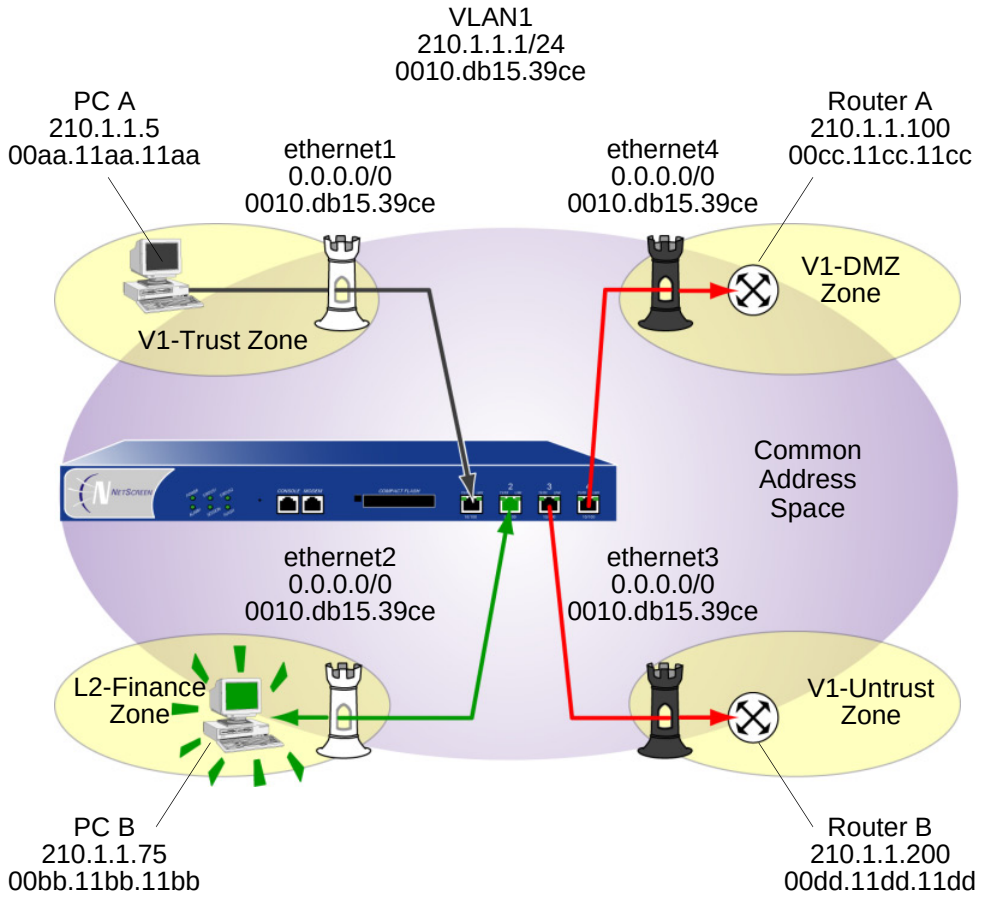

The following illustration shows how the trace-route option can locate the destination MAC when the destination IP address is in a nonadjacent subnet.

Trace-Route

Note: *Only the relevant elements of the packet header and the last four digits in the MAC addresses are shown below.*

If the following packet

Ethernet Frame			IP Datagram	
dst	src	type	src	dst
11dd	11aa	0800	210.1.1.5	195.1.1.5

arrives at ethernet1 and the forwarding table does not have an entry for MAC address 00dd.11dd.11dd, the NetScreen device floods the following trace-route packet out eth2, eth3, and eth4.

Ethernet Frame			ICMP Message		
dst	src	type	src	dst	TTL
11dd	11aa	0800	210.1.1.5	195.1.1.5	1

When the NetScreen device receives the following response at eth3,

Ethernet Frame			ICMP Message		
dst	src	type	src	dst	msg
11aa	11dd	0800	210.1.1.200	210.1.1.5	Time Exceeded

it can now associate the MAC address with the interface leading to it.

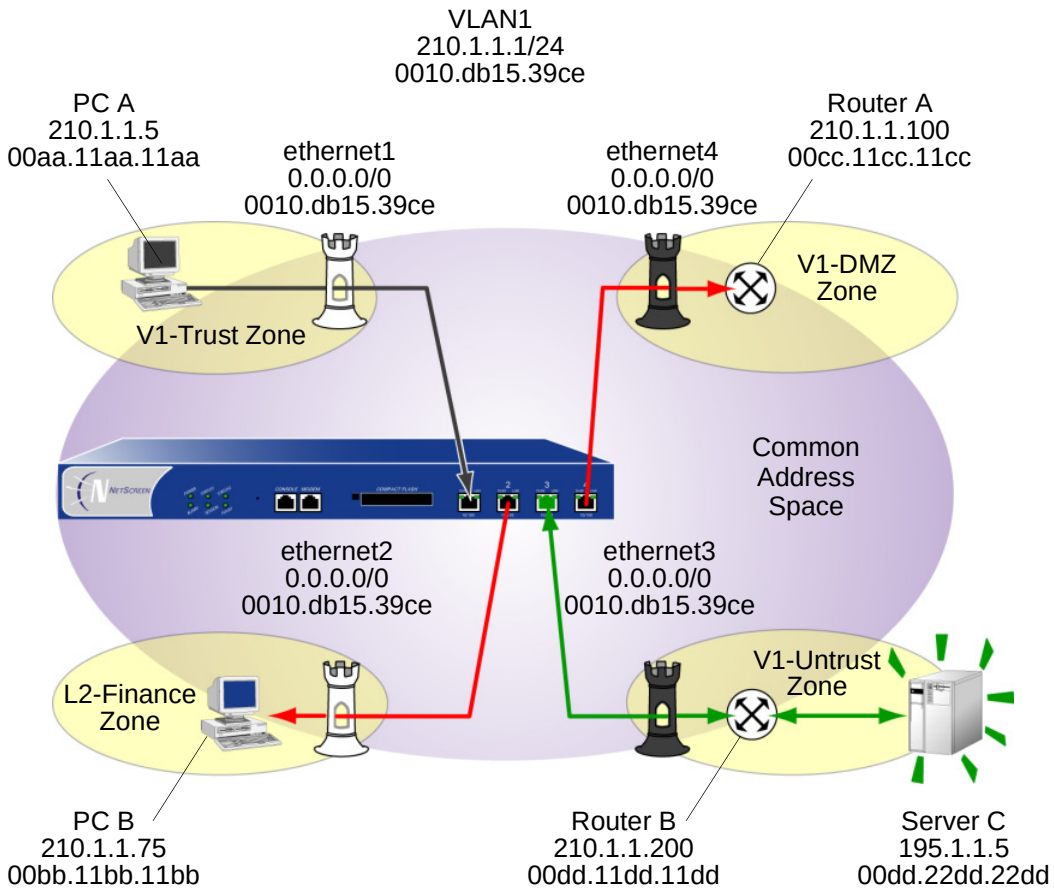

Example: VLAN1 Interface for Management

In this example, you configure the NetScreen device for management to its VLAN1 interface as follows:

- Assign the VLAN1 interface an IP address of 1.1.1.1/24.
- Enable Web, Telnet, SSH and Ping on both the VLAN1 interface and V1-Trust[5] security zone.

> ***Note:*** *To manage the device from a Layer 2 security zone, you must set the same management options for both the VLAN1 interface and the Layer 2 security zone.*

- Add a route in the trust virtual router (all Layer 2 security zones are in the trust-vr routing domain) to enable management traffic to flow between the NetScreen device and an administrative workstation beyond the immediate subnet of the NetScreen device. All security zones are in the trust-vr routing domain.

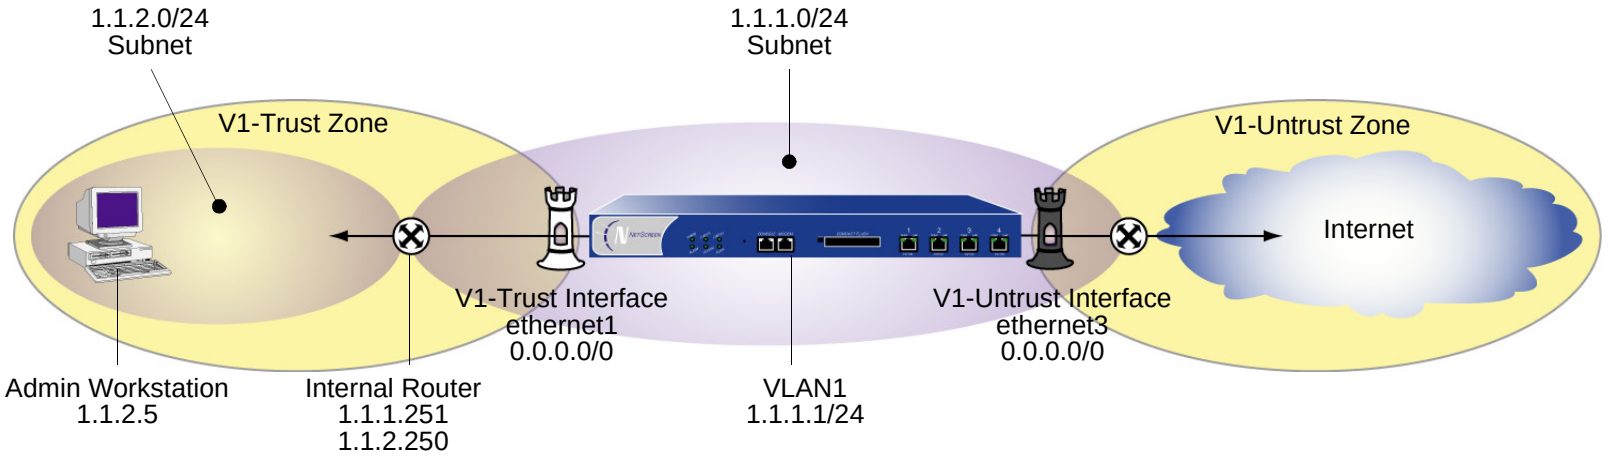

5. By default, NetScreen enables the management options for the VLAN1 interface and V1-Trust security zone. Enabling these options is included in this example for illustrative purposes only. Unless you have previously disabled them, you really do not need to enable them manually.

WebUI

1. VLAN1 Interface

Network > Interfaces > Edit (for VLAN1): Enter the following, and then click **OK**:

IP Address/Netmask: 1.1.1.1/24

Management Services: WebUI, Telnet, SSH (select)

Other Services: Ping (select)

2. V1-Trust Zone

Network > Zones > Edit (for V1-Trust): Select the following, and then click **OK**:

Management Services: WebUI, Telnet, SSH

Other Services: Ping

3. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 1.1.2.0/24

Gateway: (select)

Interface: vlan1(trust-vr)

Gateway IP Address: 1.1.1.251

Metric: 1

CLI

1. VLAN1 Interface

```
set interface vlan1 ip 1.1.1.1/24
set interface vlan1 manage web
set interface vlan1 manage telnet
set interface vlan1 manage ssh
set interface vlan1 manage ping
```

2. V1-Trust Zone

```
set zone v1-trust manage web
set zone v1-trust manage telnet
set zone v1-trust manage ssh
set zone v1-trust manage ping
```

3. Route

```
set vrouter trust-vr route 1.1.2.0/24 interface vlan1 gateway 1.1.1.251 metric 1
save
```

Example: Transparent Mode

The following example illustrates a basic configuration for a single LAN protected by a NetScreen device in Transparent mode. Policies permit outgoing traffic for all hosts in the V1-Trust zone, incoming SMTP services for the mail server, and incoming FTP-GET services for the FTP server.

To increase the security of management traffic, you change the HTTP port number for WebUI management from 80 to 5555, and the Telnet port number for CLI management from 23 to 4646. You use the VLAN1 IP address—1.1.1.1/24—to manage the NetScreen device from the V1-Trust security zone. You define addresses for the FTP and Mail servers. You also configure a default route to the external router at 1.1.1.250, so that the NetScreen device can send outbound VPN traffic to it[6]. (The default gateway on all hosts in the V1-Trust zone is also 1.1.1.250.)

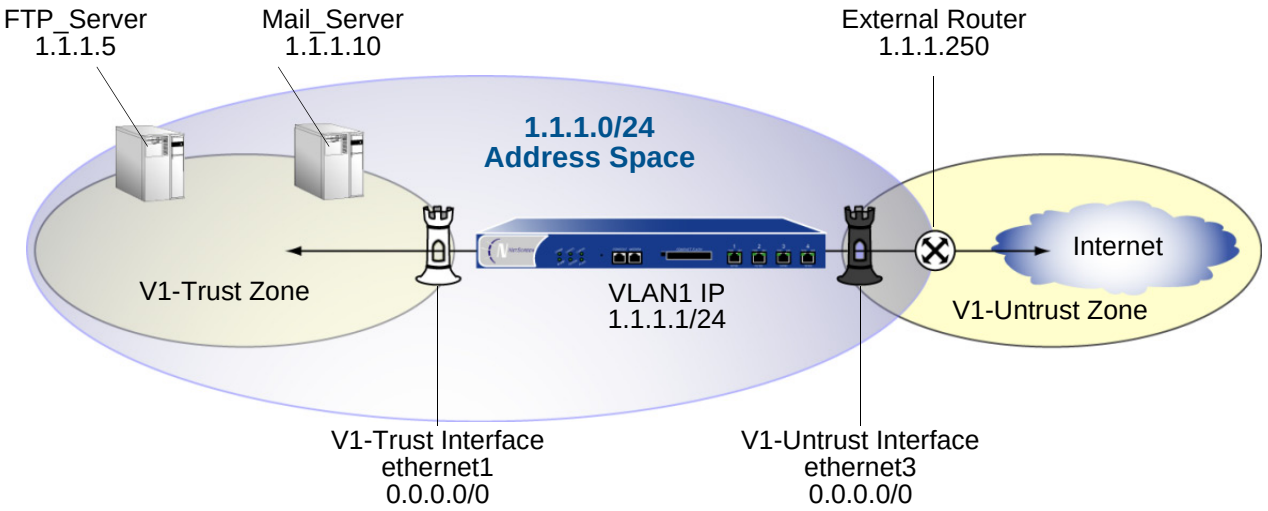

6. For an example of configuring a VPN tunnel for a NetScreen device with interfaces in Transparent mode, see “Transparent Mode VPN” on page **5**-186.

WebUI

1. VLAN1 Interface

Network > Interfaces > Edit (for the VLAN1 interface): Enter the following, and then click **OK**:

IP Address/Netmask: 1.1.1.1/24

Management Services: WebUI, Telnet (select)

Other Services: Ping (select)

2. HTTP Port

Configuration > Admin > Management: In the HTTP Port field, type 5555[7] and then click **Apply**.

3. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **OK**:

Zone Name: V1-Trust

IP Address/Netmask: 0.0.0.0/0

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: V1-Untrust

IP Address/Netmask: 0.0.0.0/0

4. V1-Trust Zone

Network > Zones > Edit (for v1-trust): Select the following, and then click **OK**:

Management Services: WebUI, Telnet

Other Services: Ping

7. The default port number is 80. Changing this to any number between 1024 and 32,767 is advised for discouraging unauthorized access to the configuration. When logging on to manage the device later, enter the following in the URL field of your Web browser: http://1.1.1.1:5555.

5. Addresses

Objects > Addresses > List > New: Enter the following and then click **OK**:

Address Name: FTP _Server

IP Address/Domain Name:

IP/Netmask: (select), 1.1.1.5/32

Zone: V1-Trust

Objects > Addresses > List > New: Enter the following and then click **OK**:

Address Name: Mail_Server

IP Address/Domain Name:

IP/Netmask: (select), 1.1.1.10/32

Zone: V1-Trust

6. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: vlan1(trust-vr)

Gateway IP Address: 1.1.1.250

Metric: 1

7. Policies

Policies > (From: V1-Trust, To: V1-Untrust) New: Enter the following and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Any

Service: Any

Action: Permit

Policies > (From: V1-Untrust, To: V1-Trust) New: Enter the following and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Mail_Server

Service: Mail

Action: Permit

Policies > (From: V1-Untrust, To: V1-Trust) New: Enter the following and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), FTP_Server

Service: FTP-GET

Action: Permit

CLI

1. VLAN1

```
set interface vlan1 ip 1.1.1.1/24
set interface vlan1 manage web
set interface vlan1 manage telnet
set interface vlan1 manage ping
```

2. Telnet

```
set admin telnet port 4646[8]
```

3. Interfaces

```
set interface ethernet1 ip 0.0.0.0/0
set interface ethernet1 zone v1-trust
set interface ethernet3 ip 0.0.0.0/0
set interface ethernet3 zone v1-untrust
```

4. V1-Trust Zone

```
set zone v1-trust manage web
set zone v1-trust manage telnet
set zone v1-trust manage ping
```

5. Addresses

```
set address v1-trust FTP_Server 1.1.1.5/32
set address v1-trust Mail_Server 1.1.1.10/32
```

6. Route

```
set vrouter trust-vr route 0.0.0.0/0 interface vlan1 gateway 1.1.1.250 metric 1
```

7. Policies

```
set policy from v1-trust to v1-untrust any any any permit
set policy from v1-untrust to v1-trust any Mail_Server mail permit
set policy from v1-untrust to v1-trust any FTP_Server ftp-get permit
save
```

8. The default port number for Telnet is 23. Changing this to any number between 1024 and 32,767 is advised for discouraging unauthorized access to the configuration. When logging on to manage the device later via Telnet, enter the following address: 1.1.1.1 4646.

NAT MODE

When an ingress interface is in Network Address Translation (NAT) mode, the NetScreen device, acting like a Layer 3 switch (or router), translates two components in the header of an outgoing IP packet destined for the Untrust zone: its source IP address and source port number. The NetScreen device replaces the source IP address of the originating host with the IP address of the Untrust zone interface. Also, it replaces the source port number with another random port number generated by the NetScreen device.

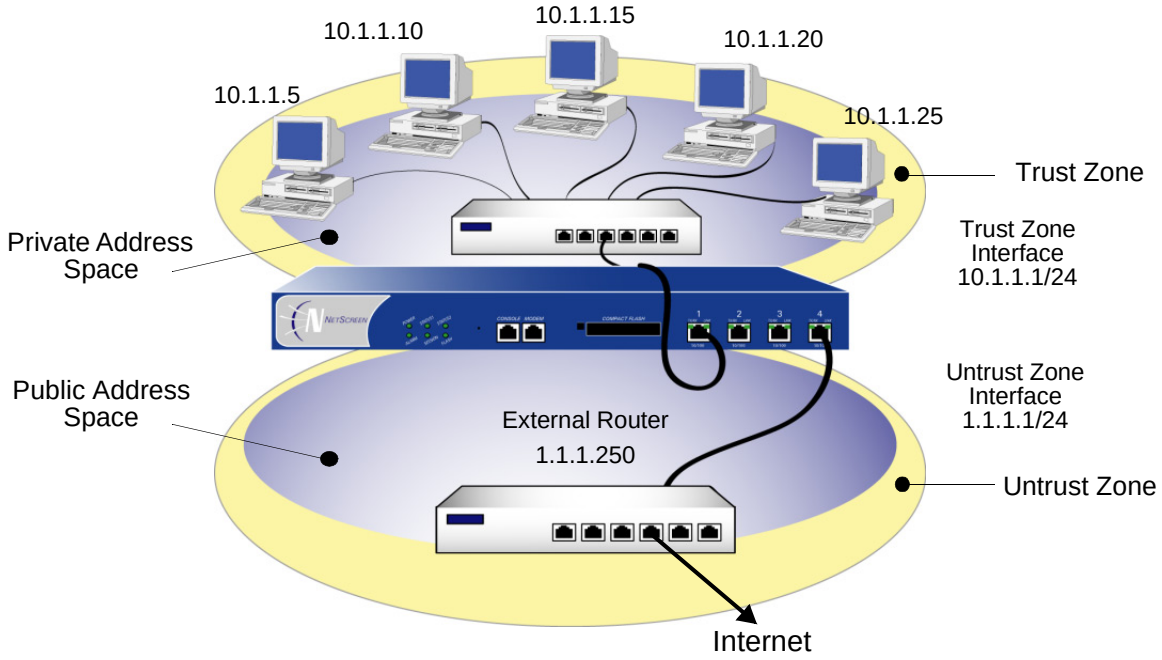

When the reply packet arrives at the NetScreen device, the device translates two components in the IP header of the incoming packet: the destination address and port number, which are translated back to the original numbers. The NetScreen device then forwards the packet to its destination.

NAT adds a level of security not provided in Transparent mode: The addresses of hosts sending traffic through an ingress interface in NAT mode (such as a Trust zone interface) are never exposed to hosts in the egress zone (such as the Untrust zone) unless the two zones are in the same virtual routing domain and the NetScreen device is advertising routes to peers through a dynamic routing protocol (DRP). Even then, the Trust zone addresses are only reachable if you have a policy permitting inbound traffic to them. (If you want to keep the Trust zone addresses hidden while using a DRP, then put the Untrust zone in the untrust-vr and the Trust zone in the trust-vr, and do not export routes for internal addresses in the trust-vr to the untrust-vr.)

If the NetScreen device uses static routing and just one virtual router, the internal addresses remain hidden when traffic is outbound, due to interface-based NAT. The policies you configure control inbound traffic. If you use only mapped IP (MIP) and virtual IP (VIP) addresses as the destinations in your inbound policies, the internal addresses still remain hidden.

Also, NAT preserves the use of public IP addresses. In many environments, resources are not available to provide public IP addresses for all devices on the network. NAT services allow many private IP addresses to have access to Internet resources through one or a few public IP addresses. The following IP address ranges are reserved for private IP networks and must not get routed on the Internet:

10.0.0.0 – 10.255.255.255

172.16.0.0 – 172.31.255.255

192.168.0.0 – 192.168.255.255

Inbound and Outbound NAT Traffic

A host in a zone sending traffic through an interface in NAT mode can initiate traffic to the Untrust zone—assuming that a policy permits it. In releases prior to ScreenOS 5.0.0, a host behind an interface in NAT mode was unable to receive traffic from the Untrust zone unless a Mapped IP (MIP), Virtual IP (VIP), or VPN tunnel was set up for it[9]. However, in ScreenOS 5.0.0, traffic to a zone with a NAT-enabled interface from any zone—including the Untrust zone—does not need to use a MIP, VIP, or VPN. If you want to preserve the privacy of addresses or if you are using private addresses that do not occur on a public network such as the Internet, you can still define a MIP, VIP, or VPN for traffic to reach them. However, if issues of privacy and private IP addresses are not a concern, traffic from the Untrust zone can reach hosts behind an interface in NAT mode directly, without the use of a MIP, VIP, or VPN.

1. Interface-based NAT on traffic from the Trust zone to the Untrust zone.
2. Interface-based NAT on traffic from the User-Defined zone to the Untrust zone.

   (Note: This is possible only if the User-Defined and Untrust zones are in different virtual routing domains.)
3. ***No*** interface-based NAT on traffic between the Trust and User-Defined zones.

4 and 5. You can use MIPs, VIPs, or VPNs for traffic from the Untrust zone to reach the Trust zone or the User-Defined zone, but they are ***not required.***

6. MIPs and VPNs are also ***not required*** for traffic between the Trust and User-Defined zones.

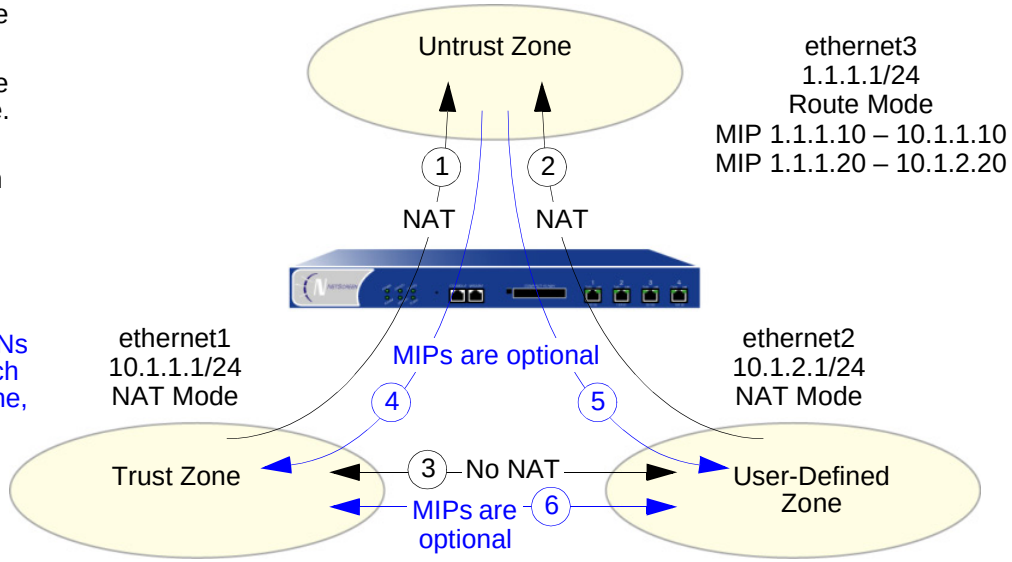

Note: *For more information about MIPs, see "Mapped IP Addresses" on page 331. For more about VIPs, see "Virtual IP Addresses" on page 356.*

9. You can define a virtual IP (VIP) address only on an interface bound to the Untrust zone.

Interface Settings

For NAT mode, define the following interface settings, where *ip_addr1* and *ip_addr2* represent numbers in an IP address, *mask* represents the numbers in a netmask, *vlan_id_num* represents the number of a VLAN tag, *zone* represents the name of a zone, and *number* represents the bandwidth size in kbps:

Zone Interfaces	Settings	Zone Subinterfaces
Trust, DMZ, and user-defined zones using NAT	IP: *ip_addr1* Netmask: *mask* Manage IP*: *ip_addr2* Traffic Bandwidth†: *number* NAT‡: (select)	IP: *ip_addr1* Netmask: *mask* VLAN Tag: *vlan_id_num* Zone Name: *zone* NAT†: (select)
Untrust**	IP: *ip_addr1* Netmask: *mask* Manage IP*: *ip_addr2* Traffic Bandwidth†: *number*	IP: *ip_addr1* Netmask: *mask* VLAN Tag: *vlan_id_num* Zone Name: *zone*

* You can set the manage IP address on a per interface basis. Its primary purpose is to provide an IP address for administrative traffic separate from network traffic. You can also use the manage IP address for accessing a specific device when it is in a high availability configuration.

† Optional setting for traffic shaping.

‡ Selecting NAT defines the interface mode as NAT. Selecting Route defines the interface mode as Route.

** Although you are able to select NAT as the interface mode on an interface bound to the Untrust zone, the NetScreen device does not perform any NAT operations on that interface.

Example: NAT Mode

The following example illustrates a simple configuration for a LAN with a single subnet in the Trust zone. The LAN is protected by a NetScreen device in NAT mode. Policies permit outgoing traffic for all hosts in the Trust zone and incoming mail for the mail server. The incoming mail is routed to the mail server through a Virtual IP address. Both the Trust and Untrust zones are in the trust-vr routing domain.

Note: *Compare this example with that for Route mode on page 120.*

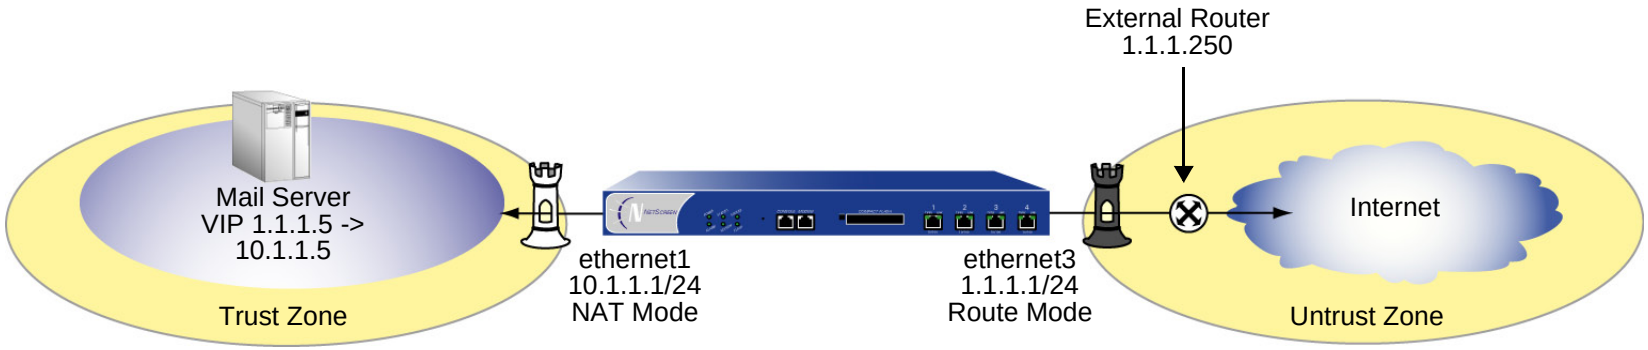

WebUI

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Enter the following, and then click **OK**:

Interface Mode: NAT[10]

10. By default, any interface that you bind to the Trust zone is in NAT mode. Consequently, this option is already enabled for interfaces bound to the Trust zone.

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask[11]: 1.1.1.1/24

Interface Mode: Route

2. VIP[12]

Network > Interfaces > Edit (for ethernet3) > VIP: Enter the following, and then click **Add**:

Virtual IP Address: 1.1.1.5

Network > Interfaces > Edit (for ethernet3) > VIP > New VIP Service: Enter the following, and then click **OK**:

Virtual Port: 25

Map to Service: Mail

Map to IP: 10.1.1.5

3. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 1.1.1.250

11. If the IP address in the Untrust zone on the NetScreen device is dynamically assigned by an ISP, leave the IP address and netmask fields empty and select **Obtain IP using DHCP**. If the ISP uses Point-to-Point Protocol over Ethernet, select **Obtain IP using PPPoE**, click the **Create new PPPoE settings** link, and enter the name and password.

12. For information about virtual IP (VIP) addresses, see "Virtual IP Addresses" on page 356.

4. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Any

Service: ANY

Action: Permit

Policies > (From: Untrust, To: Global) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), VIP(1.1.1.5)

Service: MAIL

Action: Permit

CLI

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat[13]

set interface ethernet3 zone untrust[14]
set interface ethernet3 ip 1.1.1.1/24
set interface ethernet3 route
```

2. VIP

```
set interface ethernet3 vip 1.1.1.5 25 mail 10.1.1.5
```

3. Route

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
```

4. Policies

```
set policy from trust to untrust any any any permit
set policy from untrust to global any vip(1.1.1.5) mail permit
save
```

13. The **set interface ethernet***n* **nat** command determines that the NetScreen device operates in NAT mode.

14. If the IP address in the Untrust zone on the NetScreen device is dynamically assigned by an ISP, use the following command: **set interface untrust dhcp**. If the ISP uses Point-to-Point Protocol over Ethernet, use the **set pppoe** and **exec pppoe** commands. For more information, see the *NetScreen CLI Reference Guide.*

ROUTE MODE

When an interface is in Route mode, the NetScreen device routes traffic between different zones without performing source NAT (NAT-src); that is, the source address and port number in the IP packet header remain unchanged as it traverses the NetScreen device. Unlike NAT-src, you do not need to establish mapped IP (MIP) and virtual IP (VIP) addresses to allow inbound traffic to reach hosts when the destination zone interface is in Route mode. Unlike Transparent mode, the interfaces in each zone are on different subnets.

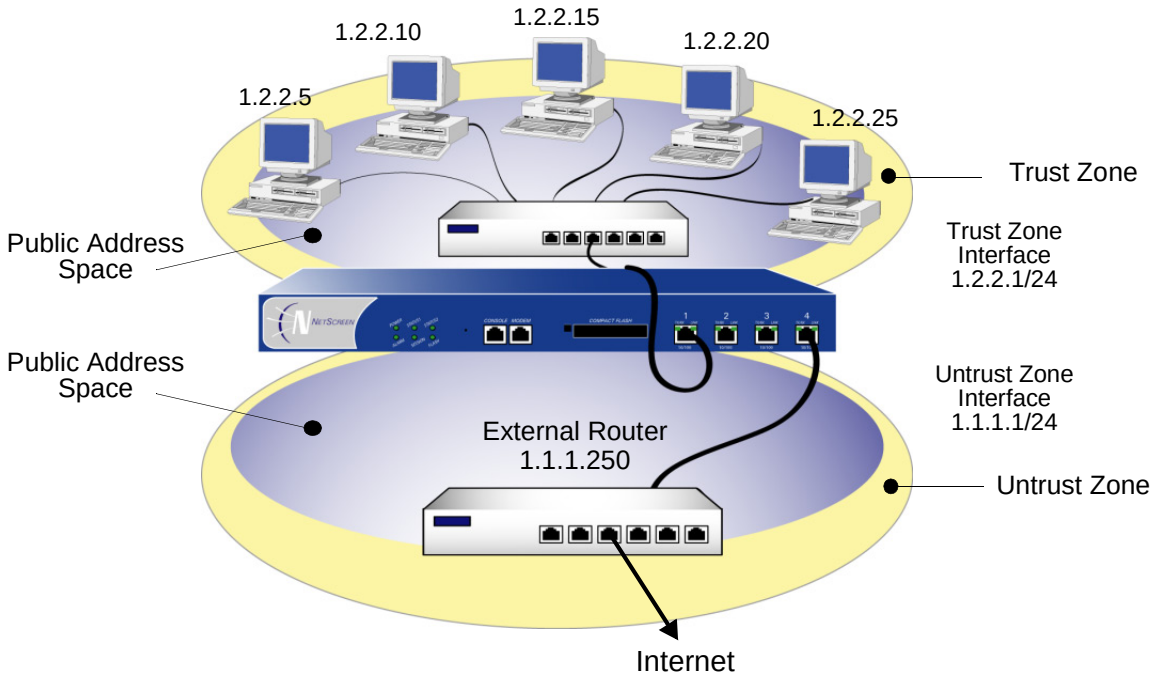

You do not have to apply source network address translation (“NAT-src”) at the interface level so that all source addresses initiating outgoing traffic get translated to the IP address of the destination zone interface. Instead, you can perform NAT-src selectively at the policy level. You can determine which traffic to route and on which traffic to perform NAT-src by creating policies that enable NAT-src for specified source addresses on either incoming or

outgoing traffic. For network traffic, NAT can use the IP address or addresses of the destination zone interface from a Dynamic IP (DIP) pool, which is in the same subnet as the destination zone interface. For VPN traffic, NAT can use a tunnel interface IP address or an address from its associated DIP pool.

Note: *For more information about configuring policy-based NAT-src, see* *"Source Network Address Translation" on page 259**.*

Interface Settings

For Route mode, define the following interface settings, where *ip_addr1* and *ip_addr2* represent numbers in an IP address, *mask* represents the numbers in a netmask, *vlan_id_num* represents the number of a VLAN tag, *zone* represents the name of a zone, and *number* represents the bandwidth size in kbps:

Zone Interfaces	Settings	Zone Subinterfaces
Trust, Untrust, DMZ, and user-defined zones	IP: *ip_addr1* Netmask: *mask* Manage IP*: *ip_addr2* Traffic Bandwidth†: *number* Route‡: (select)	IP: *ip_addr1* Netmask: *mask* VLAN Tag: *vlan_id_num* Zone Name: *zone* Route†: (select)

* You can set the manage IP address on a per interface basis. Its primary purpose is to provide an IP address for administrative traffic separate from network traffic. You can also use the manage IP address for accessing a specific device when it is in a high availability configuration.

† Optional setting for traffic shaping.

‡ Selecting Route defines the interface mode as Route. Selecting NAT defines the interface mode as NAT.

Example: Route Mode

In the previous example, “Example: NAT Mode” on page 114, the hosts in the Trust zone LAN have private IP addresses and a Mapped IP for the mail server. In the following example of the same network protected by a NetScreen device operating in Route mode, note that the hosts have public IP addresses and that a MIP is unnecessary for the mail server. Both security zones are in the trust-vr routing domain.

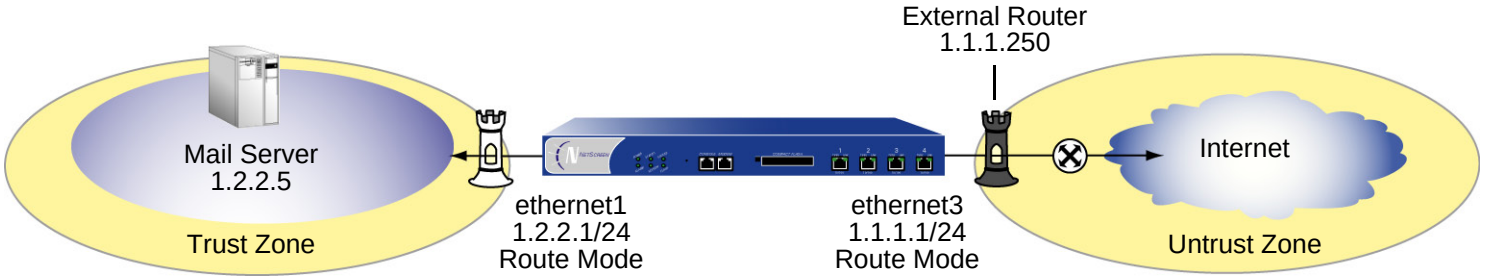

WebUI

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 1.2.2.1/24

Enter the following, and then click **OK**:

Interface Mode: Route[15]

15. Selecting **Route** determines that the NetScreen device operates in Route mode, without performing NAT on traffic entering or exiting the Trust zone.

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask[16]: 1.1.1.1/24

2. Address

Objects > Addresses > List > New: Enter the following and then click **OK**:

Address Name: Mail Server

IP Address/Domain Name:

IP/Netmask: (select), 1.2.2.5/32

Zone: Trust

3. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 1.1.1.250

16. If the IP address in the Untrust zone on the NetScreen device is dynamically assigned by an ISP, leave the IP address and netmask fields empty and select **Obtain IP using DHCP**. If the ISP uses Point-to-Point Protocol over Ethernet, select **Obtain IP using PPPoE**, click the **Create new PPPoE settings** link, and enter the name and password.

4. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Any

Service: ANY

Action: Permit

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Mail Server

Service: MAIL

Action: Permit

CLI

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 1.2.2.1/24
set interface ethernet1 route[17]

set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
set interface ethernet3 route
```

2. Address

```
set address trust mail_server 1.2.2.5/24
```

3. Route

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
```

4. Policies

```
set policy from trust to untrust any any any permit
set policy from untrust to trust any mail_server mail permit
save
```

17. The **set interface ethernet***number* **route** command determines that the NetScreen device operates in Route mode.

6

Building Blocks for Policies

This chapter discusses the components, or building blocks, that you can reference in policies. The specific topics discussed are:

- "Addresses" on page 126
  - "Address Entries" on page 127
  - "Address Groups" on page 129
- "Services" on page 134
  - "Predefined Services" on page 134
  - "Custom Services" on page 136
  - "ICMP Services" on page 139
  - "RSH ALG" on page 140
  - "H.323 Protocol for Voice-over-IP" on page 141
  - "SIP – Session Initiation Protocol" on page 156
  - "Service Groups" on page 167
- "DIP Pools" on page 171
  - "Sticky DIP Addresses" on page 174
  - "Extended Interface and DIP" on page 175
  - "Loopback Interface and DIP" on page 183
  - "DIP Groups" on page 189
- "Schedules" on page 193

> ***Note:*** *For information about user authentication, see Chapter 9, "User Authentication" on page 371.*

ADDRESSES

The NetScreen ScreenOS classifies the addresses of all other devices by location and netmask. Each zone possesses it's own list of addresses and address groups.

Individual hosts have only a single IP address defined and therefore, must have a netmask setting of 255.255.255.255 (which masks out all but this host).

Subnets have an IP address and a netmask (for example, 255.255.255.0 or 255.255.0.0).

Before you can configure policies to permit, deny, or tunnel traffic to and from individual hosts and subnets, you must make entries for them in NetScreen address lists, which are organized by zones.

Note: *You do not have to make address entries for "Any". This term automatically applies to all devices physically located within their respective zones.*

Address Entries

Before you can set up many of the NetScreen firewall, VPN, and traffic shaping features, you need to define addresses in one or more address lists. The address list for a security zone contains the IP addresses or domain names[1] of hosts or subnets whose traffic is either allowed, blocked, encrypted, or user-authenticated.

Note: *For information regarding ScreenOS naming conventions—which apply to the names you create for addresses—see* *"Naming Conventions and Character Types" on page xiv**.*

Example: Adding Addresses

In this example, you add the subnet "Sunnyvale_Eng" with the IP address 10.1.10.0/24 as an address in the Trust zone, and the address www.firenet.com as an address in the Untrust zone.

WebUI

Objects > Addresses > List > New: Enter the following information, and then click **OK**:

Address Name: Sunnyvale_Eng

IP Address/Domain Name:

IP/Netmask: (select), 10.1.10.0/24

Zone: Trust

Objects > Addresses > List > New: Enter the following information, and then click **OK**:

Address Name: FireNet

IP Address/Domain Name:

Domain Name: (select), www.firenet.com

Zone: Untrust

1. Before you can use domain names for address entries, you must configure the NetScreen device for Domain Name System (DNS) services. For information on DNS configuration, see "Domain Name System Support" on page 495.

CLI

```
set address trust Sunnyvale_Eng 10.1.10.0/24
set address untrust FireNet www.firenet.com
save
```

Example: Modifying Addresses

In this example, you change the address entry for the address “Sunnyvale_Eng” to reflect that this department is specifically for software engineering and has a different IP address—10.1.40.0/24.

WebUI

Objects > Addresses > List > Edit (for Sunnyvale_Eng): Change the name and IP address to the following, and then click **OK**:

Address Name: Sunnyvale_SW_Eng

IP Address/Domain Name:

IP/Netmask: (select), 10.1.40.0/24

Zone: Trust

CLI

```
unset address trust Sunnyvale_Eng
set address trust Sunnyvale_SW_Eng 10.1.40.0/24
save
```

Note: *After you define an address—or an address group—and associate it with a policy, you cannot change the address location to another zone (such as from Trust to Untrust). To change its location, you must first disassociate it from the underlying policy.*

Example: Deleting Addresses

In this example, you remove the address entry for the address “Sunnyvale_SW_Eng”.

WebUI

Objects > Addresses > List: Click **Remove** in the Configure column for Sunnyvale_SW_Eng.

CLI

```
unset address trust “Sunnyvale_SW_Eng”
save
```

Address Groups

The previous section explained how you create, modify, and delete address book entries for individual hosts and subnets. As you add addresses to an address list, it becomes difficult to manage how policies affect each address entry. NetScreen allows you to create groups of addresses. Rather than manage a large number of address entries, you can manage a small number of groups. Changes you make to the group are applied to each address entry in the group.

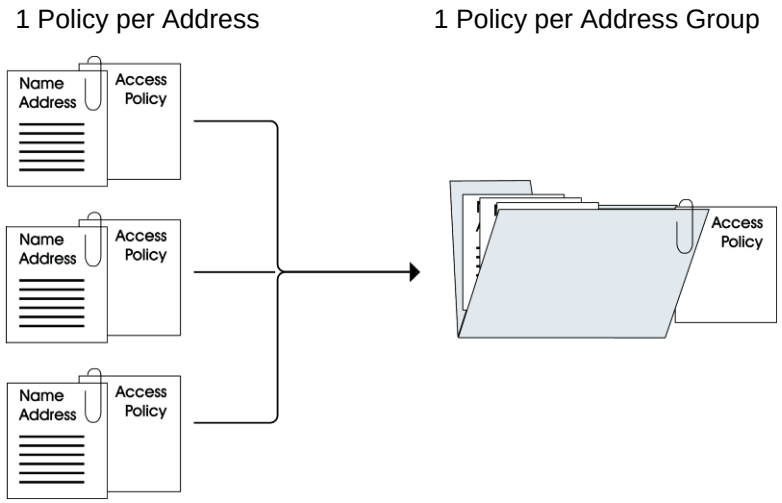

The address group option has the following features:

- You can create address groups in any zone.
- You can create address groups with existing users, or you can create empty address groups and later fill them with users.
- An address group can be a member of another address group[2].
- You can reference an address group entry in a policy like an individual address book entry.
- NetScreen applies policies to each member of the group by internally creating individual policies for each group member. While you only have to create one policy for a group, NetScreen actually creates an internal policy for each member in the group (as well as for each service configured for each user).[3]
- When you delete an individual address book entry from the address book, the NetScreen device automatically removes it from all groups to which it belonged.

The following constraints apply to address groups:

- Address groups can only contain addresses that belong to the same zone.
- Address names cannot be the same as group names. If the name “Paris” is used for an individual address entry, it cannot be used for a group name.
- If an address group is referenced in a policy, the group cannot be removed. It can, however, be edited.
- When a single policy is assigned to an address group, it is applied to each group member individually, and the NetScreen device makes an entry for each member in the access control list (ACL). If you are not vigilant, it is possible to exceed the number of available policy resources, especially if both the source and destination addresses are address groups and the specified service is a service group.
- You cannot add the predefined addresses: “Any”, “All Virtual IPs,” and “Dial-Up VPN” to groups.

2. To ensure that a group does not accidentally contain itself as a member, the NetScreen device performs a sanity check when you add one group to another. For example, if you add group A as a member to group B, the NetScreen device automatically checks that A does not already contain B as its member.

3. The automatic nature by which the NetScreen device applies policies to each address group member, saves you from having to create them one by one for each address. Furthermore, NetScreen writes these policies to ASIC which makes lookups run very fast.

Example: Creating an Address Group

In the following example, you create a group named “HQ 2nd Floor” that includes “Santa Clara Eng” and “Tech Pubs,” two addresses that you have already entered in the address book for the Trust zone.

WebUI

Objects > Addresses > Groups > (for Zone: Trust) New: Enter the following group name, move the following addresses, and then click **OK**:

Group Name: HQ 2nd Floor

Select **Santa Clara Eng** and use the **<<** button to move the address from the Available Members column to the Group Members column.

Select **Tech Pubs** and use the **<<** button to move the address from the Available Members column to the Group Members column.

CLI

```
set group address trust “HQ 2nd Floor” add “Santa Clara Eng”
set group address trust “HQ 2nd Floor” add “Tech Pubs”
save
```

Example: Editing a Group Address Entry

In this example, you add “Support” (an address that you have already entered in the address book) to the “HQ 2nd Floor” address group.

WebUI

Objects > Addresses > Groups > (for Zone: Trust) Edit (for HQ 2nd Floor): Move the following address, and then click **OK**:

Select **Support** and use the **<<** button to move the address from the Available Members column to the Group Members column.

CLI

```
set group address trust “HQ 2nd Floor” add Support
save
```

Example: Removing an Address Group Member and a Group

In this example, you remove the member “Support” from the HQ 2nd Floor address group, and delete “Sales”, an address group that you had previously created.

WebUI

Objects > Addresses > Groups > (for Zone: Trust) Edit (HQ 2nd Floor): Move the following address, and then click **OK**:

Select **support** and use the **>>** button to move the address from the Group Members column to the Available Members column.

Objects > Addresses > Groups > (Zone: Trust): Click **Remove** in the Configure column for Sales.

CLI

```
unset group address trust “HQ 2nd Floor” remove Support
unset group address trust Sales
save
```

Note: *The NetScreen device does not automatically delete a group from which you have removed all names.*

SERVICES

Services are types of IP traffic for which protocol standards exist. Each service has a port number associated with it, such as 21 for FTP and 23 for Telnet. When you create a policy, you must specify a service for it. You can select one of the predefined services from the service book, or a custom service or service group that you created. You can see which service you can use in a policy by viewing the Service drop-down List in the Policy Configuration dialog box (WebUI), or by using the **get service** command (CLI).

Predefined Services

ScreenOS supports a great number of predefined services. Later in this section, you can find more detailed information on some of these, namely:

- "ICMP Services" on page 139
- "RSH ALG" on page 140
- "H.323 Protocol for Voice-over-IP" on page 141
- "SIP – Session Initiation Protocol" on page 156

You can view the list of predefined or custom services or service groups on the NetScreen device using the WebUI or the CLI.

Using the WebUI:

Objects > Services > Predefined

Objects > Services > Custom

Objects > Services > Group

Using the CLI:

```
get service [ group | predefined | user ]
```

The output from the **get service pre-defined** CLI is similar to that shown below:

Name	Proto	Port	Group	Timeout (Minute)	Flag
ANY	0	0/65535	other	1	Pre-defined
AOL	6	5190/5194	remote	30	Pre-defined
BGP	6	179	other	30	Pre-defined
DHCP-Relay	17	67	info seeking	1	Pre-defined
DNS	17	53	info seeking	1	Pre-defined
FINGER	6	79	info seeking	30	Pre-defined
FTP	6	21	remote	30	Pre-defined
FTP-Get	6	21	remote	30	Pre-defined
FTP-Put	6	21	remote	30	Pre-defined
GOPHER	6	70	info seeking	30	Pre-defined
H.323	6	1720	remote	2160	Pre-defined
--- more ---					

Note: *Each predefined service has a source port range of 1-65535, which includes the entire set of valid port numbers. This prevents potential attackers from gaining access by using a source port outside of the range. If you need to use a different source port range for any predefined service, create a custom service. For information, see* *"Custom Services" on page 136**.*

You can set the timeout threshold (in minutes) for a predefined or custom service. You can use the service default timeout, specify a custom timeout, or use no timeout at all.

Example: Setting a Predefined Service Timeout

In this example, you change the timeout threshold for the BGP predefined service to 75 minutes:

WebUI

Objects > Services > Predefined > Edit (BGP): Enter the following and then click **OK**:

Service Timeout: Custom (select), 75 (type)

CLI

```
set service BGP timeout 75
save
```

Custom Services

Instead of using predefined services, you can easily create your own with a custom name, port number and transport protocol. The following examples describe how to add, modify and remove a custom service.

Note: *For information regarding ScreenOS naming conventions—which apply to the names you create for custom services—see* *"Naming Conventions and Character Types" on page xiv**.*

Example: Adding a Custom Service

To add a custom service to the service book, you need the following information:

- A name for the service, in this example "cust-telnet"
- A range of source port numbers: 1 – 65535
- A range of destination port numbers to receive the service request, for example: 23000 – 23000.
- Whether the service uses TCP or UDP protocol, or some other protocol as defined by the Internet specifications. In this example, the protocol is TCP.

WebUI

Objects > Services > Custom > New: Enter the following, and then click **OK**:

Service Name: cust-telnet

Service Timeout: Custom (select), 30 (type)

Transport Protocol: TCP (select)

Source Port Low: 1

Source Port High: 65535

Destination Port Low: 23000

Destination Port High: 23000

CLI

```
set service cust-telnet protocol tcp src-port 1-65535 dst-port 23000-23000
set service cust-telnet timeout 30[4]
save
```

4. The timeout value is in minutes. If you do not set it, the timeout value of a custom service is 180 minutes. If you do not want a service to time out, enter **never**.

Example: Modifying a Custom Service

In this example, you modify the custom service “cust-telnet” by changing the destination port range to 23230-23230.

Use the **set service** *service_name* **clear** command to remove the definition of a custom service without removing the service from the service book:

WebUI

Objects > Services > Custom > Edit (for cust-telnet): Enter the following, and then click **OK**:

Destination Port Low: 23230

Destination Port High: 23230

CLI

```
set service cust-telnet clear
set service cust-telnet + tcp src-port 1-65535 dst-port 23230-23230
save
```

Example: Removing a Custom Service

In this example, you remove the custom service “cust-telnet”.

WebUI

Objects > Services > Custom: Click **Remove** in the Configure column for “cust-telnet”.

CLI

```
unset service cust-telnet
save
```

ICMP Services

ScreenOS supports ICMP (Internet Control Message Protocol) as well as several ICMP messages, as predefined or custom services. When configuring a custom ICMP service, you must define a type and code[5]. There are different message types within ICMP. For example:

type 0 = Echo Request message

type 3 = Destination Unreachable message

An ICMP message type can also have a message code. The code provides more specific information on the message. For example:

Message Type	Message Code
5 = Redirect	0 = Redirect Datagram for the Network (or subnet)
	1 = Redirect Datagram for the Host
	2 = Redirect Datagram for the Type of Service and Network
	3 = Redirect Datagram for the Type of Service and Host
11 = Time Exceeded Codes	0 =Time to Live exceeded in Transit
	1 = Fragment Reassembly Time Exceeded

ScreenOS supports any type or code within the 0-255 range.

5. For more information on ICMP types and codes, refer to *RFC 792*.

Example: Defining an ICMP Service

In this example, you define a custom service named “host-unreachable” using ICMP as the transport protocol. The type is 3 (for Destination Unreachable) and the code is 1 (for Host Unreachable). You set the timeout value at 2 minutes.

WebUI

Objects > Services > Custom: Enter the following, and then click **OK**:

Service Name: host-unreachable

Service Timeout: Custom (select), 2 (type)

Transport Protocol: ICMP (select)

ICMP Type: 3

ICMP Code: 1

CLI

```
set service host-unreachable protocol icmp type 5 code 0
set service host-unreachable timeout 2
save
```

RSH ALG

RSH ALG (Remote Shell application-layer gateway) allows authenticated users to run shell commands on remote hosts. NetScreen devices support the RSH service in Transparent (L2), Route (L3) and NAT modes; but the devices do not support port translation of RSH traffic.

H.323 Protocol for Voice-over-IP

To allow secure Voice-over-IP (VoIP) communication between terminal hosts, NetScreen devices support H.323 protocol. In such a telephony system, gatekeeper devices manage call registration, admission, and call status for VoIP calls. Gatekeepers can reside in the two different zones, or in the same zone.

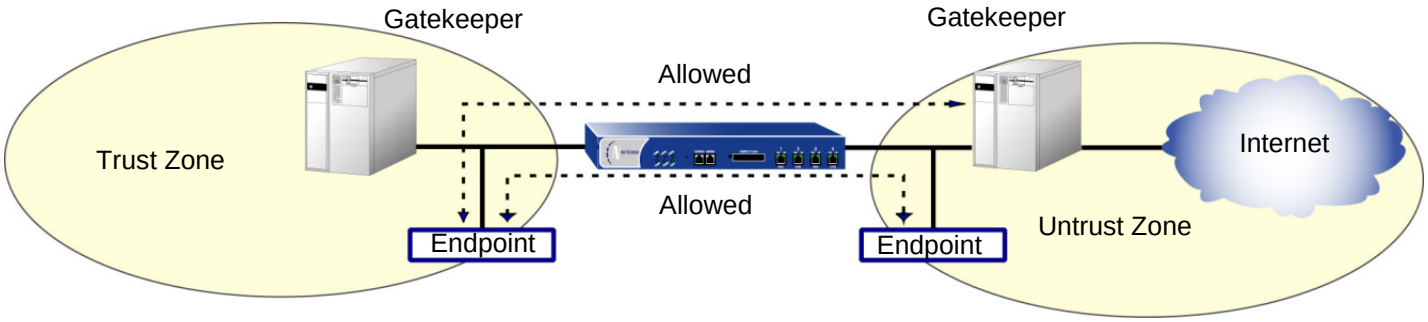

Example: Gatekeeper in the Trust Zone (Transparent or Route Mode)

In the following example, you set up two policies. Together, these policies allow H.323 traffic to pass between IP phone hosts and a gatekeeper in the Trust zone, and an IP phone host (2.2.2.5) in the Untrust zone. In this example, the NetScreen device can be in either Transparent mode or Route mode. Both the Trust and Untrust security zones are in the trust-vr routing domain.

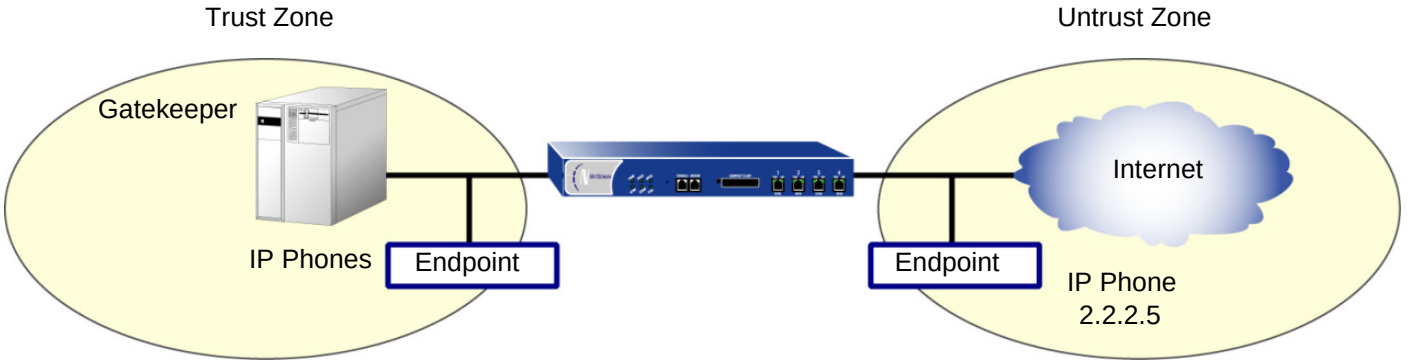

WebUI

1. Address

   Objects > Addresses > List > New: Enter the following, and then click **OK**:

   Address Name: IP_Phone

   IP Address/Domain Name:

   IP/Netmask: (select), 2.2.2.5/32

   Zone: Untrust

2. Policies

   Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

   Source Address:

   Address Book Entry: (select), Any

   Destination Address:

   Address Book Entry: (select), IP_Phone

   Service: H.323

   Action: Permit

   Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

   Source Address:

   Address Book Entry: (select), IP_Phone

   Destination Address:

   Address Book Entry: (select), Any

   Service: H.323

   Action: Permit

CLI

1. Address

```
set address untrust IP_Phone 2.2.2.5/32
```

2. Policies

```
set policy from trust to untrust any IP_Phone h.323 permit
set policy from untrust to trust IP_Phone any h.323 permit
save
```

Example: Gatekeeper in the Trust Zone (NAT Mode)

When the NetScreen device is in NAT mode, a gatekeeper or endpoint device is said to be *private* when it resides in the Trust zone, and *public* when it resides in the Untrust zone. When you set a NetScreen device in NAT mode, you must map a public IP address to each private device.

In this example, the devices in the Trust zone include the endpoint host (10.1.1.5/32) and the gatekeeper device (10.1.1.25/32). IP_Phone2 (2.2.2.5/32) is in the Untrust zone. You configure the NetScreen device to allow traffic between the endpoint host IP_Phone1 and the gatekeeper in the Trust zone and the endpoint host IP_Phone2 in the Untrust zone. Both the Trust and Untrust security zones are in the trust-vr routing domain.

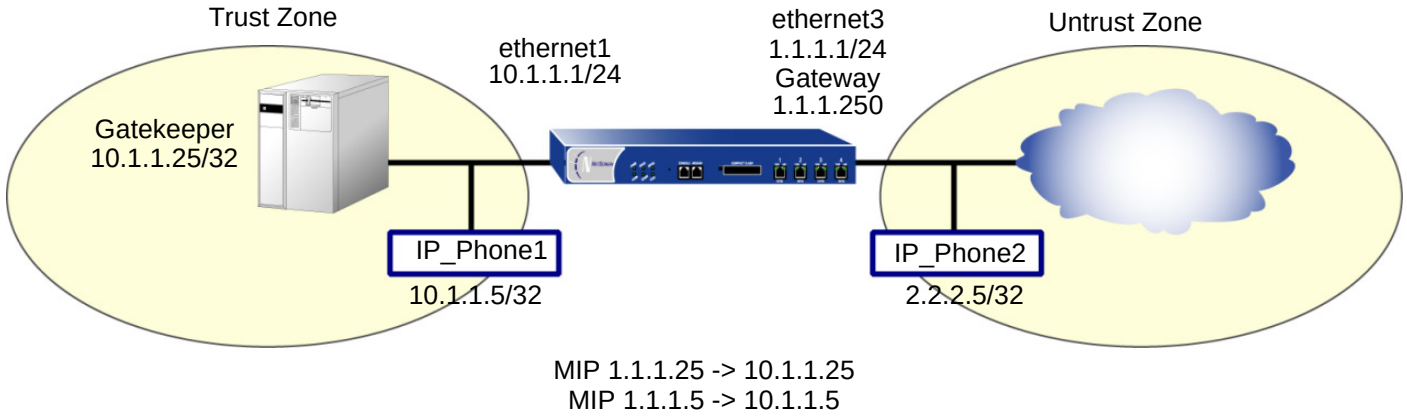

WebUI

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Enter the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

2. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: IP_Phone1

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.5/32

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Gatekeeper

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.25/32

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: IP_Phone2

IP Address/Domain Name:

IP/Netmask: (select), 2.2.2.5/32

Zone: Untrust

3. Mapped IP Addresses

Network > Interfaces > Edit (for ethernet3) > MIP > New: Enter the following, and then click **OK**:

Mapped IP: 1.1.1.5

Netmask: 255.255.255.255

Host IP Address: 10.1.1.5

Host Virtual Router Name: trust-vr

Network > Interfaces > Edit (for ethernet3) > MIP > New: Enter the following, and then click **OK**:

Mapped IP: 1.1.1.25

Netmask: 255.255.255.255

Host IP Address: 10.1.1.25

Host Virtual Router Name: trust-vr

4. Route

Network > Routing > Routing Entries> trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 1.1.1.250

5. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), IP_Phone1

Destination Address:

Address Book Entry: (select), Phone2

Service: H.323

Action: Permit

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Gatekeeper

Destination Address:

Address Book Entry: (select), Phone2

Service: H.323

Action: Permit

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), IP_Phone2

Destination Address:

Address Book Entry: (select), MIP(1.1.1.5)

Service: H.323

Action: Permit

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), IP_Phone2

Destination Address:

Address Book Entry: (select), MIP(1.1.1.25)

Service: H.323

Action: Permit

CLI

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
```

2. Addresses

```
set address trust IP_Phone1 10.1.1.5/32
set address trust gatekeeper 10.1.1.25/32
set address untrust IP_Phone2 2.2.2.5/32
```

3. Mapped IP Addresses

```
set interface ethernet3 mip 1.1.1.5 host 10.1.1.5
set interface ethernet3 mip 1.1.1.25 host 10.1.1.25
```

4. Route

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
```

5. Policies

```
set policy from trust to untrust IP_Phone1 IP_Phone2 h.323 permit
set policy from trust to untrust gatekeeper IP_Phone2 h.323 permit
set policy from untrust to trust IP_Phone2 mip(1.1.1.5) h.323 permit
set policy from untrust to trust IP_Phone2 mip (1.1.1.25) h.323 permit
save
```

Example: Gatekeeper in the Untrust Zone (Transparent or Route Mode)

Because Transparent mode and Route mode do not require address mapping of any kind, NetScreen device configuration for a gatekeeper in the Untrust zone is usually identical to the configuration for a gatekeeper in the Trust zone.

In the following example, you set up two policies to allow H.323 traffic to pass between IP phone hosts (and the gatekeeper) in the Trust zone, and the IP phone at IP address 2.2.2.5 in the Untrust zone. The device can be in Transparent or Route mode. Both the Trust and Untrust security zones are in the trust-vr routing domain.

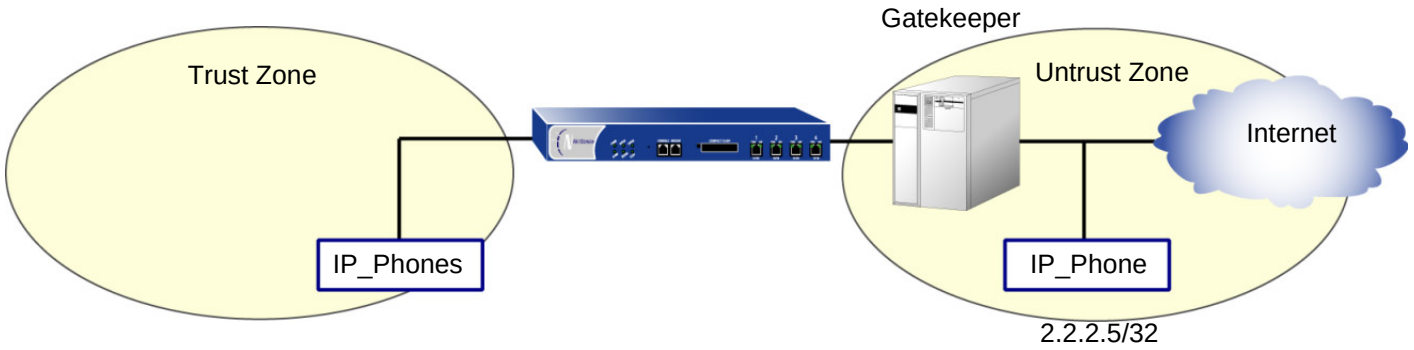

WebUI

1. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: IP_Phone

IP Address/Domain Name:

IP/Netmask: (select), 2.2.2.5/32

Zone: Untrust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Gatekeeper

IP Address/Domain Name:

IP/Netmask: (select), 2.2.2.10/32

Zone: Untrust

2. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), IP_Phone

Service: H.323

Action: Permit

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), IP_Phone

Destination Address:

Address Book Entry: (select), Any

Service: H.323

Action: Permit

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Gatekeeper

Service: H.323

Action: Permit

CLI

1. Addresses

```
set address untrust IP_Phone 2.2.2.5/32
set address untrust gatekeeper 2.2.2.10/32
```

2. Policies

```
set policy from trust to untrust any IP_Phone h.323 permit
set policy from untrust to trust IP_Phone any h.323 permit
set policy from trust to untrust any gatekeeper h.323 permit
save
```

Example: Gatekeeper in the Untrust Zone (NAT Mode)

In this example, the gatekeeper device (2.2.2.25) and host IP_Phone2 (2.2.2.5) are in the Untrust zone and host IP_Phone1 (10.1.1.5) is in the Trust zone. You configure the NetScreen device to allow traffic between host IP_Phone1 in the Trust zone, and host IP_Phone2 (and the gatekeeper) in the Untrust zone. Both the Trust and Untrust security zones are in the trust-vr routing domain.

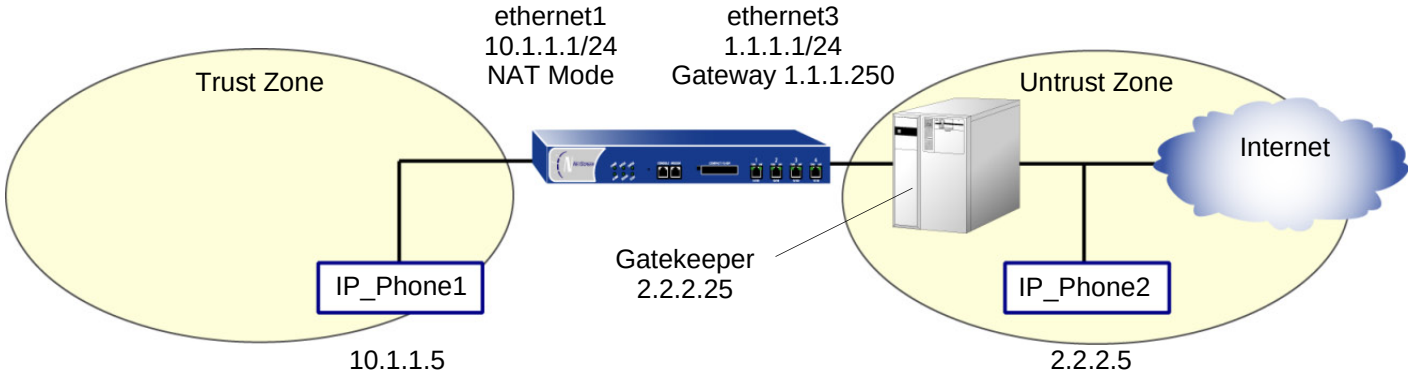

WebUI

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Enter the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

2. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: IP_Phone1

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.5/32

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Gatekeeper

IP Address/Domain Name:

IP/Netmask: (select), 2.2.2.25/32

Zone: Untrust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: IP_Phone2

IP Address/Domain Name:

IP/Netmask: (select), 2.2.2.5/32

Zone: Untrust

3. Mapped IP Address

Network > Interfaces > Edit (for ethernet3) > MIP > New: Enter the following, and then click **OK**:

Mapped IP: 1.1.1.5

Netmask: 255.255.255.255

Host IP Address: 10.1.1.5

4. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 1.1.1.250

5. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), IP_Phone1

Destination Address:

Address Book Entry: (select), IP_Phone2

Service: H.323

Action: Permit

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), IP_Phone1

Destination Address:

Address Book Entry: (select), Gatekeeper

Service: H.323

Action: Permit

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), IP_Phone2

Destination Address:

Address Book Entry: (select), MIP(1.1.1.5)

Service: H.323

Action: Permit

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Gatekeeper

Destination Address:

Address Book Entry: (select), MIP(1.1.1.5)

Service: H.323

Action: Permit

CLI

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
```

2. Addresses

```
set address trust IP_Phone1 10.1.1.5/32
set address untrust gatekeeper 2.2.2.25/32
set address untrust IP_Phone2 2.2.2.5/32
```

3. Mapped IP Addresses

```
set interface ethernet3 mip 1.1.1.5 host 10.1.1.5
```

4. Route

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
```

5. Policies

```
set policy from trust to untrust IP_Phone1 IP_Phone2 h.323 permit
set policy from trust to untrust IP_Phone1 gatekeeper h.323 permit
set policy from untrust to trust IP_Phone2 mip(1.1.1.5) h.323 permit
set policy from untrust to trust gatekeeper mip(1.1.1.5) h.323 permit
save
```

SIP – Session Initiation Protocol

SIP (Session Initiation Protocol) is an IETF (Internet Engineering Task Force)-standard protocol for initiating, modifying, and terminating multimedia sessions over the Internet. Such sessions might include conferencing, telephony, or multimedia, with features such as instant messaging and application-level mobility in network environments.

NetScreen devices support SIP as a service and can screen SIP traffic, allowing and denying it based on a policy that you configure. SIP is a predefined service in ScreenOS and uses port 5060 as the destination port. Note that NetScreen devices currently do not support SIP with NAT (network address translation).

Essentially, SIP is used to distribute the session description and, during the session, to negotiate and modify the parameters of the session. SIP is also used to terminate a multimedia session.

A user includes the session description either in an INVITE or an ACK request. A session description indicates the multimedia type of the session, for example, voice or video. SIP can use different description protocols to describe the session; NetScreen supports SDP (Session Description Protocol) only.

SDP provides information that a system can use to join a multimedia session. SDP might include information such as IP addresses, port numbers, times and dates. Note that the IP address and port number in the SDP header (the “c=” and “m=” fields respectively) are the address and port where the client wants to receive the media streams, and not the IP address and port number from which the SIP request originates (although they can be the same). See “SDP” on page 160 for more information.

SIP messages consist of requests from a client to a server and responses to the requests from a server to a client with the purpose of establishing a session (or a call). A UA (User Agent) is an application that runs at the endpoints of the call and consists of two parts: the UAC (User Agent Client) that sends SIP requests on behalf of the user, and a UAS (User Agent Server) who listens to the responses and notifies the user when they arrive. Examples of User Agents are SIP proxy servers and SIP phones.

SIP Request Methods

There are mainly six types of SIP requests that each fulfill a different purpose. Every SIP request contains a *method* field, which denotes the purpose of the request. The following lists the six different methods.

INVITE – A user sends an INVITE request to invite another user to participate in a session. The body of an INVITE request may contain the description of the session.

ACK – The user from whom the INVITE originated sends an ACK request to confirm reception of the final response to the INVITE. If the original INVITE request did not contain the session description, then the ACK request must include it.

OPTIONS – A user sends an OPTIONS request to a server to get information on its capabilities. A server replies with information such as which methods, session description protocols, and message encoding it supports.

BYE – A user sends a BYE request to abandon a session. A BYE request from either one of the users automatically terminates the session.

CANCEL – A user can send a CANCEL request to cancel a pending INVITE request. A CANCEL request has no effect if the SIP server processing the INVITE already sent a final response for the INVITE before it received the CANCEL.

REGISTER – A user sends a REGISTER request to a SIP *registrar* server to inform it of their current location. A SIP registrar server records all the information it receives in REGISTER requests and makes this information available to any SIP server trying to locate a user.

Classes of SIP Responses

SIP responses indicate the status of the transaction. They consist of codes grouped into the following classes:

100 to 199 – Informational: request received, continuing to process the request

200 to 299 – Success: the action was successfully received, understood, and accepted

300 to 399 – Redirection: further action needs to be taken in order to complete the request

400 to 499 – Client Error: the request contains bad syntax or cannot be fulfilled at this server

500 to 599 – Server Error: the server failed to fulfill an apparently valid request

600 to 699 – Global Failure: the request cannot be fulfilled at any server

The following is the complete list of current SIP response codes. NetScreen supports all of them.

1xx	100 Trying	180 Ringing	181 Call is being forwarded
	182 Queued	183 Session progress	
2xx	200 OK	202 Accepted	
3xx	300 Multiple choices	301 Moved permanently	302 Moved temporarily
	305 Use proxy	380 Alternative service	
4xx	400 Bad request	401 Unauthorized	402 Payment required
	403 Forbidden	404 Not found	405 Method not allowed
	406 Not acceptable	407 Proxy authentication required	408 Request time-out
	409 Conflict	410 Gone	411 Length required
	413 Request entity too large	414 Request-URL too large	415 Unsupported media type
	420 Bad extension	480 Temporarily not available	481 Call leg/transaction does not exist
	482 Loop detected	483 Too many hops	484 Address incomplete
	485 Ambiguous	486 Busy here	487 Request cancelled
	488 Not acceptable here		
5xx	500 Server internal error	501 Not implemented	502 Bad gateway
	502 Service unavailable	504 Gateway time-out	505 SIP version not supported
6xx	600 Busy everywhere	603 Decline	604 Does not exist anywhere
	606 Not acceptable		

ALG – Application-Layer Gateway

There are two types of SIP traffic, the signaling and the media stream. SIP signaling traffic consists of the requests and responses between client and server and uses transport protocols such as UDP or TCP. The media stream carries the data (for example, audio data), and uses application layer protocols such as RTP (Real-time Transport Protocol) over UDP.

NetScreen devices support SIP signaling messages on port 5060. You can simply create a policy that permits SIP service and the NetScreen device filters SIP signaling traffic like any other type of traffic, permitting or denying it. The media stream, however, uses dynamically assigned port numbers that can change several times during the course of a call. Without fixed ports, it is impossible to create a static policy to control the media traffic. In this case, the NetScreen device invokes the SIP ALG. The SIP ALG reads SIP messages and their SDP content and extracts the port number information it needs to dynamically open pinholes[6] and let the media stream traverse the NetScreen device.

The SIP ALG monitors SIP transactions and dynamically creates and manages pinholes based on the information it extracts from these transactions. The NetScreen SIP ALG supports all SIP methods and responses (see “SIP Request Methods” on page 157 and “Classes of SIP Responses” on page 157). You can allow SIP transactions to traverse the NetScreen firewall by creating a static policy that permits SIP service. This policy enables the NetScreen device to intercept SIP traffic and do one of the following actions: permit or deny the traffic or enable the SIP ALG to open pinholes to pass the media stream. The SIP ALG needs to open pinholes only for the SIP requests and responses that contain media information (SDP). For SIP messages that do not contain SDP, the NetScreen device simply lets them through.

The SIP ALG intercepts SIP messages that contain SDP, and using a parser, extracts the information it requires to create pinholes. The SIP ALG examines the SDP portion of the packet and a parser extracts information such as IP addresses and port numbers, which the SIP ALG records in a pinhole table. The SIP ALG uses the IP addresses and port numbers recorded in the pinhole table to open pinholes and allow media streams to traverse the NetScreen device.

Note: *NetScreen devices do not support encrypted SDP. If a NetScreen device receives a SIP message in which SDP is encrypted, the SIP ALG permits it through the firewall anyway, but generates a log message informing the user that it cannot process the packet. If SDP is encrypted, the SIP ALG cannot extract the information it needs from SDP to open pinholes. As a result, the media content that SDP describes cannot traverse the NetScreen device.*

6. We refer to a pinhole as the limited opening of a port to allow exclusive traffic.

SDP

An SDP session description is text-based and consists of a set of lines. It can contain session-level and media-level information. The session-level information applies to the whole session, while the media-level information applies to a particular media stream. An SDP session description always contains session-level information, which appears at the beginning of the description, and might contain media-level information[7], which comes after.

Of the many fields in the SDP description, two are particularly useful to the SIP ALG because they contain transport layer information. The two fields are the following:

- **c=** for connection information

  This field can appear at the session or media level. It displays in this format:

  c=<network type><address type><connection address>

  Currently, the NetScreen device supports only “IN” (for Internet) as the network type, “IP4” as the address type, and a unicast IP address[8] or domain name as the destination (connection) IP address.

  If the destination IP address is a unicast IP address, the SIP ALG creates pinholes using the IP address and port numbers specified in the media description field m=.

- **m=** for media announcement

  This field appears at the media level and contains the description of the media. It displays in this format:

  m=<media><port><transport><fmt list>

  Currently, the NetScreen device supports only “audio” as the media and “RTP” as the application layer protocol (transport). The port number indicates the destination of the media stream (and not the origin of the media stream). The format list (fmt list) provides information on the application layer protocol that the media uses.

  In this release of ScreenOS, the NetScreen device opens ports only for RTP and RTCP. Every RTP session has a corresponding RTCP[9] (Real-time Transport Control Protocol) session. Therefore, whenever a media stream uses RTP, the SIP ALG must reserve ports (create pinholes) for both RTP and RTCP traffic. By default, the port number for RTCP is one higher than the RTP port number.

7. In the SDP session description, the media-level information begins with the m= field.
8. Generally, the destination IP address can also be a multicast IP address, but NetScreen does not currently support multicast with SIP.
9. RTCP provides media synchronization and information about the members of the session and the quality of the communication.

Pinhole Creation

Both pinholes for the RTP and RTCP traffic share the same destination IP address. The IP address comes from the c= field in the SDP session description. Because the c= field can appear in either the session-level or media-level portion of the SDP session description, the parser determines the IP address based on the following rules (in accordance with SDP conventions):

- First, the SIP ALG parser verifies if there is a c= field containing an IP address in the media level. If there is one, the parser extracts that IP address and the SIP ALG uses it to create a pinhole for the media.
- If there is no c= field in the media level, the SIP ALG parser extracts the IP address from the c= field in the session level and the SIP ALG uses it to create a pinhole for the media. If the session description does not contain a c= field in either level, this indicates an error in the protocol stack and the NetScreen device drops the packet and logs the event.

The following lists the information the SIP ALG needs to create a pinhole. This information comes from the SDP session description and parameters on the NetScreen device.

- Protocol: UDP
- Source IP: unknown
- Source port: unknown
- Destination IP: The parser extracts the destination IP address from the c= field in the media or session level.
- Destination port: The parser extracts the destination port number for RTP from the m= field in the media level and calculates the destination port number for RTCP using this formula: *RTP port number + one*.
- Lifetime: This value indicates the length of time (in seconds), during which a pinhole is open to allow a packet through. A packet must go through the pinhole before the lifetime expires. When the lifetime expires, the SIP ALG removes the pinhole.

  When a packet goes through the pinhole within the lifetime period, immediately after, the SIP ALG removes the pinhole for the direction from which the packet came.

The following illustration describes a call setup between two SIP clients and how the SIP ALG creates pinholes to allow RTP and RTCP traffic. The illustration assumes that the NetScreen device has a policy that permits SIP, thus opening port 5060 for SIP signaling messages.

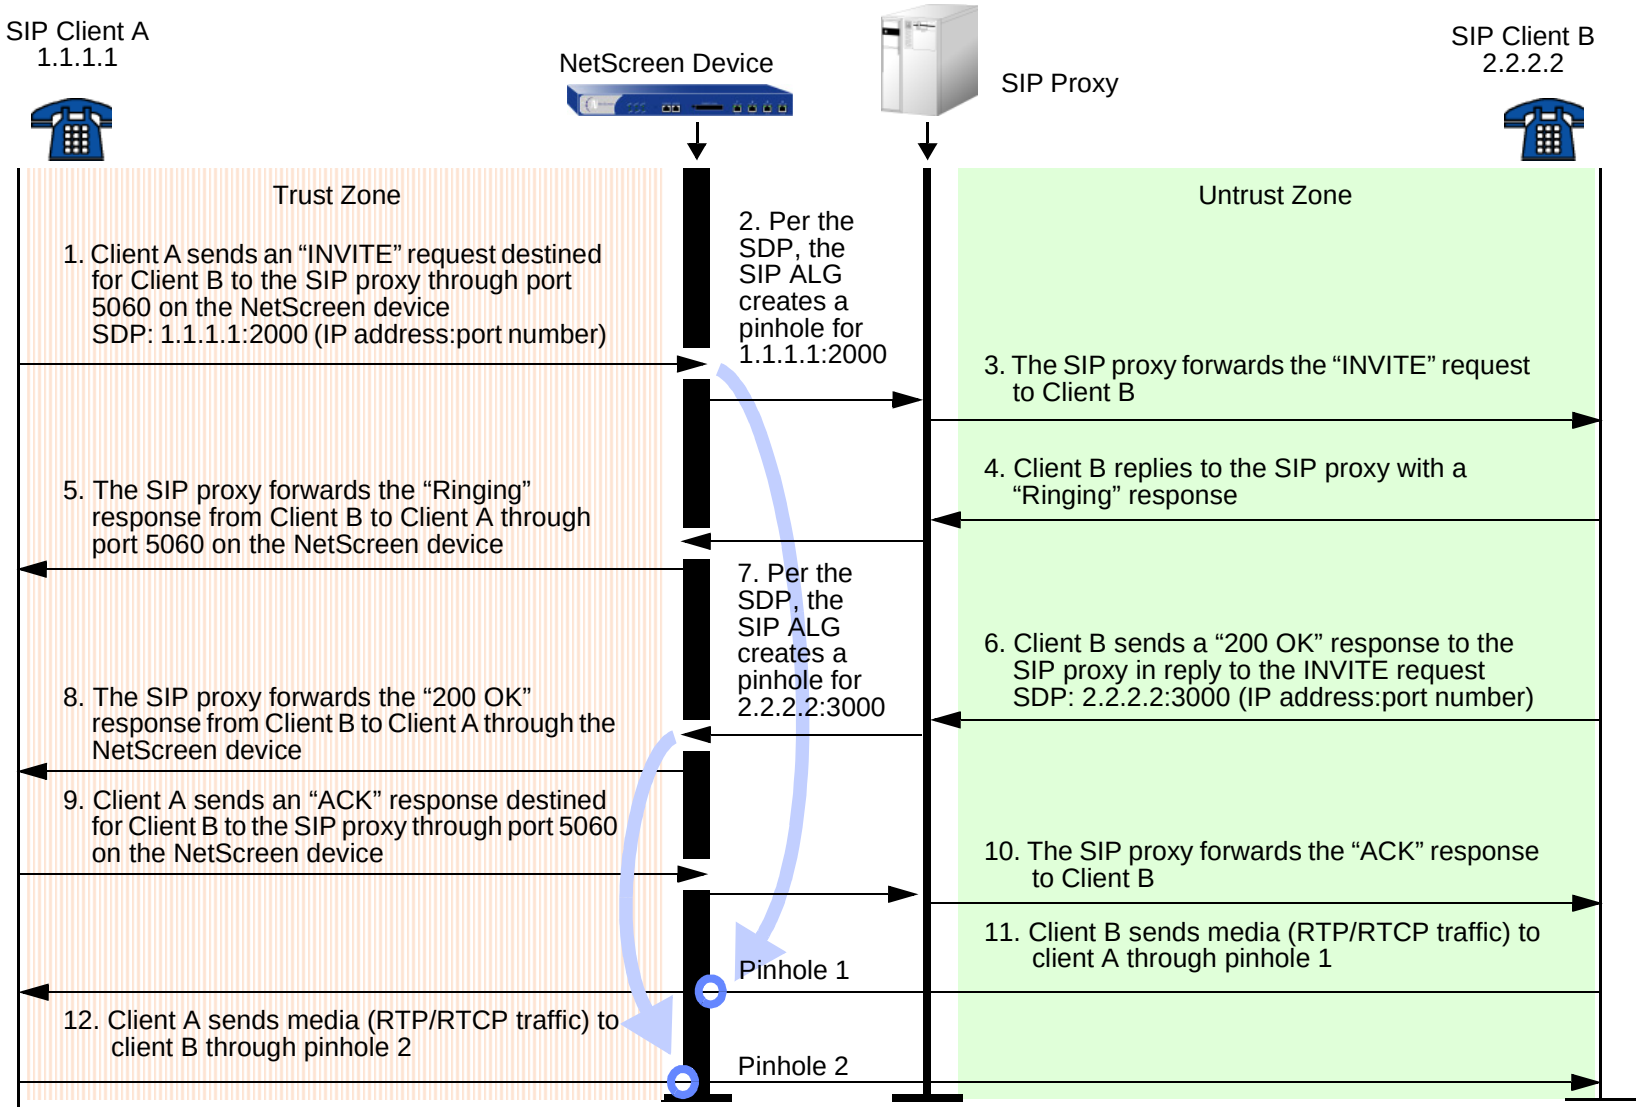

> ***Note:*** *The SIP ALG does not create pinholes for RTP and RTCP traffic when the destination IP address is 0.0.0.0, which indicates that the session is on hold. To put a session on hold, for example, during a telephone communication, a user (User A) sends the other user (User B) a SIP message in which the destination IP address is 0.0.0.0. Doing so indicates to User B not to send any media until further notice. If User B sends media anyway, the NetScreen device drops the packets.*

Session Inactivity Timeout

Typically a call ends when one of the clients sends a BYE or a CANCEL request. The SIP ALG intercepts the BYE or CANCEL request and removes all media sessions for that call. There could be reasons or problems preventing clients in a call from sending BYE or CANCEL requests, for example, a power failure. In this case, the call might go on indefinitely and indefinitely consume resources on the NetScreen device. The inactivity timeout feature helps the NetScreen device to monitor the liveliness of the call and terminate it if there is no activity for a specific period of time.

A call can have one or more voice channels. Each voice channel has two sessions (or two media streams), one for RTP and one for RTCP. When managing the sessions, the NetScreen device considers the sessions in each voice channel as one group. Settings such as the inactivity timeout apply to a group as opposed to each session.

There are two types of inactivity timeouts that determine the lifetime of a group:

- Signaling Inactivity Timeout: This parameter indicates the maximum length of time (in seconds) a call can remain active without any SIP signaling traffic. Each time a SIP signaling message occurs within a call, this timeout resets. The default setting is 43200 seconds (12 hours).
- Media Inactivity Timeout: This parameter indicates the maximum length of time (in seconds) a call can remain active without any media (RTP or RTCP) traffic within a group. Each time a RTP or RTCP packet occurs within a call, this timeout resets. The default setting is 120 seconds.

If either of these timeouts expire, the NetScreen device removes all sessions for this call from its table, thus terminating the call.

Example: Creating a Policy to Permit SIP

In this example, you create two policies to allow bidirectional traffic. One policy permits SIP traffic from SIP Client B in the Untrust zone to SIP Client A in the Trust zone. The other policy permits SIP traffic from SIP Client A in the Trust zone to SIP Client B in the Untrust zone.

WebUI

1. Address

   Objects > Addresses > List > New: Enter the following, and then click **OK**:

   Address Name: Client-A

   IP Address/Domain Name: IP/Netmask: 1.1.1.1/32

   Zone: Trust

   Objects > Addresses > List > New: Enter the following, and then click **OK**:

   Address Name: Client-B

   IP Address/Domain Name: IP/Netmask: 2.2.2.2/32

   Zone: Untrust

2. Policies

Policies > (From: Untrust, To: Trust) > New: Enter the following, and then click **OK**:

Source Address:

New Address: 2.2.2.2/32

Destination Address:

New Address: 1.1.1.1/32

Service: SIP

Action: Permit

Policies > (From: Trust, To: Untrust) > New: Enter the following, and then click **OK**:

Source Address:

New Address: 1.1.1.1/32

Destination Address:

New Address: 2.2.2.2/32

Service: SIP

Action: Permit

CLI

1. Addresses

```
set address trust client-a 1.1.1.1/32
set address untrust client-b 2.2.2.2/32
```

2. Policies

```
set policy from untrust to trust 2.2.2.2 1.1.1.1 sip permit
set policy from trust to untrust 1.1.1.1 2.2.2.2 sip permit
save
```

Example: Signaling and Media Inactivity Timeouts

In this example, you configure the signaling inactivity timeout to 30,000 seconds and the media inactivity timeout to 90 seconds.

WebUI

> ***Note:*** *You cannot configure this feature using the WebUI.*

CLI

```
set sip signaling-inactivity-timeout 30000
set sip media-inactivity-timeout 90
save
```

Service Groups

A service group is a set of services that you have gathered together under one name. After you create a group containing several services, you can then apply services at the group level to policies, thus simplifying administration.

The NetScreen service group option has the following features:

- Each service book entry can be referenced by one or more service groups.
- Each service group can contain predefined and user-defined service book entries.

Service groups are subject to the following limitations:

- Service groups cannot have the same names as services; therefore, if you have a service named “FTP,” you cannot have a service group named “FTP.”
- If a service group is referenced in a policy, you can edit the group but you cannot remove it until you have first removed the reference to it in the policy.
- If a custom service book entry is deleted from the service book, the entry is also removed from all the groups in which it was referenced.
- One service group cannot contain another service group as a member.
- The all-inclusive service term “ANY” cannot be added to groups.
- A service can be part of only one group at a time.

Example: Creating a Service Group

In this example, you create a service group named grp1 that includes IKE, FTP, and LDAP services.

WebUI

Objects > Services > Groups > New: Enter the following group name, move the following services, and then click **OK**:

Group Name: grp1

Select **IKE** and use the **<<** button to move the service from the Available Members column to the Group Members column.

Select **FTP** and use the **<<** button to move the service from the Available Members column to the Group Members column.

Select **LDAP** and use the **<<** button to move the service from the Available Members column to the Group Members column.

CLI

```
set group service grp1
set group service grp1 add ike
set group service grp1 add ftp
set group service grp1 add ldap
save
```

Note: *If you try to add a service to a service group that does not exist, the NetScreen device creates the group. Also, ensure that groups referencing other groups do not include themselves in the reference list.*

Example: Modifying a Service Group

In this example, you change the members in the service group named grp1 that you created in “Example: Creating a Service Group” on page 168. You remove IKE, FTP, and LDAP services, and add HTTP, FINGER, and IMAP.

WebUI

Objects > Services > Groups > Edit (for grp1): Move the following services, and then click **OK**:

Select **IKE** and use the **>>** button to move the service from the Group Members column to the Available Members column.

Select **FTP** and use the **>>** button to move the service from the Group Members column to the Available Members column.

Select **LDAP** and use the **>>** button to move the service from the Group Members column to the Available Members column.

Select **HTTP** and use the **<<** button to move the service from the Available Members column to the Group Members column.

Select **Finger** and use the **<<** button to move the service from the Available Members column to the Group Members column.

Select **IMAP** and use the **<<** button to move the service from the Available Members column to the Group Members column.

CLI

```
unset group service grp1 clear
set group service grp1 add http
set group service grp1 add finger
set group service grp1 add imap
save
```

Example: Removing a Service Group

In this example, you delete the service group named “grp1”.

WebUI

Objects > Services > Groups: Click **Remove** (for grp1).

CLI

```
unset group service grp1
save
```

Note: *The NetScreen device does not automatically delete a group from which you have removed all members.*

DIP POOLS

A dynamic IP (DIP) pool is a range of IP addresses from which the NetScreen device can dynamically or deterministically take addresses to use when performing network address translation on the source IP address (NAT-src) in IP packet headers. (For information about deterministic source address translation, see "NAT-Src from a DIP Pool with Address Shifting" on page 267.) If the range of addresses in a DIP pool is in the same subnet as the interface IP address, the pool must exclude the interface IP address, router IP addresses, and any mapped IP (MIP) or virtual IP (VIP) addresses that might also be in that subnet. If the range of addresses is in the subnet of an extended interface, the pool must exclude the extended interface IP address.

There are three kinds of interfaces that you can link to Dynamic IP (DIP) pools: physical interfaces and subinterfaces for network and VPN traffic, and tunnel interfaces for VPN tunnels only.

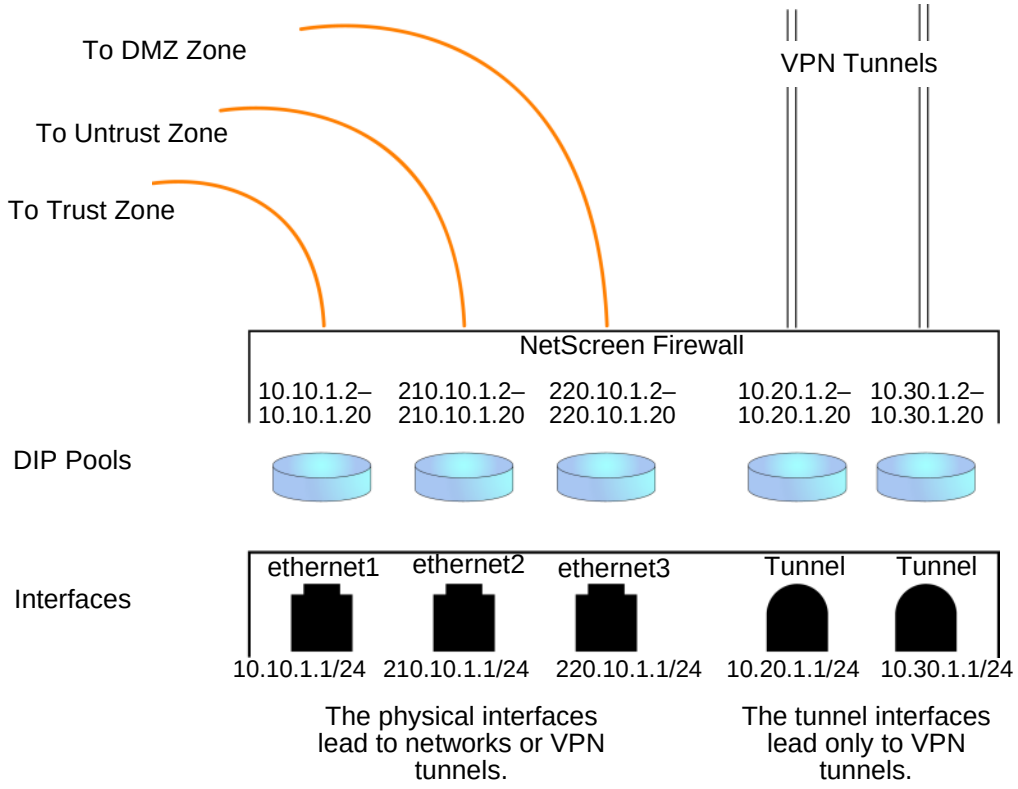

Port Address Translation

Using Port Address Translation (PAT), multiple hosts can share the same IP address, the NetScreen device maintaining a list of assigned port numbers to distinguish which session belongs to which host. With PAT enabled, up to ~64,500 hosts can share a single IP address.

Some applications, such as NetBIOS Extended User Interface (NetBEUI) and Windows Internet Naming Service (WINS), require specific port numbers and cannot function properly if PAT is applied to them. For such applications, you can specify not to perform PAT (that is, to use a fixed port) when applying DIP. For fixed-port DIP, the NetScreen device hashes the original host IP address and saves it in its host hash table, thus allowing the NetScreen device to associate the right session with each host.

Example: Creating a DIP Pool with PAT

In this example, you want to create a VPN tunnel for users at the local site to reach an FTP server at a remote site. However, the internal networks at both sites use the same private address space of 10.1.1.0/24. To solve the problem of overlapping addresses, you create a tunnel interface in the Untrust zone on the local NetScreen device, assign it IP address 10.10.1.1/24, and associate it with a DIP pool containing addresses 10.10.1.2–10.10.1.2—addresses in the neutral address space of 10.10.1.0/24. You enable port address translation for the DIP pool.

The admin at the remote site, must also create a tunnel interface with an IP address in a neutral address space, such as 10.20.2.1/24, and set up a Mapped IP (MIP) address to its FTP server, such as 10.20.2.5 to host 10.1.1.5.

Note: *This example includes only the configuration of the tunnel interface and its accompanying DIP pool. For a complete example showing all the configuration steps necessary for this scenario, see “VPN Sites with Overlapping Addresses” on page **5**-168.*

WebUI

Network > Interfaces > New Tunnel IF: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.1

Zone (VR): Untrust (trust-vr)

Fixed IP: (select)

IP Address / Netmask: 10.10.1.1/24

Network > Interfaces > Edit (for tunnel.1) > DIP > New: Enter the following, and then click **OK**:

ID: 5[10]

IP Address Range: 10.10.1.2 ~ 10.20.1.2

Port Translation: (select)

In the same subnet as the interface IP or its secondary IPs: (select)

CLI

```
set interface tunnel.1 zone untrust-tun
set interface tunnel.1 ip 10.10.1.1/24
set interface tunnel.1 dip 5 10.10.1.2 10.20.1.2
save
```

Note: *Because PAT is enabled by default, there is no argument for enabling it. To create the same DIP pool as defined above but without PAT (that is, with fixed port numbers), do the following:*

- **(WebUI)** *Network > Interfaces > Edit (for tunnel.1) > DIP > New: Clear the* **Port Translation** *check box, and then click* **OK**.
- **(CLI)** *set interface tunnel.1 dip 5 10.10.1.2 10.10.1.2 fix-port*

10. You can use the ID number displayed, which is the next available number sequentially, or type a different number.

Example: Modifying a DIP Pool

In this example, you change the range of addresses in an existing DIP pool (ID 5) from 10.20.1.2 – 10.20.1.2 to 10.20.1.2 – 10.20.1.10. This DIP pool is associated with tunnel.1. Note that to change the DIP pool range through the CLI, you must first remove (or unset) the existing dip pool and then create a new pool.

Note: *There are no policies using this particular DIP pool. If a policy uses a DIP pool, you must first delete the policy or modify it to not use the DIP pool before you can modify the DIP pool.*

WebUI

Network > Interfaces > Edit (for tunnel.1) > DIP > Edit (for ID 5): Enter the following, and then click **OK**:

IP Address Range: 10.20.1.2 ~ 10.20.1.10

CLI

```
unset interface tunnel.1 dip 5
set interface tunnel.1 dip 5 10.20.1.2 10.20.1.10
save
```

Sticky DIP Addresses

When a host initiates several sessions that match a policy requiring network address translation (NAT) and is assigned an address from a DIP pool with port translation enabled[11], the NetScreen device assigns a different source IP address for each session. Such random address assignment can be problematic for services that create multiple sessions that require the same source IP address for each session.

For example, it is important to have the same IP address for multiple sessions when using the AOL Instant Messaging (AIM) client. You create one session when you log in, and another for each chat. For the AIM server to verify that a new chat belongs to an authenticated user, it must match the source IP address of the login session with that of the chat session. If they are different—possibly because they were randomly assigned from a DIP pool during the NAT process—the AIM server rejects the chat session.

To ensure that the NetScreen device assigns the same IP address from a DIP pool to a host for multiple concurrent sessions, you can enable the “sticky” DIP address feature by entering the CLI command **set dip sticky**.

11. For DIP pools that do not perform port translation, the NetScreen device assigns one IP address for all concurrent sessions from the same host.

Extended Interface and DIP

If circumstances require that the source IP address in outbound firewall traffic be translated to an address in a different subnet from that of the egress interface, you can use the extended interface option. This option allows you to graft a second IP address and an accompanying DIP pool onto an interface that is in a different subnet. You can then enable NAT on a per-policy basis and specify the DIP pool built on the extended interface for the translation.

Example: Using DIP in a Different Subnet

In this example, two branch offices have leased lines to a central office. The central office requires them to use only the authorized IP addresses it has assigned them. However, the offices receive different IP addresses from their ISPs for Internet traffic. For communication with the central office, you use the extended interface option to configure the NetScreen device in each branch office to translate the source IP address in packets it sends to the central office to the authorized address. The authorized and assigned IP addresses for branch offices A and B are as follows:

	Assigned IP Address (from ISP) Used for Untrust Zone Physical Interface	**Authorized IP Address (from Central Office) Used for Untrust Zone Extended Interface DIP**
Office A	195.1.1.1/24	211.10.1.1/24
Office B	201.1.1.1/24	211.20.1.1/24

The NetScreen devices at both sites have a Trust zone and an Untrust zone. All security zones are in the trust-vr routing domain. You bind ethernet1 to the Trust zone and assign it IP address 10.1.1.1/24. You bind ethernet3 to the Untrust zone and give it the IP address assigned by the ISPs: 195.1.1.1/24 for Office A and 201.1.1.1/24 for Office B. You then create an extended interface with a DIP pool containing the authorized IP address on ethernet3:

- Office A: extended interface IP 211.10.1.10/24; DIP pool 211.10.1.1 – 211.10.1.1; PAT enabled
- Office B: extended interface IP 211.20.1.10/24; DIP pool 211.20.1.1 – 211.20.1.1; PAT enabled

You set the Trust zone interface in NAT mode. It uses the Untrust zone interface IP address as its source address in all outbound traffic except for traffic sent to the central office. You configure a policy to the central office that translates the source address to an address in the DIP pool in the extended interface. (The DIP pool ID number is 5. It contains one IP address, which, with port address translation, can handle sessions for ~64,500 hosts.) The MIP address that the central office uses for inbound traffic is 200.1.1.1, which you enter as “HQ” in the Untrust zone address book on each NetScreen device.

Note: *Leased lines connect branch offices A and B directly to the central office.*

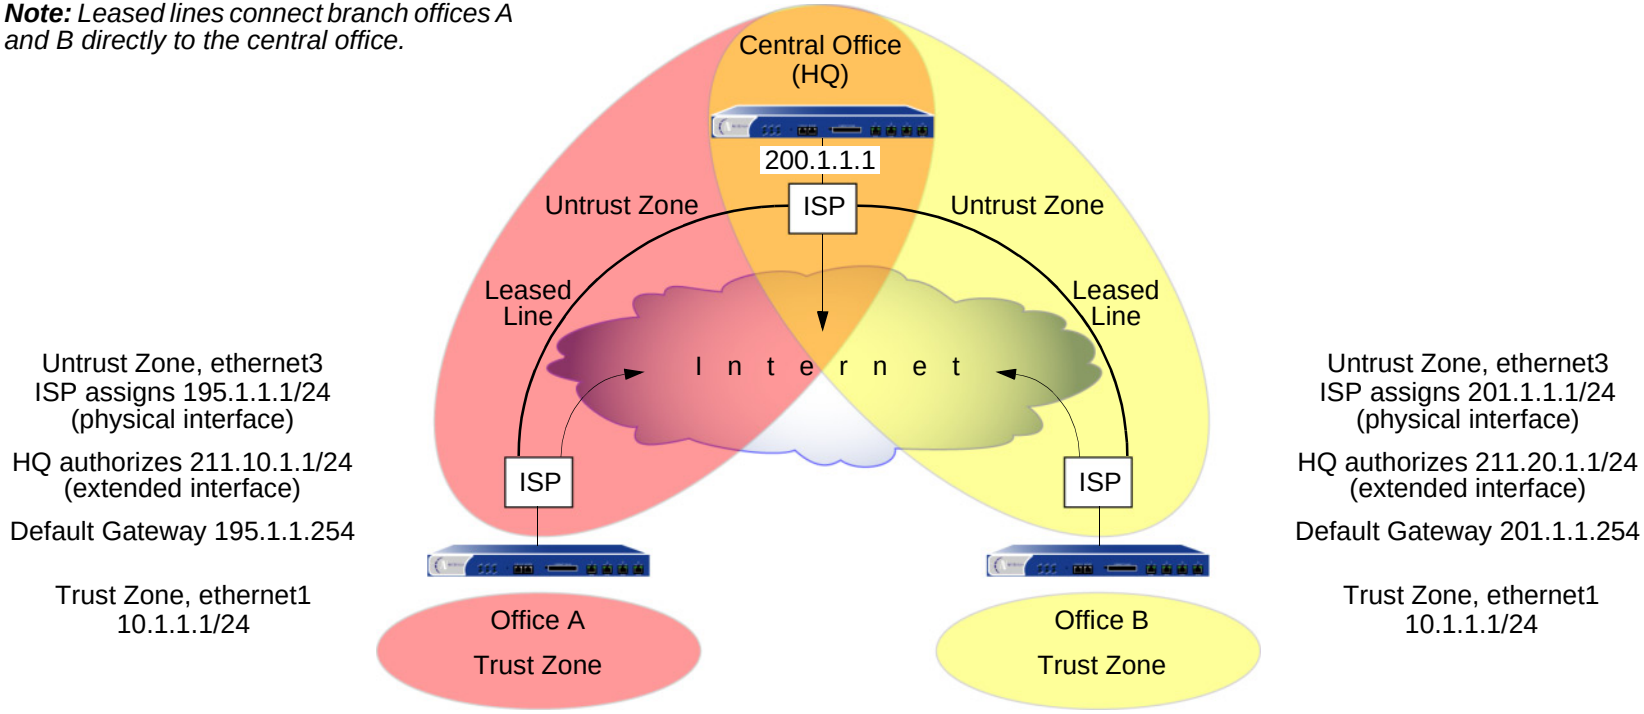

Note: *Each ISP must set up a route for traffic destined to a site at the end of a leased line to use that leased line. The ISPs route any other traffic they receive from a local NetScreen device to the Internet.*

WebUI (Branch Office A)

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **OK**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 195.1.1.1/24

Interface Mode: Route

Network > Interfaces > Edit (for ethernet3) > DIP > New: Enter the following, and then click **OK**:

ID: 5

IP Address Range: 211.10.1.1 ~ 211.10.1.1

Port Translation: (select)

Extended IP/Netmask: 211.10.1.10/255.255.255.0

2. Address

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: HQ

IP Address/Domain Name:

IP/Netmask: (select), 200.1.1.1/32

Zone: Untrust

3. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP address: 195.1.1.254

4. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Any

Service: ANY

Action: Permit

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), HQ

Service: ANY

Action: Permit

Position at Top: (select)

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

NAT:

Source Translation: (select)

(DIP on): 5 (211.10.1.1-211.10.1.1)/X-late

WebUI (Branch Office B)

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **OK**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 201.1.1.1/24

Interface Mode: Route

Network > Interfaces > Edit (for ethernet3) > DIP > New: Enter the following, and then click **OK**:

ID: 5

IP Address Range: 211.20.1.1 ~ 211.20.1.1

Port Translation: (select)

Extended IP/Netmask: 211.20.1.10/255.255.255.0

2. Address

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: HQ

IP Address/Domain Name:

IP/Netmask: (select), 200.1.1.1/32

Zone: Untrust

3. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP address: 201.1.1.254

4. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Any

Service: ANY

Action: Permit

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), HQ

Service: ANY

Action: Permit

Position at Top: (select)

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

NAT:

Source Translation: (select)

DIP On: (select), 5 (211.20.1.1-211.20.1.1)/X-late

CLI (Branch Office A)

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 195.1.1.1/24
set interface ethernet3 rout
set interface ethernet3 ext ip 211.10.1.10 255.255.255.0 dip 5 211.10.1.1
```

2. Address

```
set address untrust hq 200.1.1.1/32
```

3. Route

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 195.1.1.254
```

4. Policies

```
set policy from trust to untrust any any any permit
set policy top from trust to untrust any hq any nat src dip 5 permit
save
```

CLI (Branch Office B)

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 201.1.1.1/24
set interface ethernet3 route
set interface ethernet3 ext ip 211.20.1.10 255.255.255.0 dip 5 211.20.1.1
```

2. Address

```
set address untrust hq 200.1.1.1/32
```

3. Route

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 201.1.1.254
```

4. Policies

```
set policy from trust to untrust any any any permit
set policy top from trust to untrust any hq any nat src dip 5 permit
save
```

Loopback Interface and DIP

A loopback interface is a logical interface that is always in the up state as long as the device on which it resides is up[12]. You can create a pool of dynamic IP (DIP) addresses on a loopback interface so that it can be accessed by the group of interfaces belonging to its associated loopback interface group when performing source address translation. The addresses that the NetScreen device draws from such a DIP pool are in the same subnet as the loopback interface IP address, not in the subnet of any of the member interfaces. (Note that the addresses in the DIP pool must not overlap with the interface IP address or any MIP addresses also defined on the loopback interface.)

The primary application for putting a DIP pool on a loopback interface is to translate source addresses to the same address or range of addresses although different packets might use different egress interfaces.

Source Address Translation Using a DIP Pool on a Loopback Interface

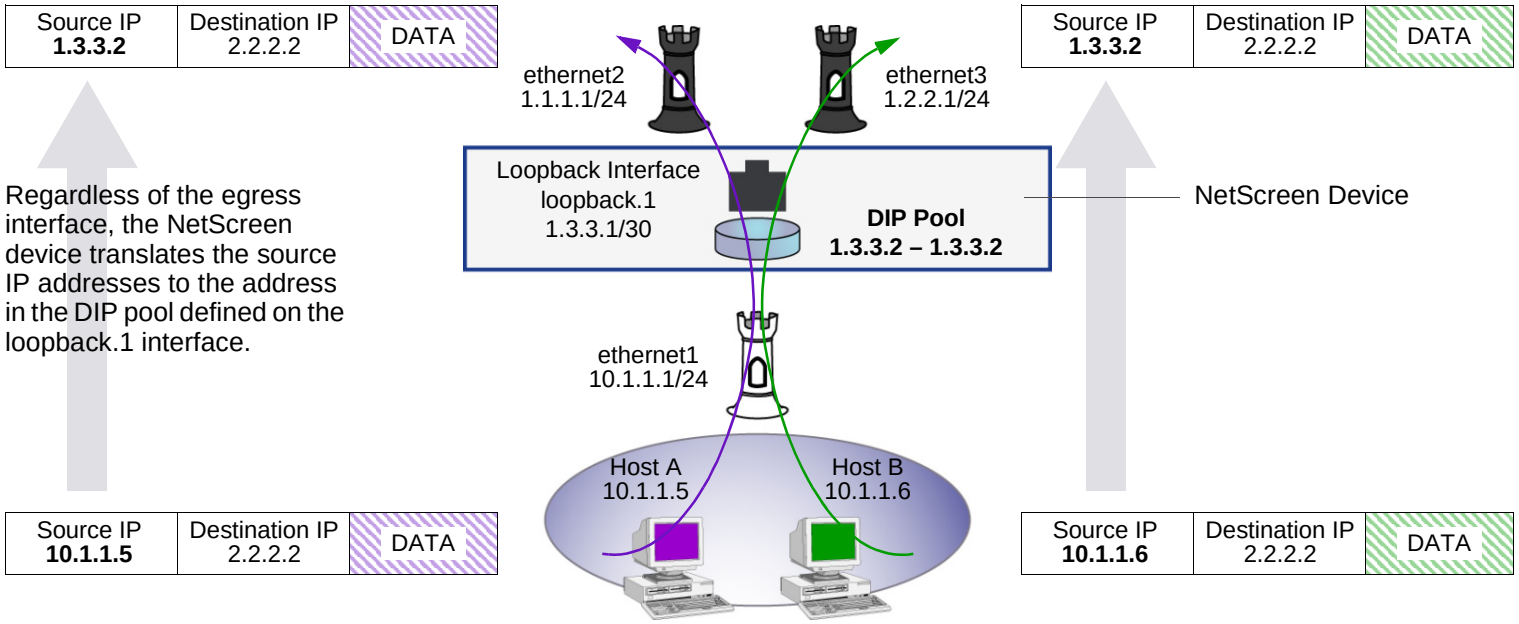

12. For information about loopback interfaces, see "Loopback Interfaces" on page 86.

Example: DIP on a Loopback Interface

In this example, the NetScreen device receives the following IP addresses for two Untrust zone interfaces from different Internet service providers (ISPs): ISP-1 and ISP-2:

- ethernet2, 1.1.1.1/24, ISP-1
- ethernet3, 1.2.2.1/24, ISP-2

You bind these interfaces to the Untrust zone and then assign them the above IP addresses. You also bind ethernet1 to the Trust zone and assign it IP address 10.1.1.1/24.

You want the NetScreen device to translate the source address in outbound traffic from the Trust zone to a remote office in the Untrust zone. The translated address must be the same IP address (1.3.3.2) because the remote office has a policy permitting inbound traffic only from that IP address. You have previously obtained the public IP addresses 1.3.3.1 and 1.3.3.2 and have notified both ISPs that you are using these addresses in addition to the addresses that they assign the device.

You configure a loopback interface loopback.1 with the IP address 1.3.3.1/30 and a DIP pool of 1.3.3.2 – 1.3.3.2 on that interface. The DIP pool has ID number 10. You then make ethernet1 and ethernet2 members of the loopback group for loopback.1.

You define an address for the remote office named “r-office” with IP address 2.2.2.2/32. You also define default routes for both ethernet1 and ethernet2 interfaces pointing to ISP-1 and ISP-2’s routers respectively.

You define routes to two gateways for outbound traffic to use. Because you do not prefer one route over the other, you do not include any metrics in the routes. Outbound traffic might follow either route[13].

Finally, you create a policy applying source network address translation (NAT-src) to outbound traffic to the remote office. The policy references DIP pool ID 10.

13. To indicate a route preference, include metrics in both routes, giving your preferred route a higher metric—that is, a value closer to 1.

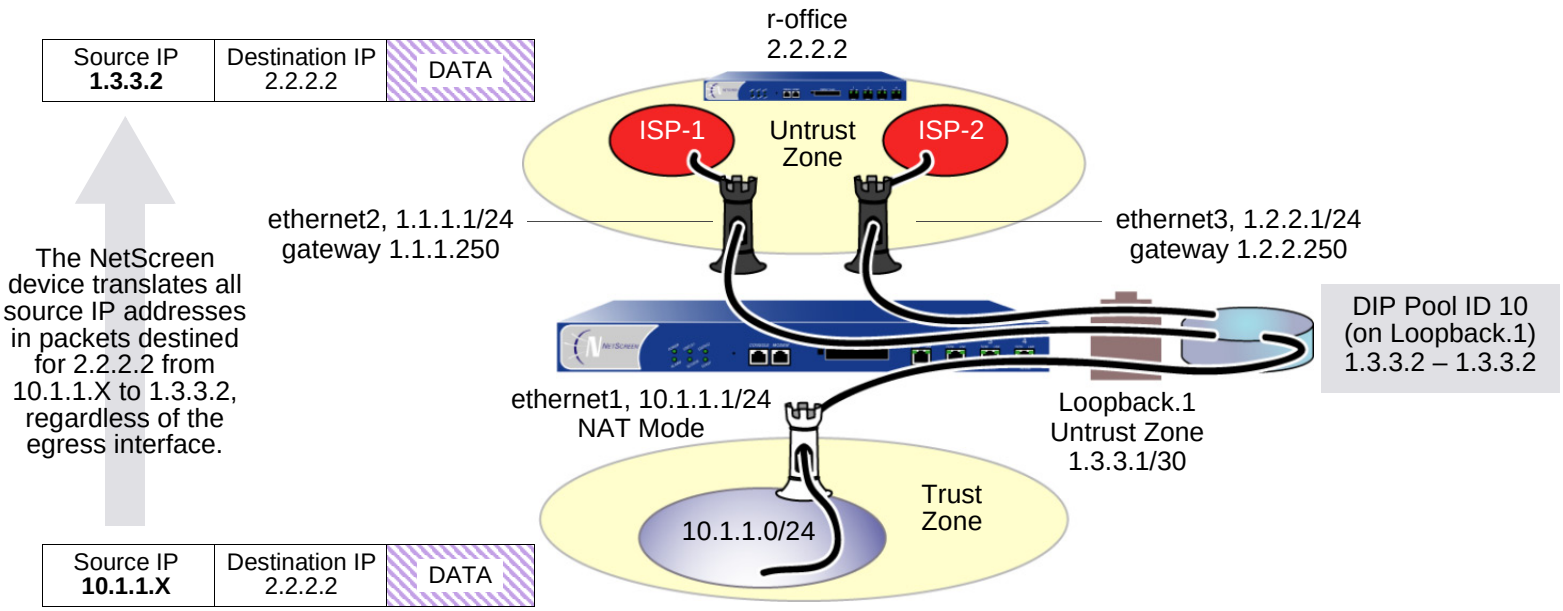

WebUI

1. Interfaces

Network > Interfaces > New Loopback IF: Enter the following, and then click **OK**:

Interface Name: loopback.1

Zone: Untrust (trust-vr)

IP Address/Netmask: 1.3.3.1/30

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **OK**:

As member of loopback group: loopback.1

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet2): Enter the following, and then click **OK**:

As member of loopback group: loopback.1

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

Interface Mode: Route

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.2.2.1/24

Interface Mode: Route

2. DIP Pool

Network > Interfaces > Edit (for loopback.1) > DIP > New: Enter the following, and then click **OK**:

ID: 5

IP Address Range: 1.3.3.2 ~ 1.3.3.2

Port Translation: (select)

3. Address

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: r-office

IP Address/Domain Name:

IP/Netmask: (select), 2.2.2.2/32

Zone: Untrust

4. Routes

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet2

Gateway IP address: 1.1.1.250

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP address: 1.2.2.250

5. Policy

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), r-office

Service: ANY

Action: Permit

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

NAT:

Source Translation: (select)

DIP On: (select), 10 (1.3.3.2-1.3.3.2)/port-xlate

CLI

1. Interfaces

```
set interface loopback.1 zone untrust
set interface loopback.1 ip 1.3.3.1/30

set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet2 zone untrust
set interface ethernet2 ip 1.1.1.1/24
set interface ethernet2 loopback-group loopback.1

set interface ethernet3 zone untrust
set interface etherenet3 ip 1.2.2.1/24
set interface ethernet3 loopback-group loopback.1
```

2. DIP Pool

```
set interface loopback.1 dip 10 1.3.3.2 1.3.3.2
```

3. Address

```
set address untrust r-office 2.2.2.2/32
```

4. Routes

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet2 gateway 1.1.1.250
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.2.2.250
```

5. Policy

```
set policy from trust to untrust any r-office any nat src dip-id 10 permit
save
```

DIP Groups

When you group two NetScreen devices into a redundant cluster to provide high availability (HA) in an active/active configuration, both devices share the same configuration and both process traffic simultaneously. A problem can arise when you define a policy to perform network address translation (NAT) using a DIP pool located on one VSI. Because that VSI is active only on the NetScreen device acting as the master of the VSD group to which the VSI is bound, any traffic sent to the other NetScreen device—the one acting as the backup of that VSD group—cannot use that DIP pool and is dropped.

Problematic use of a DIP pool in a policy when in an NSRP cluster:
set policy name out-nat from trust to untrust any any any nat src dip-id 7 permit

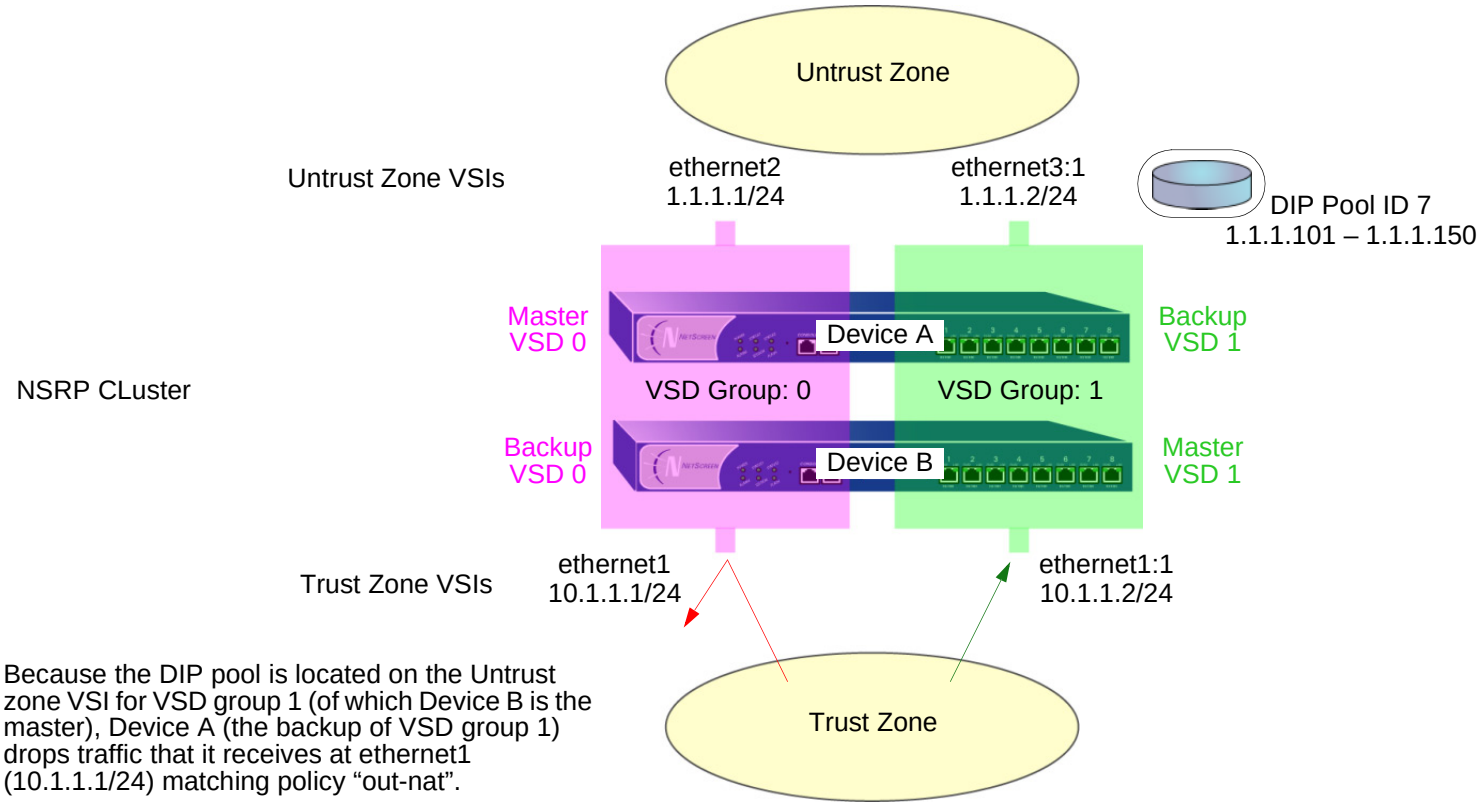

Because the DIP pool is located on the Untrust zone VSI for VSD group 1 (of which Device B is the master), Device A (the backup of VSD group 1) drops traffic that it receives at ethernet1 (10.1.1.1/24) matching policy “out-nat”.

To solve this problem, you can create two DIP pools—one on the Untrust zone VSI for each VSD group—and combine the two DIP pools into one DIP group, which you reference in the policy. Each VSI uses its own VSD pool even though the policy specifies the DIP group.

Recommended use of a DIP group in a policy when in an NSRP cluster:
set policy name out-nat from trust to untrust any any any nat dip-id 9 permit

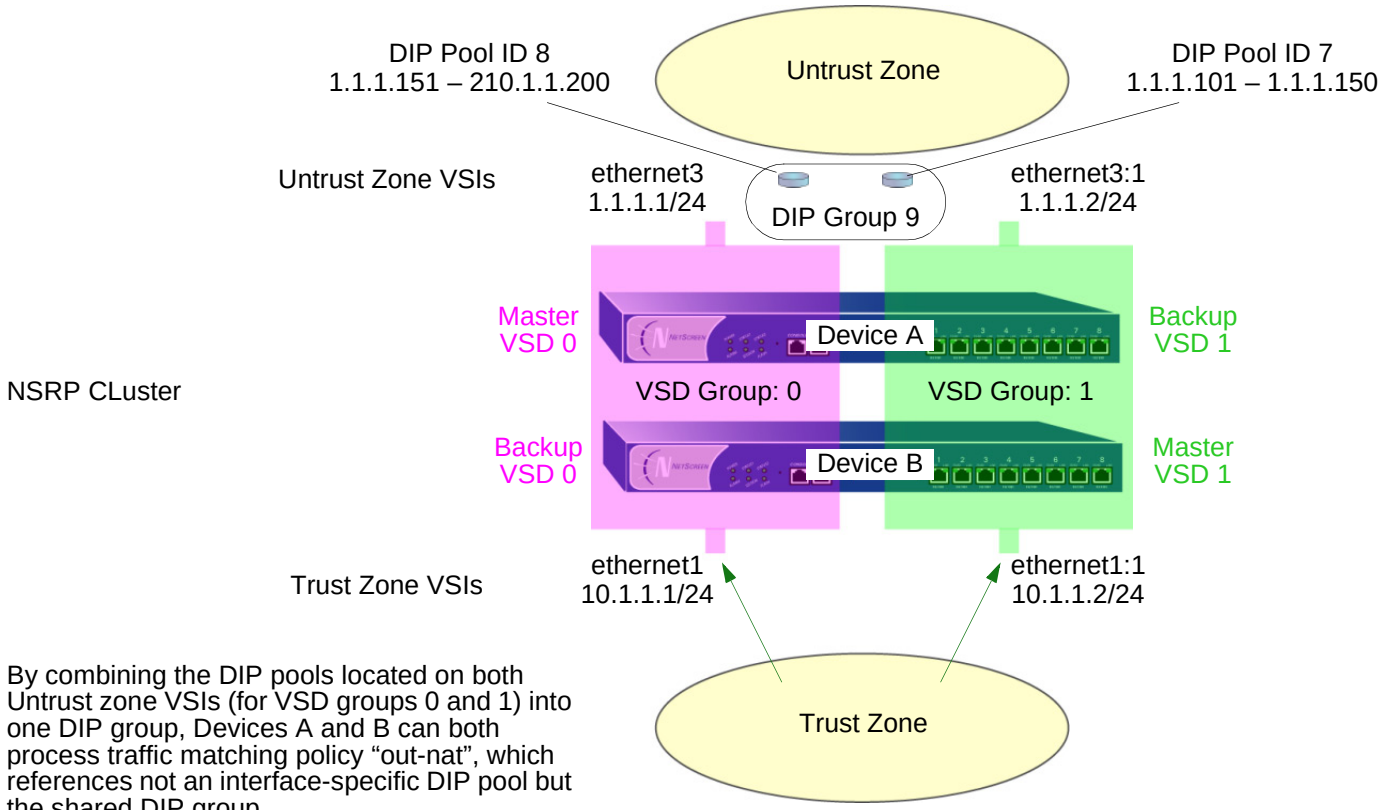

By combining the DIP pools located on both Untrust zone VSIs (for VSD groups 0 and 1) into one DIP group, Devices A and B can both process traffic matching policy “out-nat”, which references not an interface-specific DIP pool but the shared DIP group.

Note: *For more information about setting up NetScreen devices for HA, see Volume 8, “High Availability”.*

Example: DIP Group

In this example, you provide NAT services on two NetScreen devices (Devices A and B) in an active/active HA pair.

You create two DIP pools—DIP 5 (1.1.1.20 – 1.1.1.29) on ethernet3 and DIP 6 (1.1.1.30 – 1.1.1.39) on ethernet3:1. You then combine them into a DIP group identified as DIP 7, which you reference in a policy.

The VSIs for VSD groups 0 and 1 are as follows:

- Untrust zone VSI ethernet3 1.1.1.1/24 (VSD group 0)
- Untrust zone VSI ethernet3:1 1.1.1.2/24 (VSD group 1)
- Trust zone VSI ethernet1 10.1.1.1/24 (VSD group 0)
- Trust zone VSI ethernet1:1 10.1.1.1/24 (VSD group 1)

This example assumes that you have already set up Devices A and B in an NSRP cluster, created VSD group 1 (NetScreen automatically creates VSD group 0 when you put a device in an NSRP cluster), and configured the above interfaces. (For information about configuring NetScreen devices for NSRP, refer to *Volume 8, "High Availability"*.)

WebUI

1. DIP Pools

Network > Interfaces > Edit (for ethernet3) > DIP > New: Enter the following, and then click **OK**:

ID: 5

IP Address Range: 1.1.1.20 – 1.1.1.29

Port Translation: (select)

Network > Interfaces > Edit (for ethernet3:1) > DIP > New: Enter the following, and then click **OK**:

ID: 6

IP Address Range: 1.1.1.30 – 1.1.1.39

Port Translation: (select)

Note: *At the time of this release, you can only define a DIP group through the CLI.*

2. Policy

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Any

Service: ANY

Action: Permit

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

NAT:

Source Translation: (select)

DIP On: (select), 7

CLI

1. DIP Pools

```
set interface ethernet3 dip 5 1.1.1.20 1.1.1.29
set interface ethernet3:1 dip 6 1.1.1.30 1.1.1.39
```

2. DIP Groups

```
set dip group 7 member 5
set dip group 7 member 6
```

3. Policy

```
set policy from trust to untrust any any any nat src dip-id 7 permit
save
```

SCHEDULES

A schedule is a configurable object that you can associate with one or more policies to define when they are in effect. Through the application of schedules, you can control network traffic flow and enforce network security.

When you define a schedule, enter values for the following parameters:

Schedule Name: The name that appears in the Schedule drop-down list in the Policy Configuration dialog box. Choose a descriptive name to help you identify the schedule. The name must be unique and is limited to 19 characters.

Comment: Any additional information that you want to add.

Recurring: Enable this when you want the schedule to repeat on a weekly basis.

Start and End Times: You must configure both a start time and an end time. You can specify up to two time periods within the same day.

Once: Enable this when you want the schedule to start and end only once.

mm/dd/yyyy hh:mm: You must enter both start and stop dates and times.

Example: Recurring Schedule

In this example, there is a short-term employee named Tom who is using the company's Internet access for personal pursuits after work. You create a schedule for non-business hours that you can then associate with a policy to deny outbound TCP/IP traffic from that worker's computer (10.1.1.5/32) outside of regular business hours.

WebUI

1. Schedule

Objects > Schedules > New: Enter the following, and then click **OK**:

Schedule Name: After Hours

Comment: For non-business hours

Recurring: (select)

Period 1:

Week Day	Start Time	End Time
Sunday	00:00	23:59
Monday	00:00	06:00
Tuesday	00:00	06:00
Wednesday	00:00	06:00
Thursday	00:00	06:00
Friday	00:00	06:00
Saturday	00:00	23:59

Period 2:

Week Day	Start Time	End Time
Sunday	17:00	23:59
Monday	17:00	23:59
Tuesday	17:00	23:59
Wednesday	17:00	23:59
Thursday	17:00	23:59
Friday	17:00	23:59
Saturday	17:00	23:59

2. Address

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Tom

Comment: Temp

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.5/32

Zone: Trust

3. Policy

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Name: No Net

Source Address:

Address Book Entry: (select), Tom

Destination Address:

Address Book Entry: (select), Any

Service: HTTP

Action: Deny

Schedule: After Hours

CLI

1. Schedule

```
set schedule “after hours” recurrent sunday start 00:00 stop 23:59
set schedule “after hours” recurrent monday start 00:00 stop 06:00 start 17:00
    stop 23:59
set schedule “after hours” recurrent tuesday start 00:00 stop 06:00 start 17:00
    stop 23:59
set schedule “after hours” recurrent wednesday start 00:00 stop 06:00 start
    17:00 stop 23:59
set schedule “after hours” recurrent thursday start 00:00 stop 06:00 start
    17:00 stop 23:59
set schedule “after hours” recurrent friday start 00:00 stop 06:00 start 17:00
    stop 23:59
set schedule “after hours” recurrent saturday start 00:00 stop 23:59 comment
    “for non-business hours”
```

2. Address

```
set address trust tom 10.1.1.5/32 “temp”
```

3. Policy

```
set policy from trust to untrust tom any http deny schedule “after hours”
save
```

7

Policies

The default behavior of a NetScreen device is to deny all traffic between security zones (interzone traffic)[1] and—except for traffic within the Untrust zone—allow all traffic between interfaces bound to the same zone (intrazone traffic). To permit selected interzone traffic to cross a NetScreen device you must create interzone policies that override the default behavior. Similarly, to prevent selected intrazone traffic from crossing a NetScreen device, you must create intrazone policies.

This chapter describes what policies do and how the various elements that comprise a policy are related. It is divided into the following sections:

- "Basic Elements" on page 199
- "Three Types of Policies" on page 200
  - "Interzone Policies" on page 200
  - "Intrazone Policies" on page 201
  - "Global Policies" on page 201
- "Policy Set Lists" on page 202
- "Policies Defined" on page 203
  - "Policies and Rules" on page 203
  - "Anatomy of a Policy" on page 205
- "Policies Applied" on page 216
  - "Viewing Policies" on page 216
  - "Creating Policies" on page 217
  - "Entering a Policy Context" on page 235

1. By default, the NetScreen-5XP and NetScreen-5XT permit traffic from the Trust zone to the Untrust zone.

- "Multiple Items per Policy Component" on page 236
- "Address Negation" on page 237
- "Modifying and Disabling Policies" on page 241
- "Policy Verification" on page 242
- "Reordering Policies" on page 243
- "Removing a Policy" on page 244

BASIC ELEMENTS

A policy permits, denies, or tunnels[2] specified types of traffic unidirectionally between two points. The type of traffic (or "service"), the location of the two endpoints, and the invoked action compose the basic elements of a policy. Although there can be other components, the required elements, which together constitute the core section of a policy, are as follows:

- Direction – The direction of traffic between two security zones (from a source zone to a destination zone)
- Source address – The address from which traffic initiates
- Destination address – The address to which traffic is sent
- Service – The type of traffic transmitted
- Action – The action that the NetScreen device performs when it receives traffic meeting the first four criteria: permit, deny, or tunnel

For example, the policy stated in the following CLI command permits FTP traffic from any address in the Trust zone to an FTP server named "server1" in the DMZ zone:

set policy from trust to untrust any server1 ftp permit

- Direction: **from trust to untrust** (that is, from the Trust zone to the Untrust zone)
- Source Address: **any** (that is, any address in the Trust zone. The term "any" stands for a predefined address that applies to any address in a zone)
- Destination Address: **server1** (a user-defined address in the Untrust zone address book)
- Service: **ftp** (File Transfer Protocol)
- Action: **permit** (that NetScreen device permits this traffic to traverse its firewall)

2. The "tunnel" action—(VPN or L2TP tunnel)—contains the concept of "permit" implicitly.

THREE TYPES OF POLICIES

You can control the flow of traffic through the following three kinds of policies:

- Through the creation of interzone policies, you can regulate the kind of traffic that you want to permit from one security zone to another.
- Through the creation of intrazone policies, you can also control the kind of traffic that you want to permit to cross interfaces bound to the same zone.
- Through the creation of global policies, you can regulate traffic between addresses, regardless of their security zones.

Interzone Policies

Interzone policies provide traffic control between security zones. You can set interzone policies to permit, deny, or tunnel traffic from one zone to another. Using stateful inspection techniques, a NetScreen device maintains a table of active TCP sessions and active UDP “pseudo” sessions so that it can allow replies to service requests. For example, if you have a policy allowing HTTP requests from host A in the Trust zone to server B in the Untrust zone, when the NetScreen device receives HTTP replies from server B to host A, the NetScreen device checks the received packet against its table. Finding the packet to be a reply to an approved HTTP request, the NetScreen device allows the packet from server B in the Untrust zone to cross the firewall to host A in the Trust zone. To permit traffic initiated by server B to host A (not just replies to traffic initiated by host A), you must create a second policy from server B in the Untrust zone to host A in the Trust zone.

set policy from trust to untrust “host A” “server B” http permit

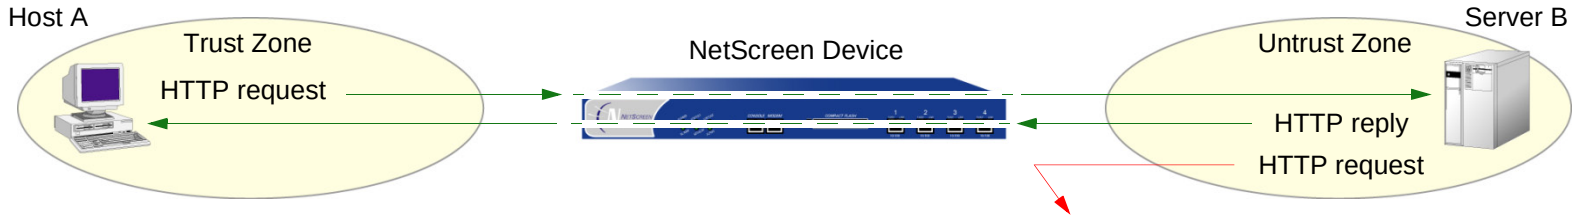

Note: *The NetScreen device rejects the HTTP request from server B because there is no policy permitting it.*

Intrazone Policies

Intrazone policies provide traffic control between interfaces bound to the same security zone. The source and destination addresses are in the same security zone, but reached via different interfaces on the NetScreen device. Like interzone policies, intrazone policies control traffic flowing unidirectionally. To allow traffic initiated at either end of a data path, you must create two policies—one policy for each direction.

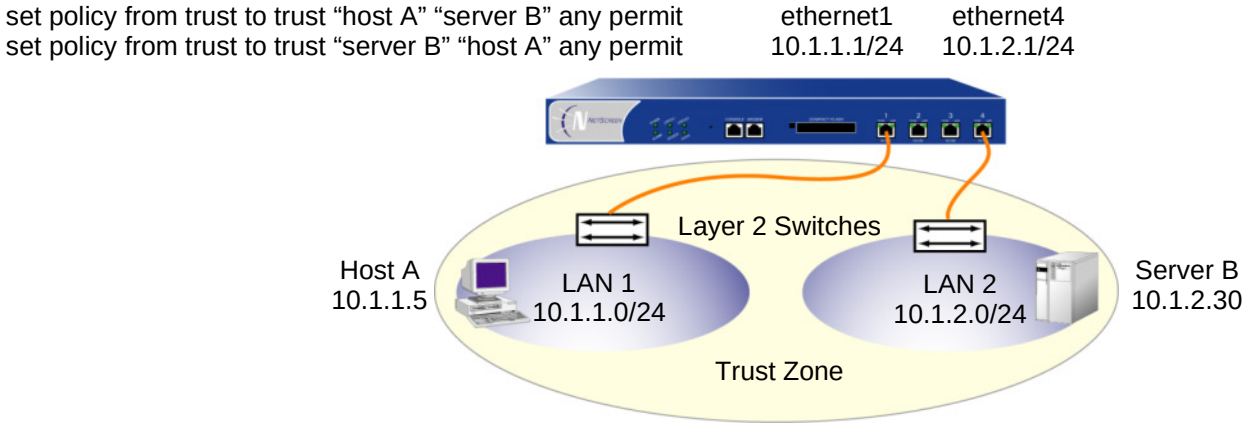

Intrazone policies do not support VPN tunnels or source network address translation (NAT-src) when it is set at the interface level (**set interface** *interface* **nat**). However, intrazone policies do support policy-based NAT-src and NAT-dst. They also support destination address translation when the policy references a mapped IP (MIP) as the destination address. (For information about NAT-src, NAT-dst, and MIPs, see “Address Translation” on page 245.)

Global Policies

Unlike interzone and intrazone policies, global policies do not reference specific source and destination zones. Global policies reference user-defined Global zone addresses or the predefined Global zone address “any”. These addresses can span multiple security zones. For example, if you want to provide access to or from multiple zones, you can create a global policy with the Global zone address “any”, which encompasses all addresses in all zones.

Note: *At the time of this release, global policies do not support source network address translation (NAT-src), VPN tunnels, or Transparent mode. You can, however, specify a MIP or VIP as the destination address in a global policy.*

POLICY SET LISTS

A NetScreen device maintains three different policy set lists, one for each of the following kinds of policies:

- Interzone policies
- Intrazone policies
- Global policies

When the NetScreen device receives a packet initiating a new session, the device notes the ingress interface, and thereby learns the source zone to which that interface is bound. The NetScreen device then performs a route lookup to determine the egress interface, and thus determines the destination zone to which that interface is bound. Using the source and destination zones, the NetScreen device can perform a policy lookup, consulting the policy set lists in the following order:

1. If the source and destination zones are different, the NetScreen device performs a policy lookup in the interzone policy set list.

   (or)

   If the source and destination zones are the same, the NetScreen device performs a policy lookup in the intrazone policy set list.

2. If the NetScreen device performs the interzone or intrazone policy lookup and does not find a match, the NetScreen device then checks the global policy set list for a match.
3. If the NetScreen device performs the interzone and global policy lookups and does not find a match, the NetScreen device then applies the default permit/deny policy to the packet: **unset/set policy default-permit-all**.

   (or)

   If the NetScreen device performs the intrazone and global policy lookups and does not find a match, the NetScreen device then applies the intrazone blocking setting for that zone to the packet: **unset/set zone** *zone* **block**.

The NetScreen device searches each policy set list from top to bottom. Therefore, you must position more specific policies above less specific policies in the list. (For information on policy order, see “Reordering Policies” on page 243.)

POLICIES DEFINED

A firewall provides a network boundary with a single point of entry and exit. Because all traffic must pass through this point, you can screen and direct that traffic through the implementation of policy set lists—for interzone policies, intrazone policies, and global policies.

Policies allow you to permit, deny, encrypt and decrypt, authenticate, prioritize, schedule, filter, and monitor the traffic attempting to cross from one security zone to another. You decide which users and what data can enter and exit, and when and where they can go.

Note: *For NetScreen devices that support virtual systems, policies set in the root system do not affect policies set in virtual systems.*

Policies and Rules

A single user-defined policy produces one or more logical rules internally, and each logical rule consists of a set of components—source address, destination address, and service. The components consume memory resources. The logical rules that reference the components do not.

Depending on the use of multiple entries or groups for the source address, destination address, and service components in a policy, the number of logical rules can be much larger than is readily apparent from the creation of the single policy. For example, the following policy produces 125 logical rules:

1 policy: 5 source addresses x 5 destination addresses x 5 services = 125 logical rules

However, the NetScreen device does not duplicate components for each logical rule. The rules make use of the same set of components in various combinations. For example, the above policy that produces 125 logical rules results in only 15 components:

5 source addresses + 5 destination addresses + 5 services = 15 components

These 15 components combine in various ways to produce the 125 logical rules generated by the single policy. By allowing multiple logical rules to use the same set of components in different combinations, the NetScreen device consumes far fewer resources than if each logical rule had a one-to-one relationship with its components.

Because the installation time of a new policy is proportional to the number of components that the NetScreen device adds, removes, or modifies, policy installation becomes faster with fewer components. Also, by allowing a large number of logical rules to share a small set of components, NetScreen allows you to create more policies—and the NetScreen device to create more rules—than would be possible if each rule required dedicated components.

Anatomy of a Policy

A policy must contain the following elements:

- ID (automatically generated, but can be user-defined in the CLI)
- Zones (source and destination)
- Addresses (source and destination)
- Services
- Action (permit, deny, tunnel)

A policy can also contain the following elements:

- Application
- Name
- VPN Tunneling
- L2TP Tunneling
- Deep Inspection
- Placement at the Top of the Policy List
- Source Address Translation
- Destination Address Translation
- User Authentication
- HA Session Backup
- URL Filtering
- Logging
- Counting
- Traffic Alarm Threshold
- Schedules
- Antivirus Scanning
- Traffic Shaping

The remainder of this section examines each of the above elements in turn.

ID

Every policy has an ID number, whether you define one or the NetScreen device automatically assigns it. You can only define an ID number for a policy through the set policy command in the CLI: **set policy id** *number* … After you know the ID number, you can enter the policy context to issue further commands to modify the policy. (For more information about policy contexts, see "Entering a Policy Context" on page 235.)

Zones

A zone can be a segment of network space to which security measures are applied (a security zone), a logical segment to which a VPN tunnel interface is bound (a tunnel zone), or either a physical or logical entity that performs a specific function (a function zone). A policy allows traffic to flow between two security zones (interzone policy) or between two interfaces bound to the same zone (intrazone policy). (For more information, see "Zones" on page 45, "Interzone Policies" on page 200, and "Intrazone Policies" on page 201.)

Addresses

Addresses are objects that identify network devices such as hosts and networks by their location in relation to the firewall—in one of the security zones. Individual hosts are specified using the mask 255.255.255.255, indicating that all 32 bits of the address are significant. Networks are specified using their subnet mask to indicate which bits are significant. To create a policy for specific addresses, you must first create entries for the relevant hosts and networks in the address book.

You can also create address groups and apply policies to them as you would to other address book entries. When using address groups as elements of policies, be aware that because the NetScreen device applies the policy to each address in the group, the number of available internal logical rules and the components that comprise those rules can become depleted more quickly than expected. This is a danger especially when you use address groups for both the source and destination. (For more information, see "Policies and Rules" on page 203.)

Services

Services are objects that identify application protocols using layer 4 information such as standard and accepted TCP and UDP port numbers for application services like Telnet, FTP, SMTP, and HTTP. The ScreenOS includes predefined core Internet services. Additionally, you can define custom services.

You can define policies that specify which services are permitted, denied, encrypted, authenticated, logged, or counted.

Action

An action is an object that describes what the firewall does to the traffic it receives.

- **Permit** allows the packet to pass the firewall.
- **Deny** blocks the packet from traversing the firewall.
- **Tunnel** encapsulates outgoing IP packets and decapsulates incoming IP packets. For an IPSec VPN tunnel, specify which VPN tunnel to use. For an L2TP tunnel, specify which L2TP tunnel to use. For L2TP-over-IPSec, specify both an IPSec VPN tunnel and an L2TP tunnel[3].

The NetScreen device applies the specified action on traffic that matches the previously presented criteria: zones (source and destination), addresses (source and destination), and service.

Application

The application option specifies the Layer 7 application that maps to the Layer 4 service that you reference in a policy. A predefined service already has a mapping to a Layer 7 application. However, for custom services, you must link the service to an application explicitly, especially if you want the policy to apply an application layer gateway (ALG[4]) or Deep Inspection to the custom service.

Applying an ALG to a custom service, involves the following two steps:

- Define a custom service with a name, timeout value, transport protocol, and source and destination ports
- When configuring a policy, reference that service and the application type for the ALG that you want to apply

For information about applying Deep Inspection to a custom service, see “Mapping Custom Services to Applications” on page **4**-152.

3. For L2TP-over-IPSec, the source and destination addresses for the IPSec VPN tunnel must be the same as those for the L2TP tunnel.
4. NetScreen supports ALGs for numerous services, including DNS, FTP, H.323, HTTP, RSH, SIP, telnet, and TFTP.

Name

You can give a policy a descriptive name to provide a convenient means for identifying its purpose.

Note: *For information regarding ScreenOS naming conventions—which apply to the names you create for policies—see "Naming Conventions and Character Types" on page xiv.*

VPN Tunneling

You can apply a single policy or multiple policies to any VPN tunnel that you have configured. In the WebUI, the VPN Tunnel option provides a drop-down list of all such tunnels. In the CLI, you can see all available tunnels with the **get vpn** command. (For more information, see "Site-to-Site VPNs" on page **5**-69 and "Dialup VPNs" on page **5**-199.)

When the VPN configurations at both ends of a VPN tunnel are using policy-based-NAT, then the administrators of both gateway devices each need to create an inbound and an outbound policy (four policies in total). When the VPN policies constitute a matching pair (that is, everything in the inbound and outbound policy configurations is the same except that the source and destination addresses are reversed), you can configure one policy and then select the **Modify matching bidirectional VPN policy** check box to create a second policy automatically for the opposite direction. For the configuration of a new policy, the matching VPN policy check box is cleared by default. For the modification of an existing policy that is a member of a matching pair, the check box is selected by default, and any changes made to one policy are propagated to the other.

Note: *This option is only available through the WebUI. It is not supported when there are multiple entries for any of the following policy components: source address, destination address, or service.*

L2TP Tunneling

You can apply a single policy or multiple policies to any Layer 2 Tunneling Protocol (L2TP) tunnel that you have configured. In the WebUI, the L2TP option provides a drop-down list of all such tunnels. In the CLI, you can see all available tunnels with the **get l2tp all** command. You can also combine a VPN tunnel and an L2TP tunnel—if both have the same endpoints—to create a tunnel combining the characteristics of each. This is called L2TP-over-IPSec.

Note: *A NetScreen device in Transparent mode does not support L2TP.*

Deep Inspection

Deep Inspection is a mechanism for filtering the traffic permitted at the Network and Transport Layers by examining not only these layers but the content and protocol characteristics at the Application Layer[5]. The goal of Deep Inspection is the detection and prevention any attacks or anomalous behavior that might be present in traffic that the NetScreen firewall permits. For more information, see “Deep Inspection” on page **4**-123.

To configure a policy for attack protection, you must make two choices: which attack group (or groups) to use and which attack action to take if an attack is detected. (For more information about Deep Inspection, see “Deep Inspection” on page **4**-123.)

Placement at the Top of the Policy List

By default, NetScreen positions a newly created policy at the bottom of a policy set list. If you need to reposition the policy, you can use either of the policy reordering methods explained in “Reordering Policies” on page 243. To avoid the extra step of repositioning a newly created policy to the top of a policy set list, you can select the **Position at Top** option in the WebUI, or use the keyword **top** in the **set policy** command (**set policy top …**) in the CLI.

5. In the Open Systems Interconnection (OSI) model, the Network Layer is Layer 3, the Transport Layer is Layer 4, and the Application Layer is Layer 7. The OSI model is a networking industry standard model of network protocol architecture. The OSI model consists of seven layers.

Source Address Translation

You can apply source address translation (NAT-src) at the policy level. With NAT-src, you can translate the source address on either incoming or outgoing network and VPN traffic. The new source address can come from either a dynamic IP (DIP) pool or the egress interface. NAT-src also supports source port address translation (PAT). To learn about all the NAT-src options that are available, see "Source Network Address Translation" on page 259.

Note: *You can also perform source address translation at the interface level, known as network address translation (NAT). For information about interface level NAT-src, or simply NAT, see "NAT Mode" on page 110.*

Destination Address Translation

You can apply destination address translation (NAT-dst) at the policy level. With NAT-dst, you can translate the destination address on either incoming or outgoing network and VPN traffic. NAT-dst can also support destination port mapping. To learn about all the NAT-dst options that are available, see "Destination Network Address Translation" on page 276.

User Authentication

Selecting this option requires the auth user at the source address to authenticate his/her identity by supplying a user name and password before traffic is allowed to traverse the firewall or enter the VPN tunnel. The NetScreen device can use the local database or an external RADIUS, SecurID, or LDAP auth server to perform the authentication check.

Note: *If a policy requiring authentication applies to a subnet of IP addresses, authentication is required for each IP address in that subnet.*

If a host supports multiple auth user accounts (as with a Unix host running Telnet), then after the NetScreen device authenticates the first user, all other users from that host can pass traffic through the NetScreen device without being authenticated, having inherited the privileges of the first user.

NetScreen provides two authentication schemes:

- Run-time authentication, in which the NetScreen device prompts an auth user to log on when it receives HTTP, FTP or Telnet traffic matching a policy that has authentication enabled
- WebAuth, in which a user must authenticate himself or herself before sending traffic through the NetScreen device

Run-Time Authentication

The run-time authentication process proceeds as follows:

1. When the auth user sends an HTTP, FTP or Telnet connection request to the destination address, the NetScreen device intercepts the packet and buffers it.
2. The NetScreen device sends the auth user a login prompt.
3. The auth user responds to this prompt with his/her user name and password.
4. The NetScreen device authenticates the auth user's login information.

   If the authentication is successful, a connection is established between the auth user and the destination address.

For the initial connection request, a policy must include one or all of the three following services: Telnet, HTTP, or FTP. Only a policy with one or all of these services is capable of initiating the authentication process. You can use any of the following services in a policy involving user authentication:

- Any (because "any" includes all three required services)
- Telnet, or FTP, or HTTP
- A service group that includes the service or services you want, plus one or more of the three services required to initiate the authentication process (Telnet, FTP, or HTTP). For example, you can create a custom service group named "Login" that supports FTP, Netmeeting, and H.323 services. Then, when you create the policy, specify "Login" as the service.

For any connection following a successful authentication, all services specified in the policy are valid.

Note: *A policy with authentication enabled does not support DNS (port 53) as the service.*

Pre-Policy Check Authentication (WebAuth)

The WebAuth authentication process proceeds as follows:

1. The auth user makes an HTTP connection to the IP address of the WebAuth server.
2. The NetScreen device sends the auth user a login prompt.
3. The auth user responds to this prompt with his/her user name and password.
4. The NetScreen device or an external auth server authenticates the auth user's login information.

   If the authentication attempt is successful, the NetScreen device permits the auth user to initiate traffic to destinations as specified in policies that enforce authentication via the WebAuth method.

Note: *For more information about these two user authentication methods, see "Auth Users and User Groups" on page 398.*

You can restrict or expand the range of auth users to which the policy applies by selecting a specific user group, local or external user, or group expression. (For information about group expressions, see "Group Expressions" on page 468.) If you do not reference an auth user or user group in a policy (in the WebUI, select the **Allow Any** option), the policy applies to all auth users in the specified auth server.

Note: *NetScreen links authentication privileges with the IP address of the host from which the auth user logs on. If the NetScreen device authenticates one user from a host behind a NAT device that uses a single IP address for all NAT assignments, then users at other hosts behind that NAT device automatically receive the same privileges.*

HA Session Backup

When two NetScreen devices are in an NSRP cluster for high availability (HA), you can specify which sessions to backup and which not to backup. For traffic whose sessions you do not want backed up, apply a policy with the HA session backup option disabled. In the WebUI, clear the **HA Session Backup** check box. In the CLI, use the **no-session-backup** argument in the **set policy** command. By default, NetScreen devices in an NSRP cluster back up sessions.

URL Filtering

NetScreen supports URL filtering using the Websense Enterprise Engine, which enables you to block or permit access to different sites based on their URLs, domain names, and IP addresses. When you enable URL filtering on a policy, the NetScreen device buffers all HTTP GET requests (in traffic to which the policy applies) and sends the URL to the Websense server. The Websense server compares the URL against its database. If it matches a restricted URL, the Websense server notifies the NetScreen device, which closes the TCP connection by sending a TCP RST to both the source and destination addresses. The NetScreen device also sends the source address a "blocked URL" message. If the URL does not match a restricted URL, the Websense server returns a "permit" message to the NetScreen device, which then forwards the buffered HTTP packet to its intended destination.

Logging

When you enable logging in a policy, the NetScreen device logs all connections to which that particular policy applies. You can view the logs through either the WebUI or CLI. In the WebUI, click **Reports > Policies >**  (for the policy whose log you want to see). In the CLI, use the **get log traffic policy** *id_num* command.

> ***Note:*** *For more information about viewing logs and graphs, see "Monitoring NetScreen Devices" on page* ***3****-65.*

Counting

When you enable counting in a policy, the NetScreen device counts the total number of bytes of traffic to which this policy applies and records the information in historical graphs. To view the historical graphs for a policy in the WebUI, click **Reports > Policies >**  (for the policy whose traffic count you want to see).

Traffic Alarm Threshold

You can set a threshold that triggers an alarm when the traffic permitted by the policy exceeds a specified number of bytes per second, bytes per minute, or both. Because the traffic alarm requires the NetScreen device to monitor the total number of bytes, you must also enable the counting feature.

> ***Note:*** *For more information about traffic alarms, see "Traffic Alarms" on page* ***3****-82.*

Schedules

By associating a schedule to a policy, you can determine when the policy is in effect. You can configure schedules on a recurring basis and as a one-time event. Schedules provide a powerful tool in controlling the flow of network traffic and in enforcing network security. For an example of the latter, if you were concerned about employees transmitting important data outside the company, you might set a policy that blocked outbound FTP-Put and MAIL traffic after normal business hours.

In the WebUI, define schedules in the **Objects > Schedules** section. In the CLI, use the **set schedule** command.

Note: *In the WebUI, scheduled policies appear with a gray background to indicate that the current time is not within the defined schedule. When a scheduled policy becomes active, it appears with a white background.*

Antivirus Scanning

Some NetScreen devices support antivirus (AV) scanning in SMTP and HTTP traffic using the Trend Micro InterScan VirusWall scanner (edition 3.6). When the NetScreen device receives traffic to which a policy with antivirus blocking applies, it sends it to the VirusWall antivirus scanner. When the scanner receives the entire content of an SMTP or HTTP packet, it examines the data for viruses by comparing it against its database of virus patterns. If it finds anything amiss, the VirusWall quarantines the infected data for further study and returns the SMTP or HTTP file—without the infected data—to the NetScreen device. The NetScreen device then forwards the file to the intended recipient.

Some NetScreen devices support an internal AV scanner that you can configure to filter POP3, SMTP, and HTTP traffic. If the embedded AV scanner detects a virus, it drops the packet and sends a message reporting the virus to the client initiating the traffic.

(For more information about antivirus scanning, see “Antivirus Scanning” on page **4**-76. For information specific to the VirusWall scanner, refer to the Trend Micro documentation.)

Traffic Shaping

You can set parameters for the control and shaping of traffic for each policy. The traffic shaping parameters include:

Guaranteed Bandwidth: Guaranteed throughput in kilobits per second (kbps). Traffic below this threshold passes with the highest priority without being subject to any traffic management or shaping mechanism.

Maximum Bandwidth: Secured bandwidth available to the type of connection in kilobits per second (kbps). Traffic beyond this threshold is throttled and dropped.

Note: *It is advised that you do not use rates less than 10 kbps. Rates below this threshold lead to dropped packets and excessive retries that defeat the purpose of traffic management.*

Traffic Priority: When traffic bandwidth falls between the guaranteed and maximum settings, the NetScreen device passes higher priority traffic first, and lower priority traffic only if there is no other higher priority traffic. There are eight priority levels.

DiffServ Codepoint Marking: Differentiated Services (DiffServ) is a system for tagging (or “marking”) traffic at a position within a hierarchy of priority. You can map the eight NetScreen priority levels to the DiffServ system. By default, the highest priority (priority 0) in the NetScreen system maps to the first three bits (0111) in the DiffServ field (see RFC 2474), or the IP precedence field in the ToS byte (see RFC 1349), in the IP packet header. The lowest priority (priority 7) in the NetScreen system maps to (0000) in the ToS DiffServ system.

Note: *For a more detailed discussion of traffic management and shaping, see* *“Traffic Shaping” on page 477**.*

To change the mapping between the NetScreen priority levels and the DiffServ system, use the following CLI command:

set traffic-shaping ip_precedence *number0 number1 number2 number3 number4 number5 number6 number7*

where *number0* is the mapping for priority 0 (the highest priority in the TOS DiffServ system), *number1* is the mapping for priority 1, and so forth.

POLICIES APPLIED

This section describes the management of policies: viewing, creating, modifying, ordering and reordering, and removing policies.

Viewing Policies

To view policies through the WebUI, click **Policies**. You can sort the displayed policies by source and destination zones by choosing zone names from the **From** and **To** drop-down lists and then clicking **Go**. In the CLI, use the **get policy** [**all** | **from** *zone* **to** *zone* | **global** | **id** *number*] command.

Policy Icons

When viewing a list of policies, the WebUI uses icons to provide you a graphical summary of policy components. The table below defines the different icons used in the policies page.

Icon	Function	Description
	Permit	The NetScreen device passes all traffic to which the policy applies.
	Deny	The NetScreen device blocks all traffic to which the policy applies.
	Policy-level NAT	The NetScreen device performs policy-based source or destination network address translation (NAT-src or NAT-dst) on all traffic to which the policy applies.
	Encapsulation and Decapsulation	The NetScreen device encapsulates all outbound VPN traffic and decapsulates all inbound VPN traffic to which the policy applies.
	Bidirectional VPN policies	A matching VPN policy exists for the opposite direction.

Icon	Function	Description
	Authentication	The user must authenticate himself/herself when initiating a connection.
	Antivirus	The NetScreen device sends all traffic to which the policy applies to a Trend Micro antivirus (AV) scanner.
	Deep Inspection	The NetScreen device performs Deep Inspection (DI) on all traffic to which the policy applies.
	Deep Inspection and Antivirus	The NetScreen device performs Deep Inspection and antivirus protection on all traffic to which the policy applies.
	Logging	All traffic is logged and made available for syslog and e-mail, if enabled.
	Counting	The NetScreen device counts (in bytes) the amount of traffic to which the policy applies.
	Traffic Alarm	Indicates that you have set traffic alarm thresholds.

Creating Policies

To allow traffic to flow between two zones, you create policies to permit, deny, or tunnel traffic between those zones. You can also create policies to control traffic within the same zone if the NetScreen device is the only network device that can route the intrazone traffic between the source and destination addresses referenced in the policy. You can also create global policies, which make use of source and destination addresses in the Global zone address book.

To allow bidirectional traffic between two zones—for example, between the Trust and Untrust zones—you need to create a policy that goes from Trust to Untrust, and then create a second policy from Untrust to Trust. Depending on your needs, the policies can use the same or different IP addresses, only the source and destination addresses are reversed.

Policy Location

You can define policies between any zones that are located within the same system—root or virtual. To define a policy between the root system and a vsys, one of the zones must be a shared zone. (For information about shared zones in relation to virtual systems, see Volume 7, “Virtual Systems”.)

Example: Interzone Policies for E-Mail Service

In this example, you create three policies to control the flow of e-mail traffic.

The first policy allows internal users in the Trust zone to send and retrieve e-mail from a local mail server in the DMZ zone. This policy permits the services MAIL (that is, SMTP) and POP3 originating from the internal users to traverse the NetScreen firewall to reach the local mail server.

The second and third policies permit the service MAIL to traverse the firewall between the local mail server in the DMZ zone and a remote mail server in the Untrust zone.

However, before creating policies to control traffic between different security zones, you must first design the environment in which to apply those policies. First, you first bind interfaces to zones and assign the interfaces IP addresses:

- Bind ethernet1 to the Trust zone and assign it IP address 10.1.1.1/24.
- Bind ethernet2 to the DMZ zone and assign it IP address 1.2.2.1/24.
- Bind ethernet3 to the Untrust zone and assign it IP address 1.1.1.1/24.

  All security zones are in the trust-vr routing domain.

Second, you create addresses for use in the policies:

- Define an address in the Trust zone named “corp_net” and assign it IP address 10.1.1.0/24.
- Define an address in the DMZ zone named “mail_svr” and assign it IP address 1.2.2.5/32.
- Define an address in the Untrust zone named “r-mail_svr” and assign it IP address 2.2.2.5/32.

Third, you create a service group named “MAIL-POP3” containing the two predefined services MAIL and POP3.

Fourth, you configure a default route in the trust-vr routing domain pointing to the external router at 1.1.1.250 through ethernet3.

After completing steps 1 – 4, you can then create the policies necessary to permit the transmission, retrieval, and delivery of e-mail in and out of your protected network.

WebUI

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Enter the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet2): Enter the following, and then click **OK**:

Zone Name: DMZ

Static IP: (select this option when present)

IP Address/Netmask: 1.2.2.1/24

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

2. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: corp_net

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.0/24

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: mail_svr

IP Address/Domain Name:

IP/Netmask: (select), 1.2.2.5/32

Zone: DMZ

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: r-mail_svr

IP Address/Domain Name:

IP/Netmask: (select), 2.2.2.5/32

Zone: Untrust

3. Service Group

Objects > Services > Group: Enter the following group name, move the following services, and then click **OK**:

Group Name: MAIL-POP3

Select **MAIL** and use the **<<** button to move the service from the Available Members column to the Group Members column.

Select **POP3** and use the **<<** button to move the service from the Available Members column to the Group Members column.

4. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 1.1.1.250

5. Policies

Policies > (From: Trust, To: Untrust) > New : Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), corp_net

Destination Address:

Address Book Entry: (select), mail_svr

Service: Mail-POP3

Action: Permit

Policies > (From: DMZ, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), mail_svr

Destination Address:

Address Book Entry: (select), r-mail_svr

Service: MAIL

Action: Permit

Policies > (From: Untrust, To: DMZ) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), r-mail_svr

Destination Address:

Address Book Entry: (select), mail_svr

Service: MAIL

Action: Permit

CLI

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet2 zone dmz
set interface ethernet2 ip 1.2.2.1/24
set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
```

2. Addresses

```
set address trust corp_net 10.1.1.0/24
set address dmz mail_svr 1.2.2.5/32
set address untrust r-mail_svr 2.2.2.5/32
```

3. Service Group

```
set group service MAIL-POP3
set group service MAIL-POP3 add mail
set group service MAIL-POP3 add pop3
```

4. Route

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
```

5. Policies

```
set policy from trust to dmz corp_net mail_svr MAIL-POP3 permit
set policy from dmz to untrust mail_svr r-mail_svr MAIL permit
set policy from untrust to dmz r-mail_svr mail_svr MAIL permit
save
```

Example: Interzone Policy Set

A small software firm, ABC Design, has divided its internal network into two subnets, and both are in the Trust zone. These two subnets are:

- Engineering (with the defined address “Eng”)
- The rest of the company (with the defined address “Office”).

It also has a DMZ zone for its Web and mail servers.

The following example presents a typical set of policies for the following users:

- “Eng” can use all the services for outbound traffic except FTP-Put, IMAP, MAIL, and POP3.
- “Office” can use e-mail and access the Internet, provided they authenticate themselves via WebAuth. (For information about WebAuth user authentication, see “Auth Users and User Groups” on page 398.)
- Everyone in the Trust zone can access the Web and mail servers in the DMZ zone.
- A remote mail server in the Untrust zone can access the local mail server in the DMZ zone.
- There is also a group of system administrators (with the user-defined address “sys-admins”) who have complete user and administrative access to the servers in the DMZ zone.

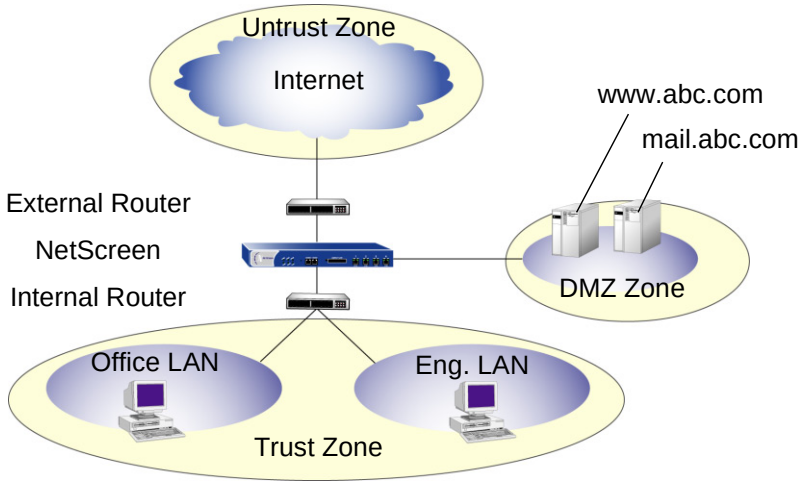

This example focuses only on policies and assumes that you have already configured the interfaces, addresses, service groups, and routes that must be in place. For more information on configuring these, see “Interfaces” on page 65, “Addresses” on page 126, “Service Groups” on page 167, and “Routing Tables and Static Routing” on page 29.

From Zone - Src Addr	To Zone - Dest Addr	Service	Action
Trust - Any	Untrust - Any	Com (service group: FTP-Put, IMAP, MAIL, POP3)	Deny
Trust - Eng	Untrust - Any	Any	Permit
Trust - Office	Untrust - Any	Internet (service group: FTP-Get, HTTP, HTTPS)	Permit (+ WebAuth)

From Zone - Src Addr	To Zone - Dest Addr	Service	Action
Untrust - Any	DMZ - mail.abc.com	MAIL	Permit
Untrust - Any	DMZ - www.abc.com	Web (service group: HTTP, HTTPS)	Permit

From Zone - Src Addr	To Zone - Dest Addr	Service	Action
Trust - Any	DMZ - mail.abc.com	e-mail (service group: IMAP, MAIL, POP3)	Permit
Trust - Any	DMZ - www.abc.com	Internet (service group: FTP-Get, HTTP, HTTPS)	Permit
Trust - sys-admins	DMZ - Any	Any	Permit

From Zone - Src Addr	To Zone - Dest Addr	Service	Action
DMZ - mail.abc.com	Untrust - Any	MAIL	Permit

Note: *The default policy is to deny all.*

WebUI

1. From Trust, To Untrust

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Eng

Destination Address:

Address Book Entry: (select), Any

Service: ANY

Action: Permit

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Office

Destination Address:

Address Book Entry: (select), Any

Service: Internet[6]

Action: Permit

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

Authentication: (select)

WebAuth: (select)

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Any

Service: Com[7]

Action: Deny

Position at Top: (select)

Note: *For traffic from the Untrust zone to the Trust zone, the default deny policy denies everything.*

6. "Internet" is a service group with the following members: FTP-Get, HTTP, and HTTPS.
7. "Com" is a service group with the following members: FTP-Put, MAIL, IMAP, and POP3.

2. From Untrust, To DMZ

Policies > (From: Untrust, To: DMZ) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), mail.abc.com

Service: MAIL

Action: Permit

Policies > (From: Untrust, To: DMZ) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), www.abc.com

Service: Web[8]

Action: Permit

8. "Web" is a service group with the following members: HTTP and HTTPS.

3. From Trust, To DMZ

Policies > (From: Trust, To: DMZ) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), mail.abc.com

Service: e-mail[9]

Action: Permit

Policies > (From: Trust, To: DMZ) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), www.abc.com

Service: Internet

Action: Permit

9. “e-mail” is a service group with the following members: MAIL, IMAP, and POP3.

Policies > (From: Trust, To: DMZ) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), sys-admins

Destination Address:

Address Book Entry: (select), Any

Service: ANY

Action: Permit

4. From DMZ, To Untrust

Policies > (From: DMZ, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), mail.abc.com

Destination Address:

Address Book Entry: (select), Any

Service: MAIL

Action: Permit

CLI

1. From Trust, To Untrust

```
set policy from trust to untrust eng any any permit
set policy from trust to untrust office any Internet[10] permit webauth
set policy top from trust to untrust any any Com[11] deny
```

2. From Untrust, To DMZ

```
set policy from untrust to dmz any mail.abc.com mail permit
set policy from untrust to dmz any www.abc.com Web[12] permit
```

3. From Trust, To DMZ

```
set policy from trust to dmz any mail.abc.com e-mail[13] permit
set policy from trust to dmz any www.abc.com Internet[10] permit
set policy from trust to dmz sys-admins any any permit
```

4. From DMZ, To Untrust

```
set policy from dmz to untrust mail.abc.com any mail permit
save
```

10. “Internet” is a service group with the following members: FTP-Get, HTTP, and HTTPS.
11. “Com” is a service group with the following members: FTP-Put, MAIL, IMAP, and POP3.
12. “Web” is a service group with the following members: HTTP and HTTPS.
13. “e-mail” is a service group with the following members: MAIL, IMAP, and POP3.

Example: Intrazone Policies

In this example, you create an intrazone policy to permit a group of accountants access to a confidential server on the corporate LAN in the Trust zone. You first bind ethernet1 to the Trust zone and give it IP address 10.1.1.1/24. You then bind ethernet2 to the Trust zone and assign it IP address 10.1.5.1/24. You enable intrazone blocking in the Trust zone. Next, you define two addresses—one for a server on which the company stores its financial records (10.1.1.100/32) and another for the subnet on which hosts for the accounting department are located (10.1.5.0/24). You then create an intrazone policy to permit access to the server from those hosts.

WebUI

1. Trust Zone – Interfaces and Blocking

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.5.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

Network > Zones > Edit (for Trust): Enter the following, and then click **OK**:

Block Intra-Zone Traffic: (select)

2. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Hamilton

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.100/32

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: accounting

IP Address/Domain Name:

IP/Netmask: (select), 10.1.5.0/24

Zone: Trust

3. Policy

Policies > (From: Trust, To: Trust) > New : Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), accounting

Destination Address:

Address Book Entry: (select), Hamilton

Service: ANY

Action: Permit

CLI

1. Trust Zone – Interfaces and Blocking

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet2 zone trust
set interface ethernet2 ip 10.1.5.1/24
set interface ethernet2 nat

set zone trust block
```

2. Addresses

```
set address trust Hamilton 10.1.1.100/32
set address trust accounting 10.1.5.0/24
```

3. Policy

```
set policy from trust to trust accounting Hamilton any permit
save
```

Example: Global Policy

In this example, you create a global policy so that every host in every zone can access the company Web site, which is www.netscreen.com[14]. Using a global policy is a convenient shortcut when there are many security zones. In this example, one global policy accomplishes what *n* interzone policies would have accomplished (where *n* = number of zones).

WebUI

1. Global Address

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: server1

IP Address/Domain Name:

Domain Name: (select), www.netscreen.com

Zone: Global

2. Policy

Policies > (From: Global, To: Global) > New : Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), server1

Service: HTTP

Action: Permit

14. To use a domain name instead of an IP address, be sure to have DNS service configured on the NetScreen device.

CLI

1. Global Address

```
set address global server1 www.netscreen.com
```

2. Policy

```
set policy global any server1 http permit
save
```

Entering a Policy Context

When configuring a policy through the CLI, after you first create a policy, you can then enter the context of the policy to make additions and modifications. For example, perhaps you first create the following policy:

```
set policy id 1 from trust to untrust host1 server1 HTTP permit attack
    HIGH:HTTP:SIGS action close
```

If you want to make some changes to the policy, such as adding another source or destination address, another service, or another attack group, you can enter the context for policy 1 and then enter the pertinent commands:

```
set policy id 1
ns(policy:1)-> set src-address host2
ns(policy:1)-> set dst-address server2
ns(policy:1)-> set service FTP
ns(policy:1)-> set attack CRITICAL:HTTP:SIGS
```

You can also remove multiple items for a single policy component as long as you do not remove them all. For example, you can remove server2 from the above configuration, but not server2 and server1 because then no destination address would remain:

You can remove either server2, ✓

```
ns(policy:1)-> unset dst-address server2
```

or you can remove server1, ✓

```
ns(policy:1)-> unset dst-address server1
```

but you cannot remove them both. ✗

```
ns(policy:1)-> unset dst-address server2
ns(policy:1)-> unset dst-address server1
```

Multiple Items per Policy Component

ScreenOS allows you to add multiple items to the following components of a policy:

- Source address
- Destination address
- Service
- Attack group

In pre-ScreenOS 5.0.0 releases, the only way to have multiple source and destination addresses or services is to first create an address or service group with multiple members and then reference that group in a policy. You can still use address and service groups in policies in ScreenOS 5.0.0. In addition, you can simply add extra items directly to a policy component.

Note: *If the first address or service referenced in a policy is "Any", you cannot logically add anything else to it. NetScreen prevents this kind of misconfiguration and displays an error message should it occur.*

To add multiple items to a policy component, do either of the following:

WebUI

To add more addresses and services, click the **Multiple** button next to the component to which you want to add more items. To add more attack groups, click the **Attack Protection** button. Select an item in the "Available Members" column, and then use the **<<** key to move it to the "Active Members" column. You can repeat this action with other items. When finished, click **OK** to return to the policy configuration page.

CLI

Enter the policy context with the following command:

```
set policy id number
```

Then use one of the following commands as appropriate:

```
ns(policy:number)-> set src-address string
ns(policy:number)-> set dst-address string
ns(policy:number)-> set service string
ns(policy:number)-> set attack string
```

Address Negation

You can configure a policy so that it applies to all addresses except the one specified as either the source or destination. For example, you might want to create a policy that permits Internet access to everyone except the "P-T_contractors" address group. To accomplish this, you can use the address negation option.

In the WebUI, this option is available on the pop-up that appears when you click the **Multiple** button next to either Source Address or Destination Address on the policy configuration page.

In the CLI, you insert an exclamation point (!) immediately before source or destination address.

Note: *Address negation occurs at the policy component level, applying to all items in the negated component.*

Example: Destination Address Negation

In this example, you create an intrazone policy that allows all addresses in the Trust zone access to all FTP servers except to an FTP server named "vulcan", which engineering uses to post functional specifications for one another.

However, before creating the policy, you must first design the environment in which to apply it. First, you enable intrazone blocking for the Trust zone. Intrazone blocking requires a policy lookup before the NetScreen device passes traffic between two interfaces bound to the same zone.

Second, you bind two interfaces to the Trust zone and assign them IP addresses:

- You bInd ethernet1 to the Trust zone and assign it IP address 10.1.1.1/24.
- You bind ethernet4 to the Trust zone and assign it IP address 10.1.2.1/24.

Third, you create an address (10.1.2.5/32) for the FTP server named "vulcan" in the Trust zone.

After completing these two steps , you can then create the intrazone policies.

Note: *You do not have to create a policy for the engineering department to reach their FTP server because the engineers are also in the 10.1.2.0/24 subnet and do not have to cross the NetScreen firewall to reach their own server.*

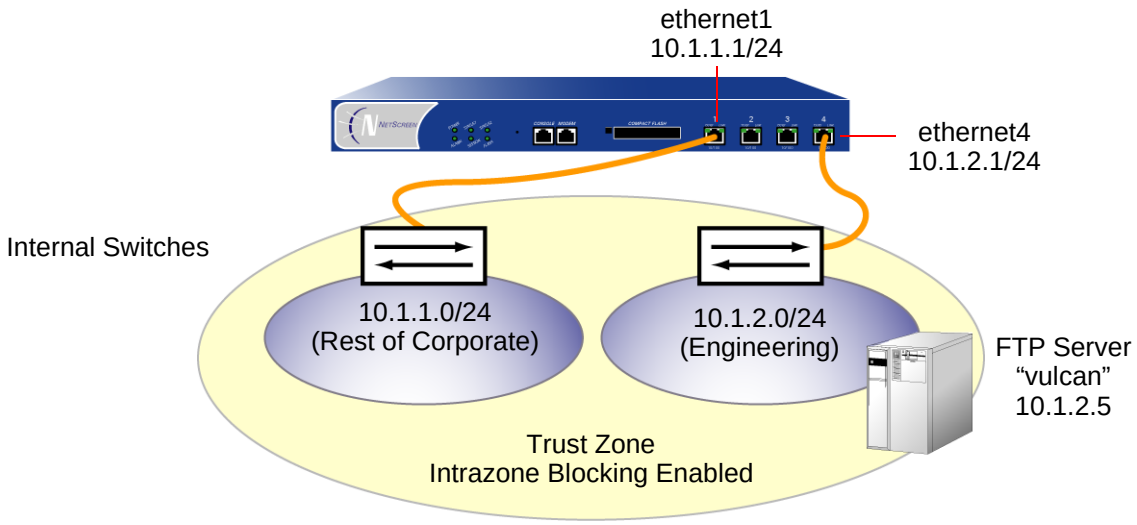

WebUI

1. Intrazone Blocking

Network > Zones > Edit (for Trust): Enter the following, and then click **OK**:

Virtual Router Name: trust-vr

Block Intra-Zone Traffic: (select)

2. Trust Zone Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet4): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.2.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

3. Address

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: vulcan

IP Address/Domain Name:

IP/Netmask: (select), 10.1.2.5/32

Zone: Trust

4. Policy

Policies > (From: Trust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), vulcan

> Click **Multiple**, select the **Negate Following** check box, and then click **OK** to return to the basic policy configuration page.

Service: FTP

Action: Permit

CLI

1. Intrazone Blocking

```
set zone trust block
```

2. Trust Zone Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet4 zone trust
set interface ethernet4 ip 10.1.2.1/24
set interface ethernet1 nat
```

3. Address

```
set address trust vulcan 10.1.2.5/32
```

4. Policy

```
set policy from trust to trust any !vulcan ftp permit
save
```

Modifying and Disabling Policies

After you create a policy, you can always return to it to make modifications. In the WebUI, click the **Edit** link in the Configure column for the policy that you want to change. In the Policy configuration page that appears for that policy, make your changes, and then click **OK**. In the CLI, use the **set policy** command.

ScreenOS also provides a means for enabling and disabling policies. By default, a policy is enabled. To disable it, do the following:

WebUI

Policies: Clear the **Enable** check box in the Configure column for the policy that you want to disable.

The row of text for a disabled policy appears as grey.

CLI

```
set policy id id_num disable
save
```

Note: *To enable the policy again, select* ***Enable*** *in the Configure column for the policy that you want to enable (WebUI), or type* ***unset policy id*** *id_num* ***disable*** *(CLI).*

Policy Verification

ScreenOS offers a tool for verifying that the order of policies in the policy list is valid. It is possible for one policy to eclipse, or "shadow", another policy. Consider the following example:

```
set policy id 1 from trust to untrust any any HTTP permit
set policy id 2 from trust to untrust any dst-A HTTP deny
```

Because the NetScreen device performs a policy lookup starting from the top of the list, when it finds a match for traffic received, it does not look any lower in the policy list. In the above example, the NetScreen device never reaches policy 2 because the destination address "any" in policy 1 includes the more specific "dst-A" address in policy 2. When an HTTP packet arrives at the NetScreen device from an address in the Trust zone bound for dst-A in the Untrust zone, the NetScreen device always first finds a match with policy 1.

To correct the above example, you can simply reverse the order of the policies, putting the more specific one first:

```
set policy id 2 from trust to untrust any dst-A HTTP deny
set policy id 1 from trust to untrust any any HTTP permit
```

Of course, this example is purposefully simple to illustrate the basic concept. In cases where there are dozens or hundreds of policies, the eclipsing of one policy by another might not be so easy to spot. To check if there is any policy shadowing[15] in your policy list, you can use the following CLI command:

```
exec policy verify
```

This command reports the shadowing and shadowed policies. It is then the admin's responsibility to correct the situation.

The policy verification tool cannot detect the case where a combination of policies shadows another policy. In the following example, no single policy shadows policy 3; however, policies 1 and 2 together do shadow it:

```
set group address trust grp1 add host1
set group address trust grp1 add host2
set policy id 1 from trust to untrust host1 server1 HTTP permit
set policy id 2 from trust to untrust host2 server1 HTTP permit
set policy id 3 from trust to untrust grp1 server1 HTTP deny
```

15. The concept of policy "shadowing" refers to the situation where a policy higher in the policy list always takes effect before a subsequent policy. Because the policy lookup always uses the first policy it finds that matches the five-part tuple of source and destination zone, source and destination address, and service type, if another policy applies to the same tuple (or a subset of the tuple), the policy lookup uses the first policy in the list and never reaches the second one.

Reordering Policies

The NetScreen device checks all attempts to traverse the firewall against policies, beginning with the first one listed in the policy set for the appropriate list (see "Policy Set Lists" on page 202) and moving through the list. Because the NetScreen device applies the action specified in the policy to the first matching policy in the list, you must arrange them from the most specific to the most general. (Whereas a specific policy does not preclude the application of a more general policy located down the list, a general policy appearing before a specific one does.)

By default, a newly created policy appears at the bottom of a policy set list. There is an option that allows you to position a policy at the top of the list instead. In the Policy configuration page in the WebUI, select the **Position at Top** check box. In the CLI, add the key word **top** to the **set policy** command: **set policy top** …

To move a policy to a different position in the list, do either of the following:

WebUI

There are two ways to reorder policies in the WebUI: by clicking the circular arrows or by clicking the single arrow in the Configure column for the policy you want to move.

If you click the circular arrows:

A User Prompt dialog box appears.

To move the policy to the very end of the list, enter <-1>. To move it up in the list, enter the ID number of the policy above which you want to move the policy in question.

Click **OK** to execute the move.

If you click the single arrow:

A Policy Move page appears displaying the policy you want to move and a table displaying the other policies.

In the table displaying the other policies, the first column, Move Location, contains arrows pointing to various locations where you can move the policy. Click the arrow that points to the location in the list where you want to move the policy.

The Policy List page reappears with the policy you moved in its new position.

CLI

```
set policy move id_num { before | after } number
save
```

Removing a Policy

In addition to modifying and repositioning a policy, you can also delete it. In the WebUI, click **Remove** in the Configure column for the policy that you want to remove. When the system message prompts for confirmation to proceed with the removal, click **Yes**. In the CLI, use the **unset policy** *id_num* command.

8

Address Translation

NetScreen provides many methods for performing source and destination IP address and source and destination port address translation. This chapter describes the various address translation methods available and is organized into the following sections:

- "Introduction to Address Translation" on page 246
  - "Policy-Based Translation Options" on page 253
  - "Directional Nature of NAT-Src and NAT-Dst" on page 257
- "Source Network Address Translation" on page 259
  - "NAT-Src from a DIP Pool with PAT Enabled" on page 260
  - "NAT-Src from a DIP Pool with PAT Disabled" on page 264
  - "NAT-Src from a DIP Pool with Address Shifting" on page 267
  - "NAT-Src from the Egress Interface IP Address" on page 273
- "Destination Network Address Translation" on page 276
  - "Packet Flow for Destination Translation" on page 278
  - "Routing for Destination Translation" on page 282
  - "NAT-Dst: One-to-One Mapping" on page 286
  - "NAT-Dst: Many-to-One Mapping" on page 295
  - "NAT-Dst: Many-to-Many Mapping" on page 300
  - "NAT-Dst with Port Mapping" on page 305
  - "NAT-Src and NAT-Dst in the Same Policy" on page 310
- "Mapped IP Addresses" on page 331
  - "MIP and the Global Zone" on page 332
  - "MIP-Same-as-Untrust" on page 342
  - "MIP and the Loopback Interface" on page 346
- "Virtual IP Addresses" on page 356
  - "VIP and the Global Zone" on page 359

INTRODUCTION TO ADDRESS TRANSLATION

NetScreen provides several mechanisms for applying network address translation (NAT). The concept of NAT comprises the translation of the IP address in an IP packet header and, optionally, the translation of the port number in the TCP segment or UDP datagram header. The translation can involve the source address (and optionally the source port number), the destination address (and optionally the destination port number), or a combination of translated elements.

Source Network Address Translation

When performing source network address translation (NAT-src), the NetScreen device translates the original source IP address to a different address. The translated address can come from a dynamic IP (DIP) pool or from the egress interface of the NetScreen device. If the NetScreen device draws the translated address from a DIP pool, it can do so either arbitrarily or deterministically; that is, it can draw any address from the DIP pool at random, or it can consistently draw a specific address in relation to the original source IP address[1]. If the translated address comes from the egress interface, the NetScreen device translates the source IP address in all packets to the IP address of that interface. You can configure the NetScreen device to apply NAT-src at either the interface level or at the policy level. If you configure a policy to apply NAT-src and the ingress interface is in NAT mode, the policy-based NAT-src settings override the interface-based NAT[2]. (This chapter focusses on policy-based NAT-src. For details on interface-based NAT-src—or just "NAT"—see "NAT Mode" on page 110. For more information about DIP pools, see "DIP Pools" on page 171.)

Source IP Address Translation

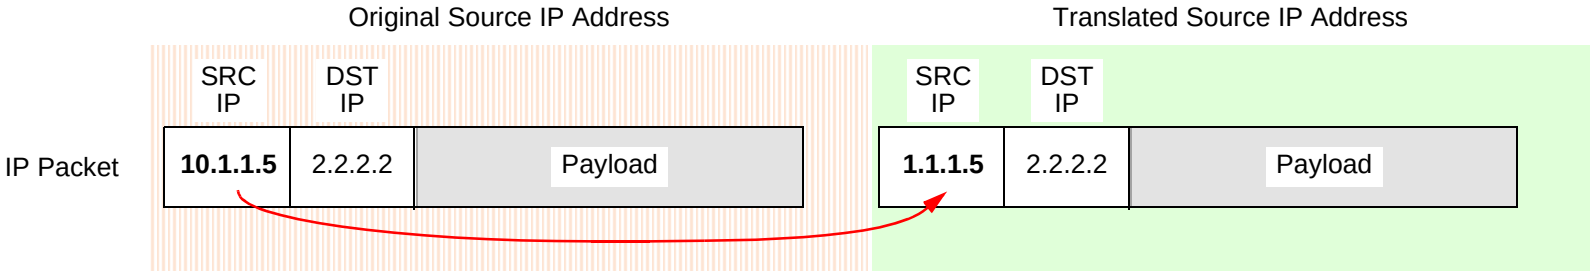

1. Deterministic address translation uses a technique called address shifting, which is explained later in this chapter. For information about address shifting that applies to NAT-src, see "NAT-Src from a DIP Pool with Address Shifting" on page 267. For information about address shifting that applies to NAT-dst, see "NAT-Src and NAT-Dst in the Same Policy" on page 310.
2. You can use policy-based NAT-src when the ingress interface is in Route or NAT mode. If it is in NAT mode, the policy-level NAT-src parameters supersede the interface-level NAT parameters.

With policy-based NAT-src, you can optionally choose to have the NetScreen device perform port address translation (PAT) on the original source port number. When PAT is enabled, the NetScreen device can translate up to ~64,500 different IP addresses to a single IP address with up to ~64,500 different port numbers[3]. The NetScreen device uses the unique, translated port number to maintain session state information for traffic to and from the same, single IP address. For interface-based NAT-src—or just "NAT"—port address translation is enabled automatically. Because the NetScreen device translates all original IP addresses to the same translated IP address (that of the egress interface), the NetScreen device uses the translated port number to identify each session to which a packet belongs. Similarly, if a DIP pool consists of only one IP address and you want the NetScreen device to apply NAT-src to multiple hosts using that address, then PAT is required for the same reason.

Source IP Address Translation and
Source Port Address Translation

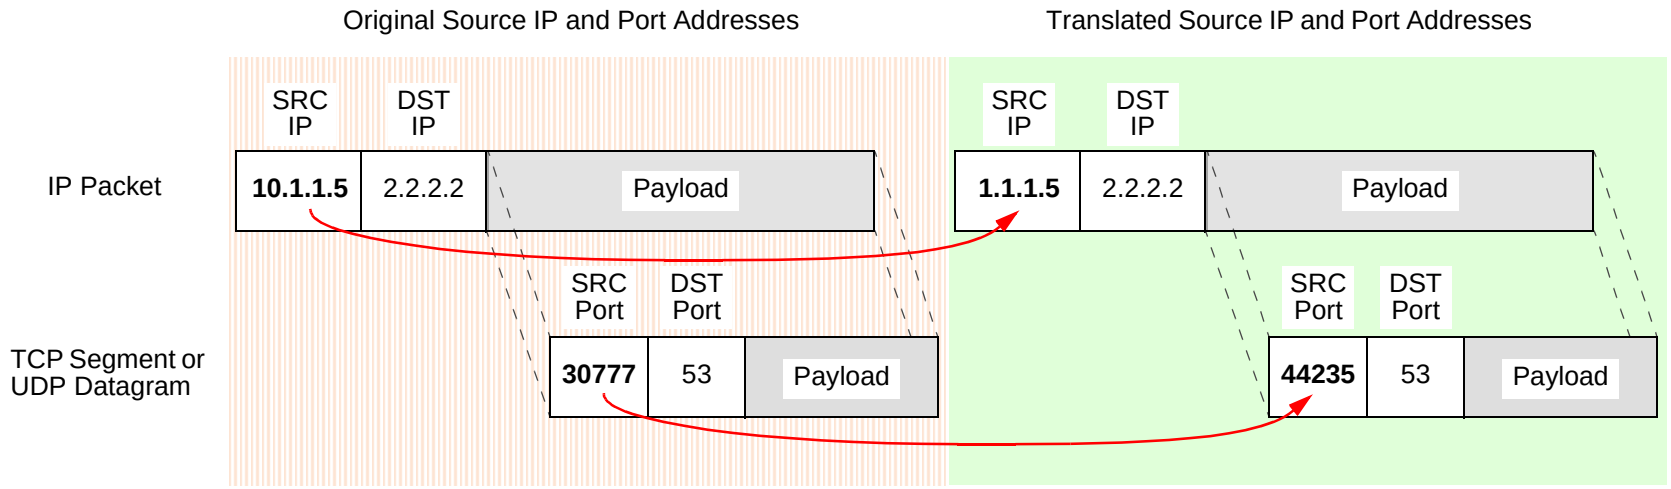

3. With PAT enabled, the NetScreen device maintains a pool of free port numbers to assign along with addresses from the DIP pool. The figure of ~64,500 is derived by subtracting 1023, the numbers reserved for the well-known ports, from the maximum number of ports, which is 65,535. Thus, when the NetScreen device performs NAT-src with a DIP pool containing a single IP address and PAT is enabled, the NetScreen device can translate the original IP addresses of up to ~64,500 hosts to a single IP address and translate each original port number to a unique port number.

For custom applications that require a specific source port number to operate properly, performing PAT causes such applications to fail. To provide for such cases, you can disable PAT.

Note: *For more information about NAT-src, see* *"Source Network Address Translation" on page 259**.*

Destination Network Address Translation

NetScreen offers the following three mechanisms for performing destination network address translation (NAT-dst):

- Policy-based NAT-dst
- Mapped IP (MIP)
- Virtual IP (VIP)

All three options translate the original destination IP address in an IP packet header to a different address. WIth policy-based NAT-dst and VIPs, you can optionally enable port mapping[4].

Note: *NetScreen does not support the use of policy-based NAT-dst in combination with MIPs and VIPs. If you have configured a MIP or VIP, the NetScreen device applies the MIP or VIP to any traffic to which a policy-based NAT-dst configuration also applies. In other words, MIPs and VIPs disable policy-based NAT-dst if the NetScreen device is accidentally configured to apply both to the same traffic.*

4. For information about port mapping, see the "Policy-Based NAT-Dst" on the next page and also "Destination Network Address Translation" on page 276.

Policy-based NAT-Dst: You can configure a policy to translate one destination IP address to another address, one IP address range to a single IP address, or one IP address range to another IP address range. When a single destination IP address translates to another IP address or an IP address range translates to a single IP address, NetScreen can support NAT-dst with or without port mapping. Port mapping is the deterministic translation of one original destination port number to another specific number, unlike PAT which translates any original source port number randomly assigned by the initiating host to another number randomly assigned by the NetScreen device.

Destination IP Address Translation without Destination Port Mapping

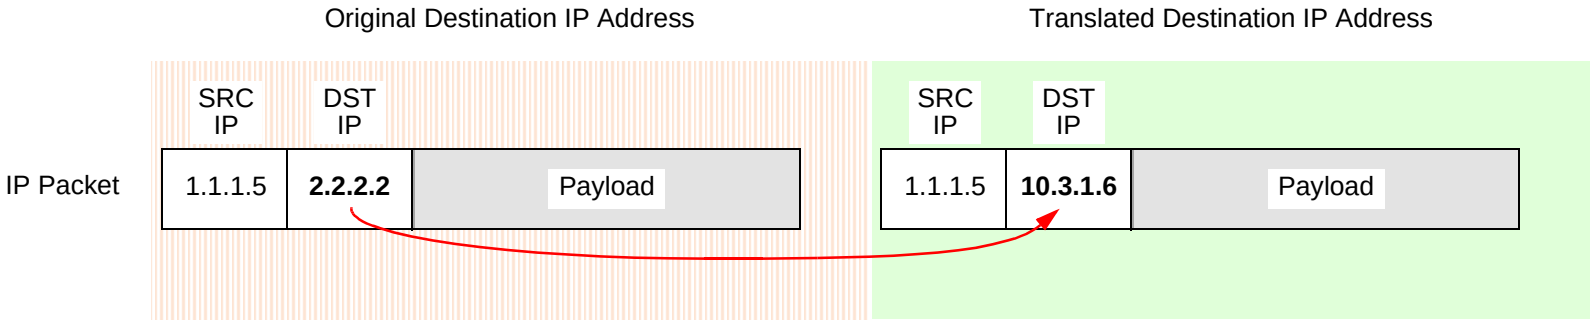

Destination IP Address Translation with Destination Port Mapping

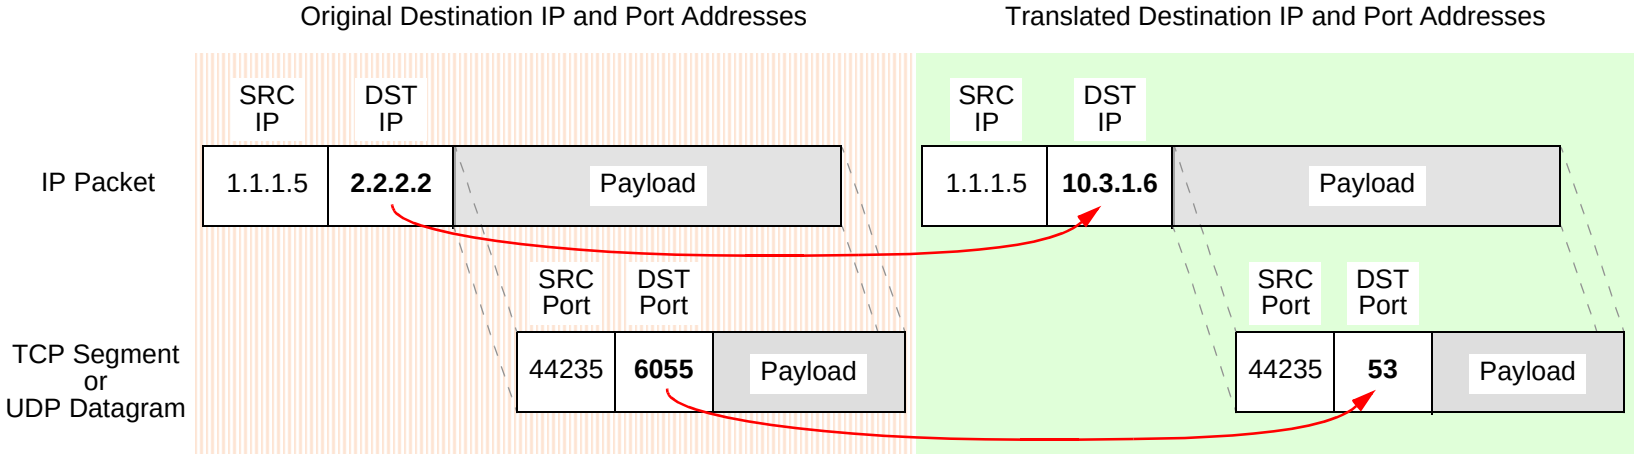

When you configure a policy to perform NAT-dst to translate an address range to a single address, the NetScreen device translates any destination IP address from within the user-defined range of original destination addresses to a single address. You can also enable port mapping.

Destination IP Address Translation from an IP Address Range to a Single IP Address

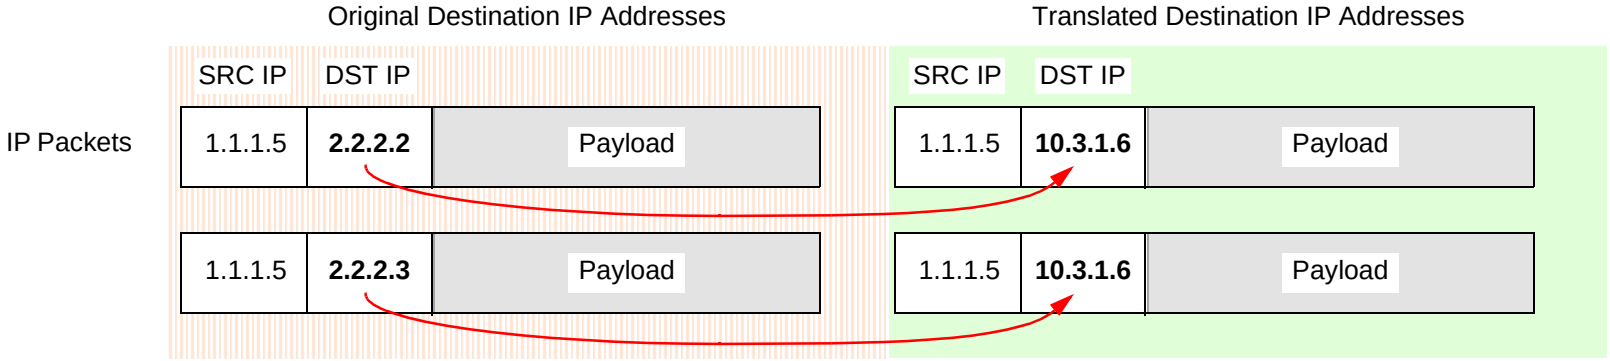

Destination IP Address Translation from an IP Address Range to a Single IP Address with Destination Port Mapping

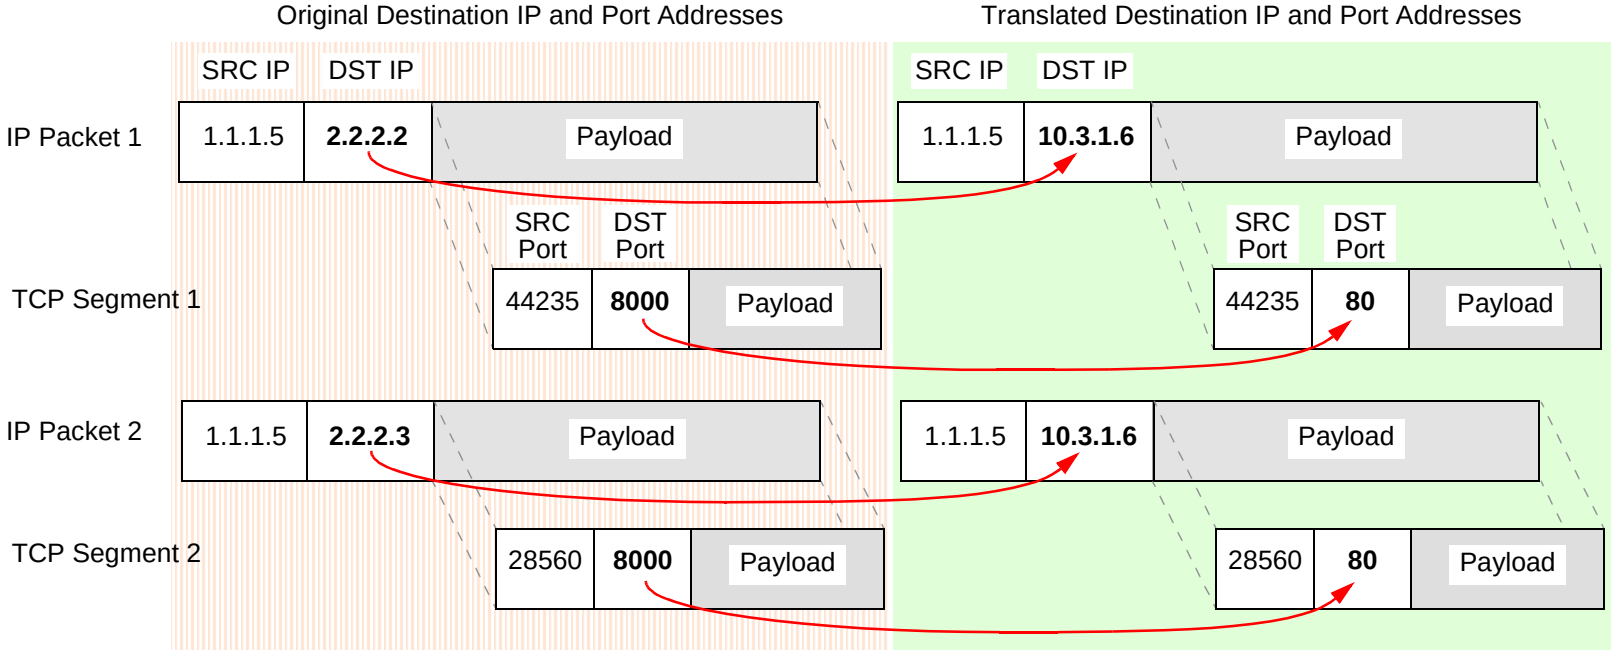

When you configure a policy to perform NAT-dst for an address range, the NetScreen device uses address shifting to translate a destination IP address from within a range of original destination addresses to a known address in another range of addresses.

Destination IP Address Translation
with Address Shifting

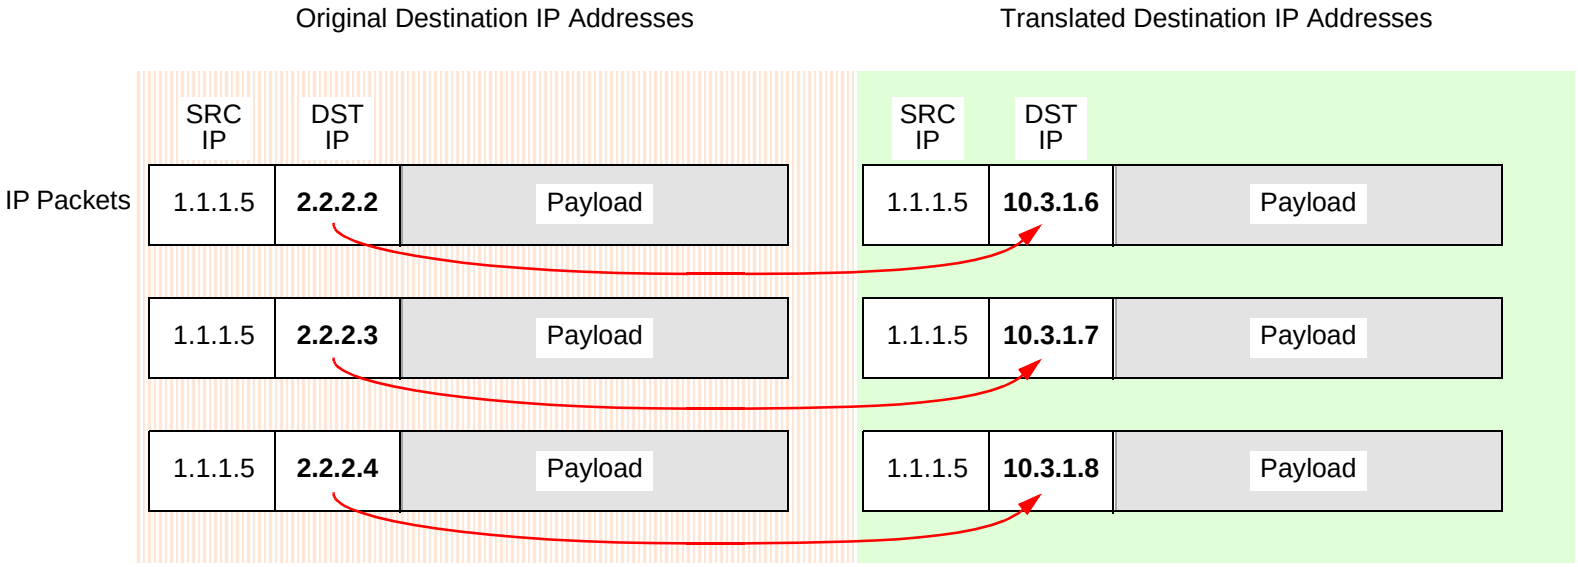

When performing NAT-dst for a range of IP addresses, the NetScreen device maintains a mapping of each IP address in one address range to a corresponding IP address in another address range.

Note: *You can combine NAT-src and NAT-dst within the same policy. Each translation mechanism operates independently and unidirectionally. That is, if you enable NAT-dst on traffic from zone1 to zone2, the NetScreen device does not perform NAT-src on traffic originating from zone2 and destined to zone1 unless you specifically configure it to do so. For more information, see* *"Directional Nature of NAT-Src and NAT-Dst" on page 257**. For more information about NAT-dst, see* *"Destination Network Address Translation" on page 276**.*

MIPs: A MIP is a mapping of one IP address to another IP address. You define one address in the same subnet as an interface IP address. The other address belongs to the host to which you want to direct traffic. Address translation for a MIP behaves bidirectionally, so that the NetScreen device translates the destination IP address in all traffic coming to a MIP to the host IP address and source IP address in all traffic originating from the host IP address to the MIP address. MIPs do not support port mapping. For more information about MIPs, see "Mapped IP Addresses" on page 331.

VIPs: A VIP is a mapping of one IP address to another IP address based on the destination port number. A single IP address defined in the same subnet as an interface can host mappings of several services—identified by various destination port numbers—to as many hosts[5]. VIPs also support port mapping. Like MIPs, address translation for a VIP behaves bidirectionally. The NetScreen device translates the destination IP address in all traffic coming to a VIP to a host IP address and source IP address in all traffic originating from the host IP address to the VIP address. For more information about VIPs, see "Virtual IP Addresses" on page 356.

Whereas the address translation mechanisms for MIPs and VIPs are bidirectional, the capabilities provided by policy-based NAT-src and NAT-dst separate address translation for inbound and outbound traffic, providing better control and security. For example, if you use a MIP to a Web server, whenever that server initiates outbound traffic to get an update or patch, its activity is exposed, which might provide information for a vigilant attacker to exploit. The policy-based address translation methods allow you to define a different address mapping when the Web server receives traffic (using NAT-dst) than when it initiates traffic (using NAT-src). By thus keeping its activities hidden, you can better protect the server from anyone attempting to gather information in preparation for an attack. In this ScreenOS release, policy-based NAT-src and NAT-dst offer a single approach that can duplicate and surpass the functionality of interface-based MIPs and VIPs.

5. On some NetScreen devices, you can define a VIP to be the same as an interface IP address. This ability is convenient when the NetScreen device only has one assigned IP address, and when the IP address is assigned dynamically.

Policy-Based Translation Options

NetScreen provides the following ways to apply source and destination network address translation (NAT-src) and (NAT-dst). Note that you can always combine NAT-src with NAT-dst within the same policy.

NAT-Src from a DIP Pool with PAT – The NetScreen device translates the original source IP address to an address drawn from a dynamic IP (DIP) pool. The NetScreen device also applies source port address translation (PAT). For more information, see "NAT-Src from a DIP Pool with PAT Enabled" on page 260.

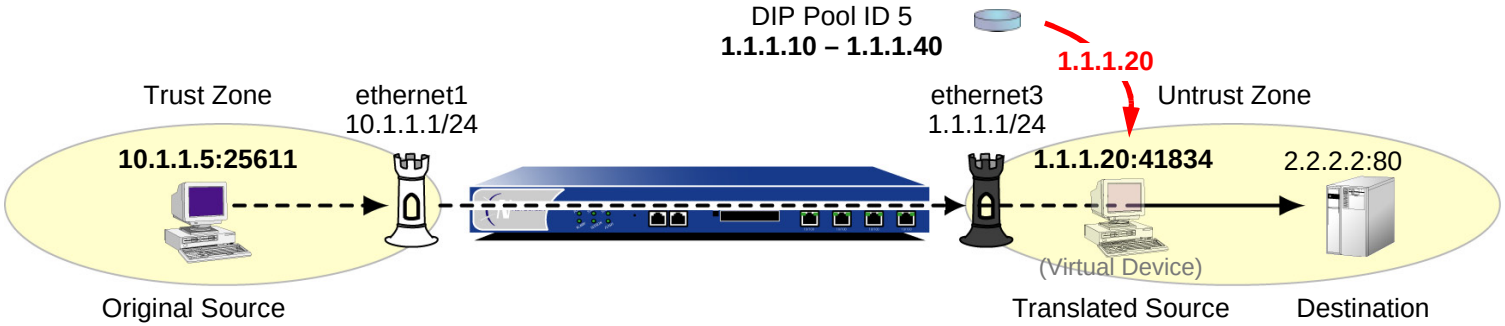

Note: *In this and in subsequent illustrations a "virtual device" is used to represent a translated source or destination address when that address does not belong to an actual device.*

NAT-Src from a DIP Pool without PAT – The NetScreen device translates the original source IP address to an address drawn from a DIP pool. The NetScreen device does not apply source PAT. For more information, see "NAT-Src from a DIP Pool with PAT Disabled" on page 264.

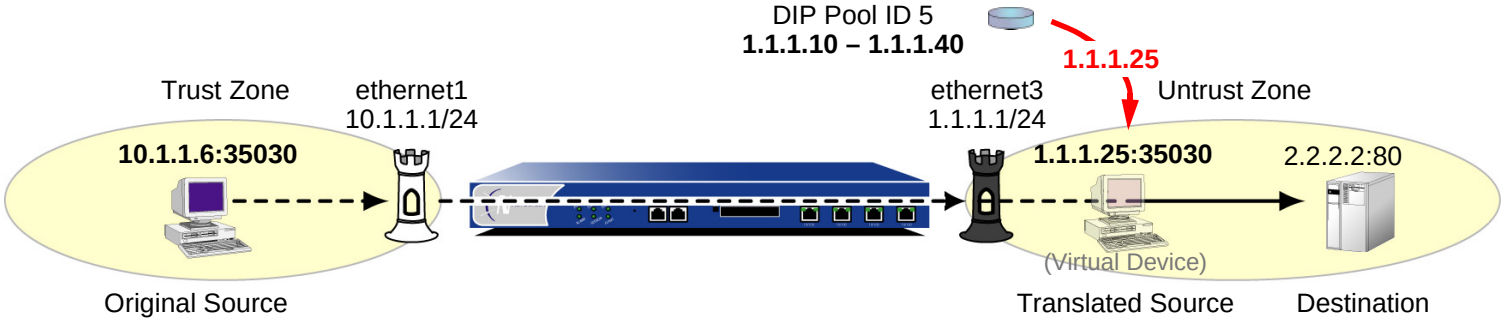

NAT-Src from a DIP Pool with Address Shifting – The NetScreen device translates the original source IP address to an address drawn from a dynamic IP (DIP) pool, consistently mapping each original address to a particular translated address. The NetScreen device does not apply source port address translation (PAT). For more information, see "NAT-Src from a DIP Pool with Address Shifting" on page 267.

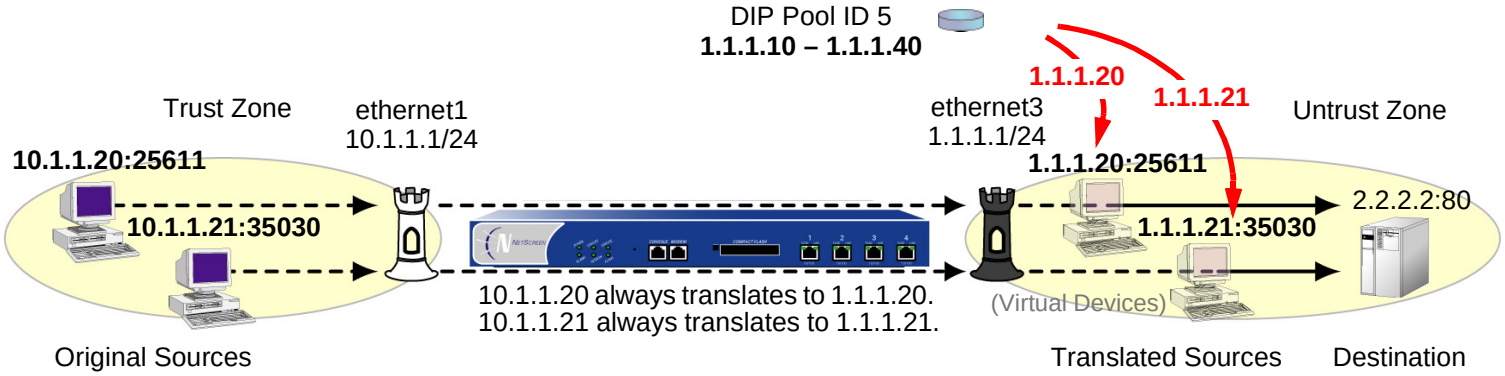

NAT-Src from the Egress Interface IP Address – The NetScreen device translates the original source IP address to the address of the egress interface. The NetScreen device applies source PAT as well. For more information, see "NAT-Src from the Egress Interface IP Address" on page 273.

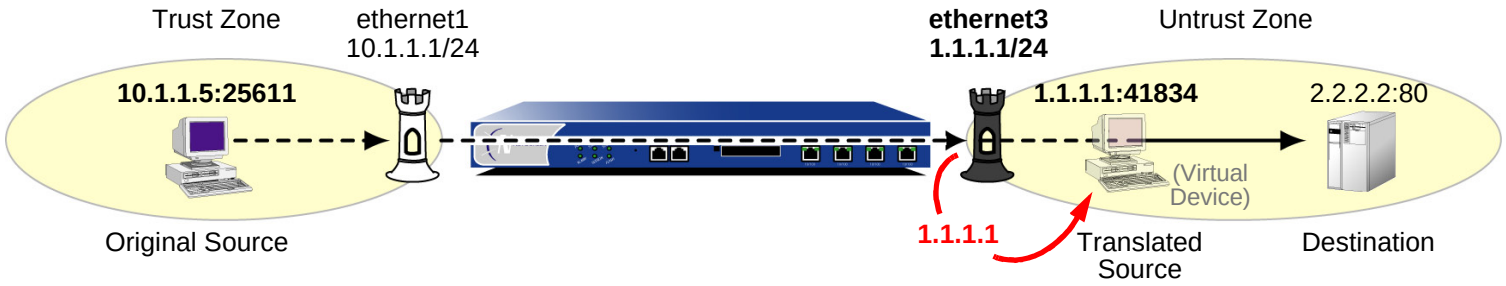

NAT-Dst to a Single IP Address with Port Mapping – The NetScreen device performs destination network address translation (NAT-dst) and destination port mapping. For more information, see "NAT-Dst with Port Mapping" on page 305.

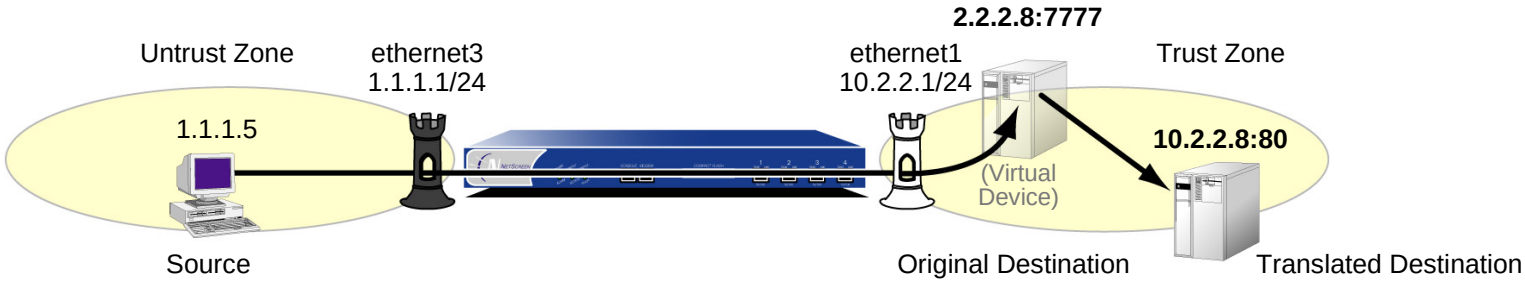

NAT-Dst to a Single IP Address without Port Mapping– The NetScreen device performs NAT-dst but does not change the original destination port number. For more information, see "Destination Network Address Translation" on page 276.

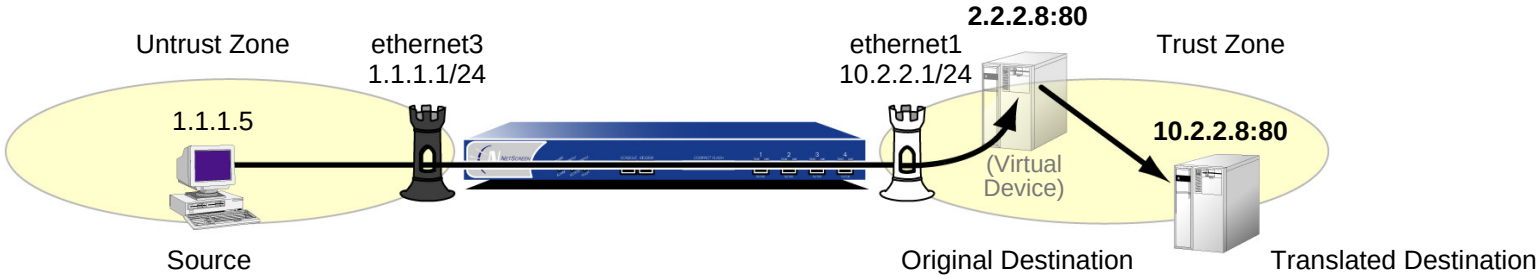

NAT-Dst from an IP Address Range to a Single IP Address– The NetScreen device performs NAT-dst to translate a range of IP addresses to a single IP address. If you also enable port mapping, the NetScreen device translates the original destination port number to another number. For more information, see "NAT-Dst: Many-to-One Mapping" on page 295.

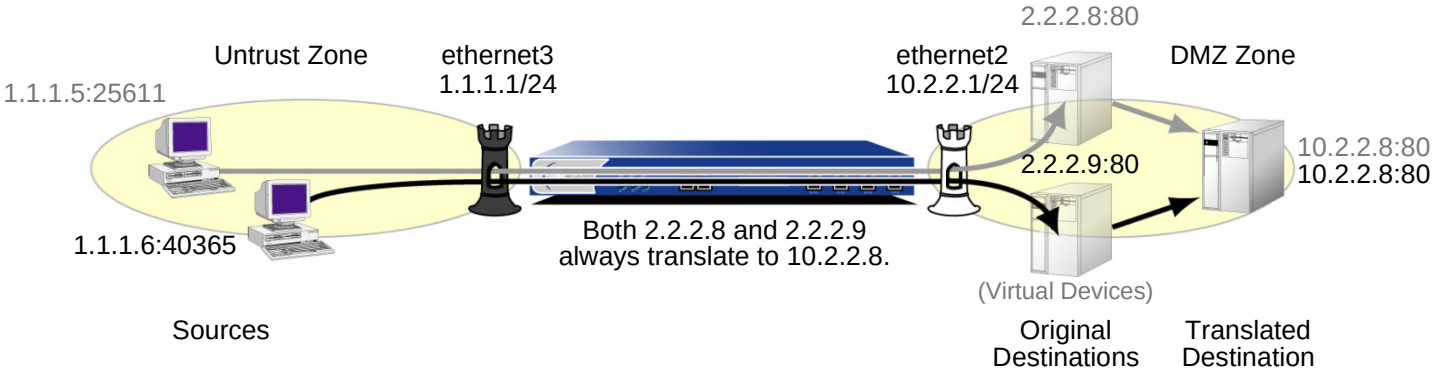

NAT-Dst between IP Address Ranges – When you apply NAT-dst for a range of IP addresses, the NetScreen device maintains a consistent mapping of an original destination address to a translated address within the specified range using a technique called address shifting. Note that address shifting does not support port mapping. For more information, see "NAT-Dst: Many-to-Many Mapping" on page 300.

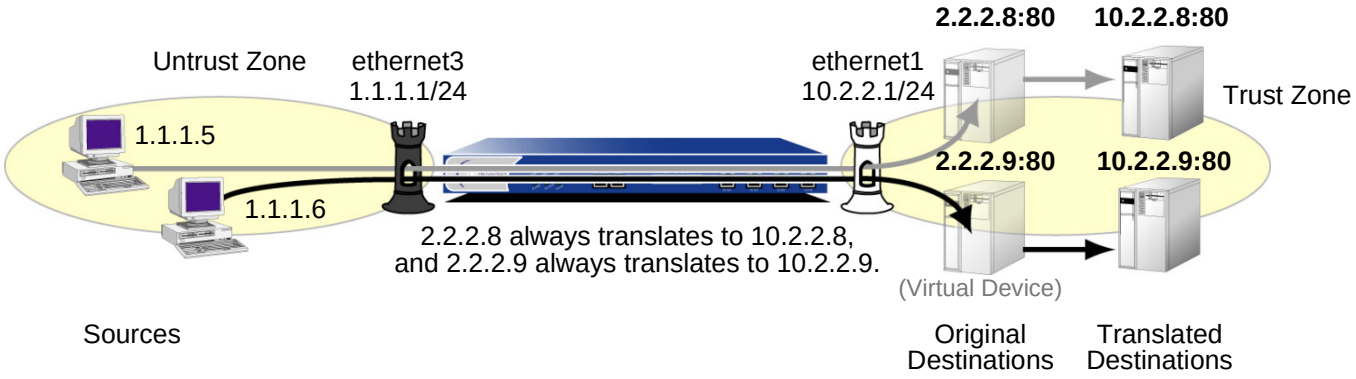

Directional Nature of NAT-Src and NAT-Dst

The application of NAT-src is separate from that of NAT-dst. You determine their applications on traffic by the direction indicated in a policy. For example, if the NetScreen device applies a policy requiring NAT-dst for traffic sent from host A to virtual host B, the NetScreen device translates the original destination IP address from 2.2.2.2 to 3.3.3.3. (It also translates the source IP address from 3.3.3.3 to 2.2.2.2 in responding traffic.)

set policy from “zone A” to “zone B” “host A” “virtual host B” any nat dst ip 3.3.3.3 permit
set vrouter trust-vr route 2.2.2.2/32 interface ethernet1

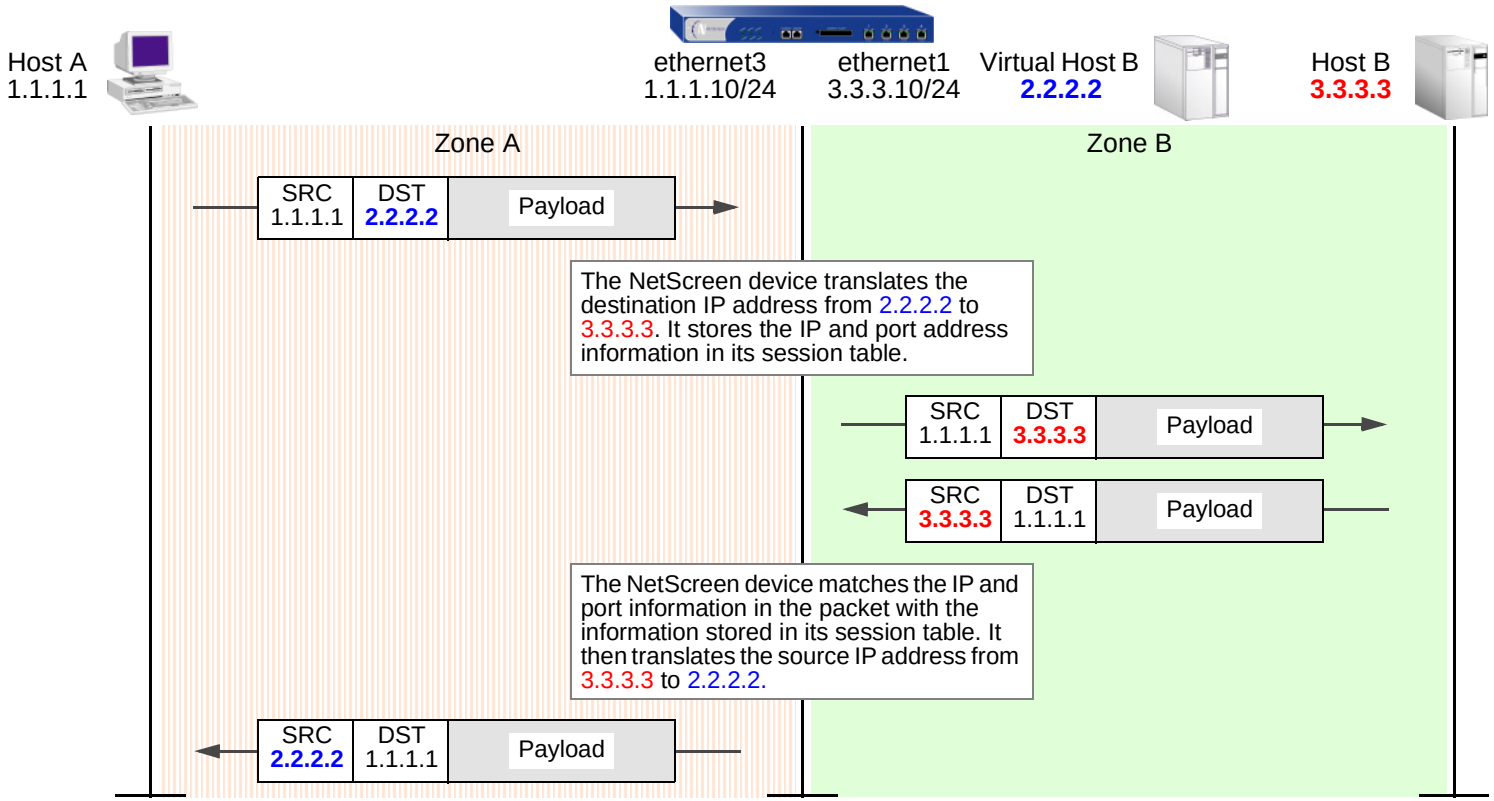

Note: *You must set a route to 2.2.2.2/32 (virtual host B) so the NetScreen device can do a route lookup to determine the destination zone. For more about NAT-dst routing issues, see* *“Routing for Destination Translation” on page 282**.*

However, if you only create the above policy specifying NAT-dst from host A to host B, the NetScreen device does not translate the original source IP address of host B if host B initiates traffic to host A, rather than responding to traffic from host A. For the NetScreen device to do translate the source IP address of host B when it initiates traffic to host A, you must configure a second policy from host B to host A specifying NAT-src[6]. (This behavior differs from that of MIPs and VIPs. See "Mapped IP Addresses" on page 331 and "Virtual IP Addresses" on page 356.)

```
set interface ethernet1 dip-id 7 1.1.1.3 1.1.1.3
set policy from "zone B" to "zone A" "host B" "host A" any nat src dip-id 7 permit
```

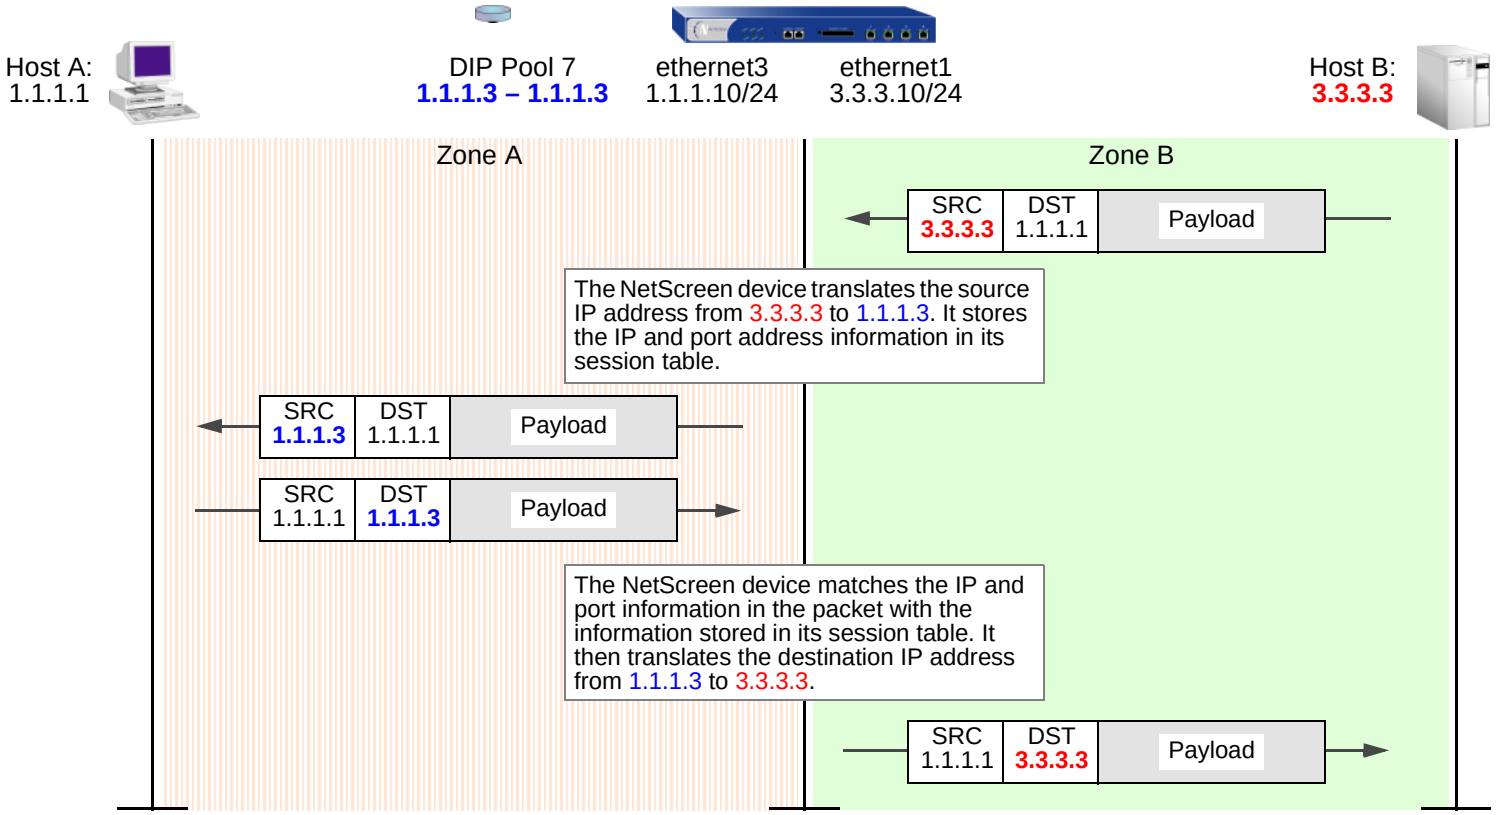

6. To retain focus on the IP address translation mechanisms, port address translation (PAT) is not shown. If you specify fixed port numbers for a DIP pool consisting of a single IP address, then only one host can use that pool at a time. The policy above specifies only "host B" as the source address. If "host B" is the only host that uses DIP pool 7, then it is unnecessary to enable PAT.

SOURCE NETWORK ADDRESS TRANSLATION

It is sometimes necessary for the NetScreen device to translate the original source IP address in an IP packet header to another address. For example, when hosts with private IP addresses initiate traffic to a public address space, the NetScreen device must translate the private source IP address to a public one[7]. Also, when sending traffic from one private address space through a VPN to a site using the same addresses, the NetScreen devices at both ends of the tunnel must translate the source and destination IP addresses to mutually neutral addresses.

A dynamic IP (DIP) address pool provides the NetScreen device with a supply of addresses from which to draw when performing source network address translation (NAT-src). When a policy requires NAT-src and references a specific DIP pool, the NetScreen device draws addresses from that pool when performing the translation.

Note: *The DIP pool must use addresses within the same subnet as the default interface in the destination zone referenced in the policy. If you want to use a DIP pool with addresses outside the subnet of the destination zone interface, you must define a DIP pool on an extended interface. For more information, see* *"Extended Interface and DIP" on page 175**.*

The DIP pool can be as small as a single IP address, which, if you enable port address translation (PAT), can support up to ~64,500 hosts concurrently[8]. Although all packets receiving a new source IP address from that pool get the same address, they each get a different port number. By maintaining a session table entry that matches the original address and port number with the translated address and port number, the NetScreen device can track which packets belong to which session, and which sessions belong to which hosts.

If you use NAT-src but do not specify a DIP pool in the policy, the NetScreen device translates the source address to that of the egress interface in the destination zone. In such cases, PAT is required and automatically enabled.

For applications requiring that a particular source port number remain fixed, you must disable PAT and define a DIP pool with a range of IP addresses large enough for each concurrently active host to receive a different translated address. For fixed-port DIP, the NetScreen device assigns one translated source address to the same host for all its concurrent sessions. In contrast, when the DIP pool has PAT enabled, the NetScreen device might assign a single host different addresses for different concurrent sessions—unless you define the DIP as sticky (see "Sticky DIP Addresses" on page 174).

7. For information about public and private IP addresses, see "Public IP Addresses" on page 77 and "Private IP Addresses" on page 78.

8. When PAT is enabled, the NetScreen device also maintains a pool of free port numbers to assign along with addresses from the DIP pool. The figure of ~64,500 is derived by subtracting 1023, the numbers reserved for the well-known ports, from the maximum number of ports, which is 65,535.

NAT-Src from a DIP Pool with PAT Enabled

When applying source network address translation (NAT-src) with port address translation (PAT), the NetScreen device translates IP addresses and port numbers, and performs stateful inspection as illustrated below (note that only the elements in the IP packet and TCP segment headers relevant to NAT-src are shown):

set policy from trust to untrust any any http nat src dip-id 5 permit

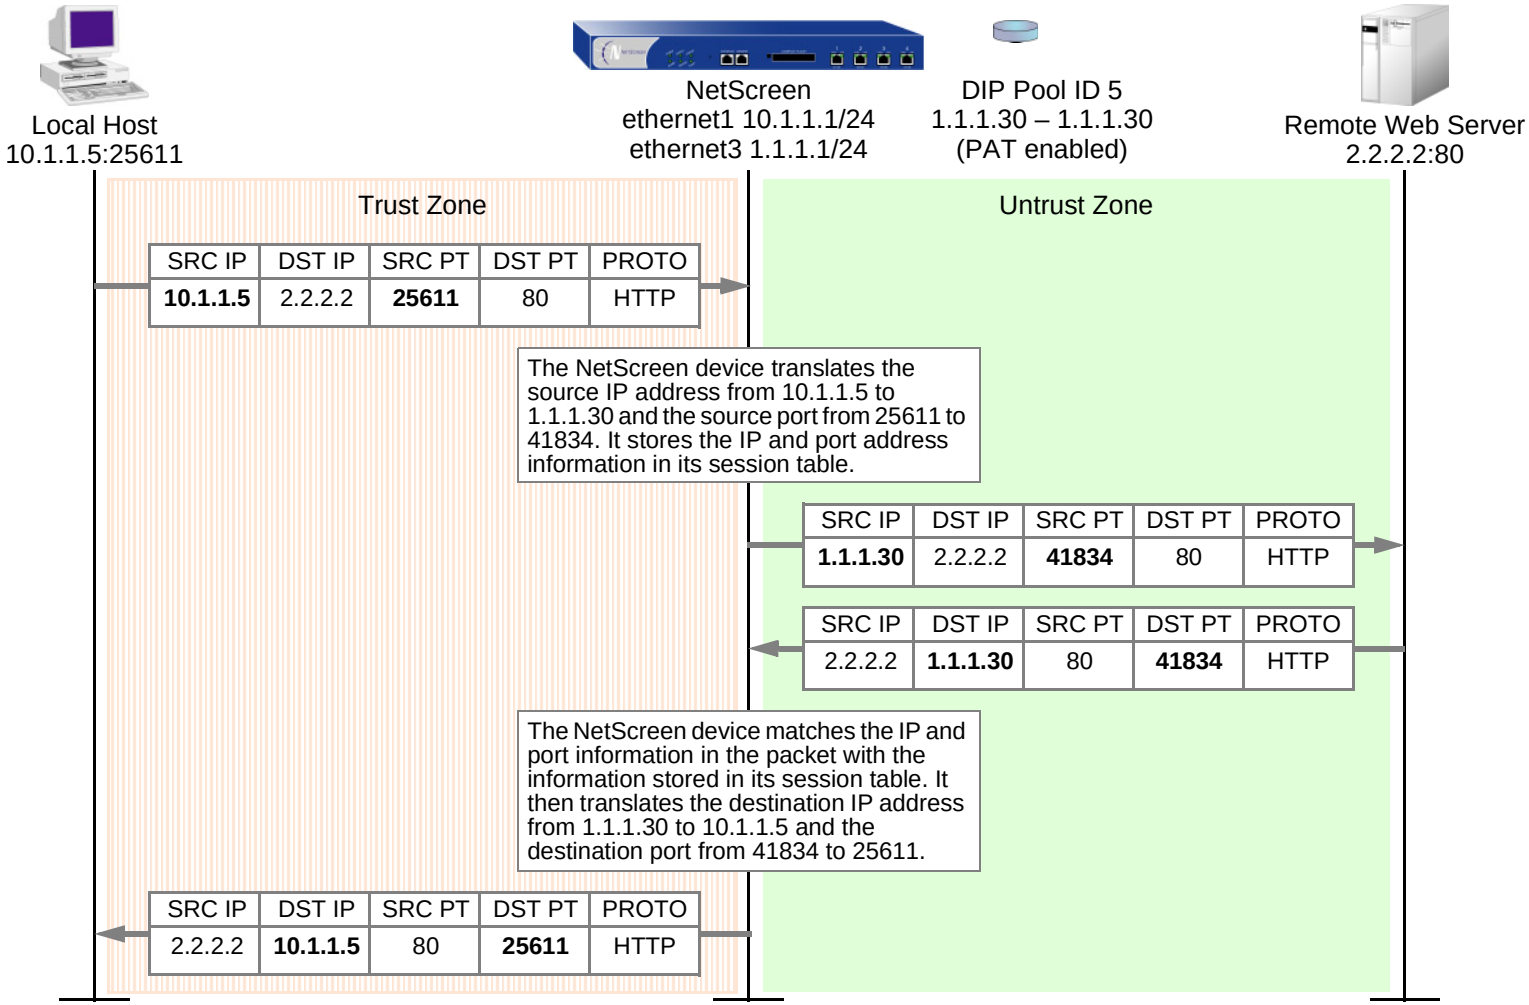

Example: NAT-Src with PAT Enabled

In this example, you define a DIP pool 5 on ethernet3, an interface bound to the Untrust zone. The DIP pool contains a single IP address—1.1.1.30—and has PAT enabled, which it is by default[9]. You then set a policy that instructs the NetScreen device to perform the following tasks:

- Permit HTTP traffic from any address in the Trust zone to any address in the Untrust zone
- Translate the source IP address in the IP packet header to 1.1.1.30, which is the sole entry in DIP pool 5
- Translate the original source port number in the TCP segment header or UDP datagram header to a new, unique number
- Send HTTP traffic with the translated source IP address and port number out ethernet3 to the Untrust zone

Source Network Address Translation (NAT-src)
(Translates source IP address to a one-address DIP pool—1.1.1.30. PAT is enabled.)

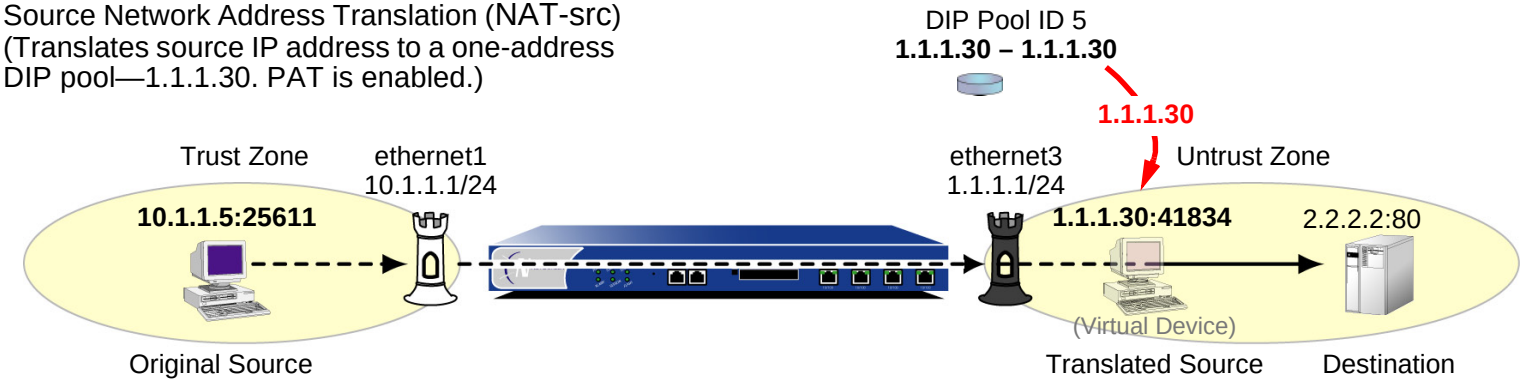

SRC IP	DST IP	SRC PT	DST PT	PROTO
10.1.1.5	2.2.2.2	**25611**	80	HTTP

SRC IP	DST IP	SRC PT	DST PT	PROTO
1.1.1.30	2.2.2.2	**41834**	80	HTTP

9. When you define a DIP pool, the NetScreen device enables PAT by default. To disable PAT, you must add the key word fix-port to the end of the CLI command, or clear the Port Translation option on the DIP configuration page in the WebUI. For example, **set interface etherenet3 dip 5 1.1.1.30 1.1.1.30 fix-port**, or Network > Interfaces > Edit (for ethernet3) > DIP: ID: 5; Start: 1.1.1.30; End: 1.1.1.30; Port Translation: (clear).

WebUI

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Select the following, and then click **OK**:
Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

2. DIP

Network > Interfaces > Edit (for ethernet3) > DIP > New: Enter the following, and then click **OK**:

ID: 5

IP Address Range: (select), 1.1.1.30 ~ 1.1.1.30

Port Translation: (select)

In the same subnet as the interface IP or its secondary IPs: (select)

3. Policy

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Any

Service: HTTP

Action: Permit

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

NAT:

Source Translation: (select)

(DIP on): 5 (1.1.1.30 - 1.1.1.30)/X-late

CLI

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
```

2. DIP

```
set interface ethernet3 dip 5 1.1.1.30 1.1.1.30
```

3. Policy

```
set policy from trust to untrust any any http nat src dip-id 5 permit
save
```

NAT-Src from a DIP Pool with PAT Disabled

The occasion can arise when you want to perform source network address translation (NAT-src) for the IP address but not port address translation (PAT) for the source port number. Perhaps a custom application requires a specific that the source port address be a specific number. Perhaps the target host requires that the source IP address and port address be certain numbers to uniquely identify the host. In such cases, you can define a policy instructing the NetScreen device to perform NAT-src without PAT.

Example: NAT-Src with PAT Disabled

In this example, you define a DIP pool 6 on ethernet3, an interface bound to the Untrust zone. The DIP pool contains a range of IP addresses from 1.1.1.50 to 1.1.1.150. You disable PAT. You then set a policy that instructs the NetScreen device to perform the following tasks:

- Permit traffic for a user-defined service named “e-stock” from any address in the Trust zone to any address in the Untrust zone[10]
- Translate the source IP address in the IP packet header to any available address in DIP pool 6
- Retain the original source port number in the TCP segment header or UDP datagram header
- Send e-stock traffic with the translated source IP address and original port number out ethernet3 to the Untrust zone

WebUI

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

10. It is assumed that you have previously defined the user-defined service “e-stock”. This fictional service requires that all e-stock transactions originate from specific source port numbers. For this reason, you must disable PAT for DIP pool 6.

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

2. DIP

Network > Interfaces > Edit (for ethernet3) > DIP > New: Enter the following, and then click **OK**:

ID: 6

IP Address Range: (select), 1.1.1.50 ~ 1.1.1.150

Port Translation: (clear)

In the same subnet as the interface IP or its secondary IPs: (select)

3. Policy

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Any

Service: e-stock

Action: Permit

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

NAT:

Source Translation: (select)

DIP on: (select), 6 (1.1.1.50 - 1.1.1.150)

CLI

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
```

2. DIP

```
set interface ethernet3 dip 6 1.1.1.50 1.1.1.150 fix-port
```

3. Policy

```
set policy from trust to untrust any any e-stock nat src dip-id 6 permit
save
```

NAT-Src from a DIP Pool with Address Shifting

You can define a one-to-one mapping from an original source IP address to a translated source IP address for a range of IP addresses. Such a mapping ensures that the NetScreen device always translates a particular source IP address from within that range to the same translated address within a DIP pool. There can be any number of addresses in the range. You can even map one subnet to another subnet, with a consistent one-to-one mapping of each original address in one subnet to its translated counterpart in the other subnet.

One possible use for performing NAT-src with address shifting is to provide greater policy granularity on another NetScreen device that receives traffic from the first one. For example, the admin for NetScreen-A at site A defines a policy that translates the source addresses of its hosts when communicating with NetScreen-B at site B through a site-to-site VPN tunnel. If NetScreen-A applies NAT-src using addresses from a DIP pool without address shifting, the NetScreen-B admin can only configure generic policies regarding the traffic it can allow from site A. Unless the NetScreen-B admin knows the specific translated IP addresses, he can only set inbound policies for the range of source addresses drawn from the NetScreen-A DIP pool. On the other hand, if the NetScreen-B admin knows what the translated source addresses are (because of address shifting), the NetScreen-B admin can now be more selective and restrictive with the policies he sets for inbound traffic from site A.

Note that it is possible to use a DIP pool with address shifting enabled in a policy that applies to source addresses beyond the range specified in the pool. In such cases, the NetScreen device passes traffic from all source addresses permitted in the policy, applying NAT-src with address shifting to those addresses that fall within the DIP pool range but leaving those addresses that fall outside the DIP pool range unchanged. If you want the NetScreen device to apply NAT-src to all source addresses, make sure that the range of source addresses is smaller or the same size as the range of the DIP pool.

Note: *The NetScreen device does not support source port address translation (PAT) with address shifting.*

Example: NAT-Src with Address Shifting

In this example, you define DIP pool 10 on ethernet3, an interface bound to the Untrust zone. You want to translate five addresses between 10.1.1.11 and 10.1.1.15 to five addresses between 1.1.1.101 and 1.1.1.105, and you want the relationship between each original and translated address to be consistent:

Original Source IP Address		Translated Source IP Address
10.1.1.11	→	1.1.1.101
10.1.1.12	→	1.1.1.102
10.1.1.13	→	1.1.1.103
10.1.1.14	→	1.1.1.104
10.1.1.15	→	1.1.1.105

You define addresses for five hosts in the Trust zone and added them to an address group named “group1”. The addresses for these hosts are 10.1.1.11, 10.1.1.12, 10.1.1.13, 10.1.1.14, and 10.1.1.15. You configure a policy from the Trust zone to the Untrust zone that references that address group in a policy to which you apply NAT-src with DIP pool 10. The policy instructs the NetScreen device to perform NAT-src whenever a member of goup1 initiates HTTP traffic to an address in the Untrust zone. Furthermore, the NetScreen device always performs NAT-src from a particular IP address—such as 10.1.1.13—to the same translated IP address—1.1.1.103.

You then set a policy that instructs the NetScreen device to perform the following tasks:

- Permit HTTP traffic from group1 in the Trust zone to any address in the Untrust zone
- Translate the source IP address in the IP packet header to its corresponding address in DIP pool 10
- Send HTTP traffic with the translated source IP address and port number out ethernet3 to the Untrust zone

WebUI

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Select the following, and then click **OK**:
Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

2. DIP

Network > Interfaces > Edit (for ethernet3) > DIP > New: Enter the following, and then click **OK**:

ID: 10

IP Shift: (select)

From: 10.1.1.11

To: 1.1.1.101 ~ 1.1.1.105

In the same subnet as the interface IP or its secondary IPs: (select)

3. Addresses

Objects > Addresses > List > New: Enter the following information, and then click **OK**:

Address Name: host1

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.11/32

Zone: Trust

Objects > Addresses > List > New: Enter the following information, and then click **OK**:

Address Name: host2

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.12/32

Zone: Trust

Objects > Addresses > List > New: Enter the following information, and then click **OK**:

Address Name: host3

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.13/32

Zone: Trust

Objects > Addresses > List > New: Enter the following information, and then click **OK**:

Address Name: host4

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.14/32

Zone: Trust

Objects > Addresses > List > New: Enter the following information, and then click **OK**:

Address Name: host5

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.15/32

Zone: Trust

Objects > Addresses > Group > (for Zone: Trust) New: Enter the following group name, move the following addresses, and then click **OK**:

Group Name: group1

Select **host1** and use the **<<** button to move the address from the Available Members column to the Group Members column.

Select **host2** and use the **<<** button to move the address from the Available Members column to the Group Members column.

Select **host3** and use the **<<** button to move the address from the Available Members column to the Group Members column.

Select **host4** and use the **<<** button to move the address from the Available Members column to the Group Members column.

Select **host5** and use the **<<** button to move the address from the Available Members column to the Group Members column.

4. Policy

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), group1

Destination Address:

Address Book Entry: (select), Any

Service: HTTP

Action: Permit

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

NAT:

Source Translation: (select)

(DIP on): 10 (1.1.1.101 - 1.1.1.105)

CLI

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
```

2. DIP

```
set interface ethernet3 dip 10 shift-from 10.1.1.11 to 1.1.1.101 1.1.1.105
```

3. Addresses

```
set address trust host1 10.1.1.11/32
set address trust host2 10.1.1.12/32
set address trust host3 10.1.1.13/32
set address trust host4 10.1.1.14/32
set address trust host5 10.1.1.15/32

set group address trust group1 add host1
set group address trust group1 add host2
set group address trust group1 add host3
set group address trust group1 add host4
set group address trust group1 add host5
```

4. Policy

```
set policy from trust to untrust group1 any http nat src dip-id 10 permit
save
```

NAT-Src from the Egress Interface IP Address

If you apply NAT-src to a policy but do not specify a DIP pool, then the NetScreen device translates the source IP address to the address of the egress interface. In such cases, the NetScreen device always applies PAT.

Example: NAT-Src without DIP

In this example, you define a policy that instructs the NetScreen device to perform the following tasks:

- Permit HTTP traffic from any address in the Trust zone to any address in the Untrust zone
- Translate the source IP address in the IP packet header to 1.1.1.1, which is the IP address of ethernet3, the interface bound to the Untrust zone, and thus the egress interface for traffic sent to any address in the Untrust zone
- Translate the original source port number in the TCP segment header or UDP datagram header to a new, unique number
- Send traffic with the translated source IP address and port number out ethernet3 to the Untrust zone

Source Network Address Translation (NAT-src)
(Translates the source IP address to the egress interface IP address—1.1.1.1—in the destination zone. PAT is enabled.)

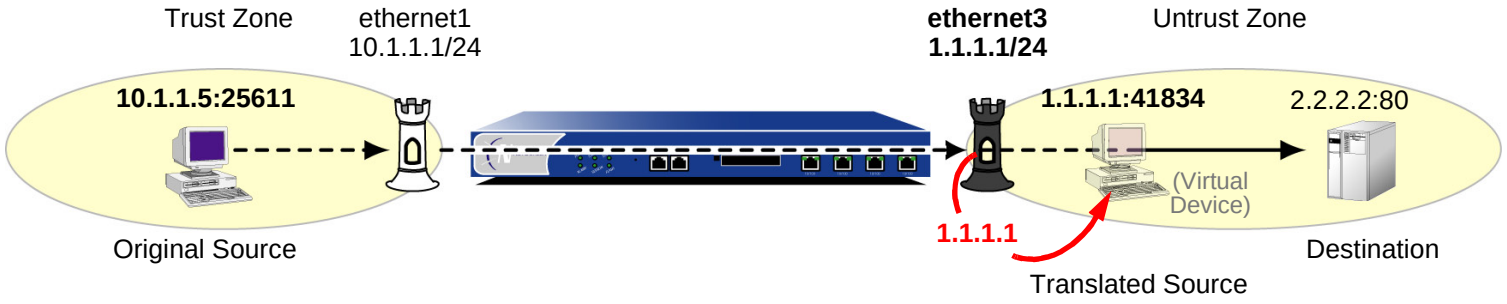

SRC IP	DST IP	**SRC PT**	DST PT	PROTO
10.1.1.5	2.2.2.2	**25611**	80	HTTP

SRC IP	DST IP	**SRC PT**	DST PT	PROTO
1.1.1.1	2.2.2.2	**41834**	80	HTTP

WebUI

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

2. Policy

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Any

Service: HTTP

Action: Permit

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

NAT:

Source Translation: (select)

(DIP on): None (Use Egress Interface IP)

CLI

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
```

2. Policy

```
set policy from trust to untrust any any http nat src permit
save
```

DESTINATION NETWORK ADDRESS TRANSLATION

You can define policies to translate the destination address from one IP address to another. Perhaps you need the NetScreen device to translate one or more public IP addresses to one or more private addresses. The relationship of the original destination address to the translated destination address can be a one-to-one relationship, a many-to-one relationship, or a many-to-many relationship. The following illustration depicts the concepts of one-to-one and many-to-one NAT-dst relationships.

"NAT-Dst: One-to-One Mapping" on page 286

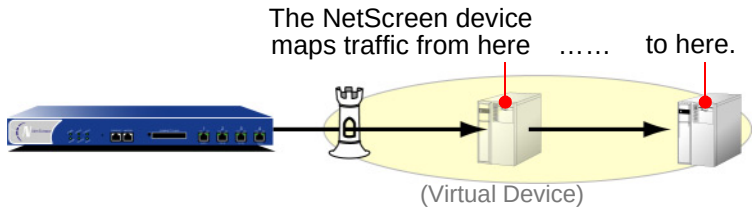

"NAT-Dst: Many-to-One Mapping" on page 295

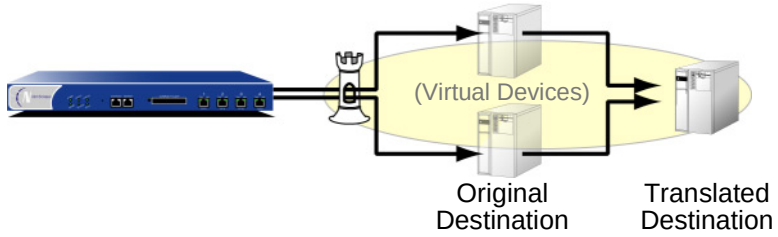

Note: *The original destination IP address and the translated destination IP address must be in the same security zone.*

Both of the above configurations support destination port mapping. Port mapping is the deterministic translation of one original destination port number to another specific number. The relationship of the original-to-translated number in port mapping differs from port address translation (PAT). With port mapping, the NetScreen device translates a predetermined original port number to another predetermined port number. With PAT, the NetScreen device translates a randomly assigned original source port number to another randomly assigned number.

You can translate a range of destination addresses to another range—such as one subnet to another—with address shifting, so that the NetScreen device consistently maps each original destination address to a specific translated destination address. Note that NetScreen does not support port mapping with address shifting. The following illustration depicts the concept of a many-to-many relationship for NAT-dst.

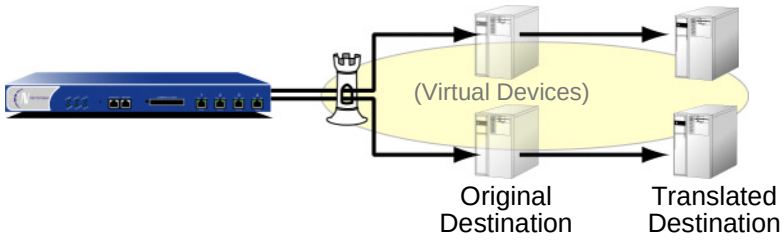

"NAT-Dst: Many-to-Many Mapping" on page 300

There must be entries in the route table for both the original destination IP address and the translated destination IP address. The NetScreen device performs a route lookup using the original destination IP address to determine the destination zone for a subsequent policy lookup. It then performs a second route lookup using the translated address to determine where to send the packet. To ensure that the routing decision is in accord with the policy, both the original destination IP address and the translated IP address must be in the same security zone. (For more information about the relationship of the destination IP address, route lookup, and policy lookup, see "Packet Flow for Destination Translation" on page 278.)

Packet Flow for Destination Translation

The following steps describe the path of a packet through a NetScreen device and the various operations that it performs when applying destination network address translation.

1. An HTTP packet with source IP address:port number 1.1.1.5:32455 and destination IP address:port number 5.5.5.5:80 arrives at ethernet1, which is bound to the Untrust zone.

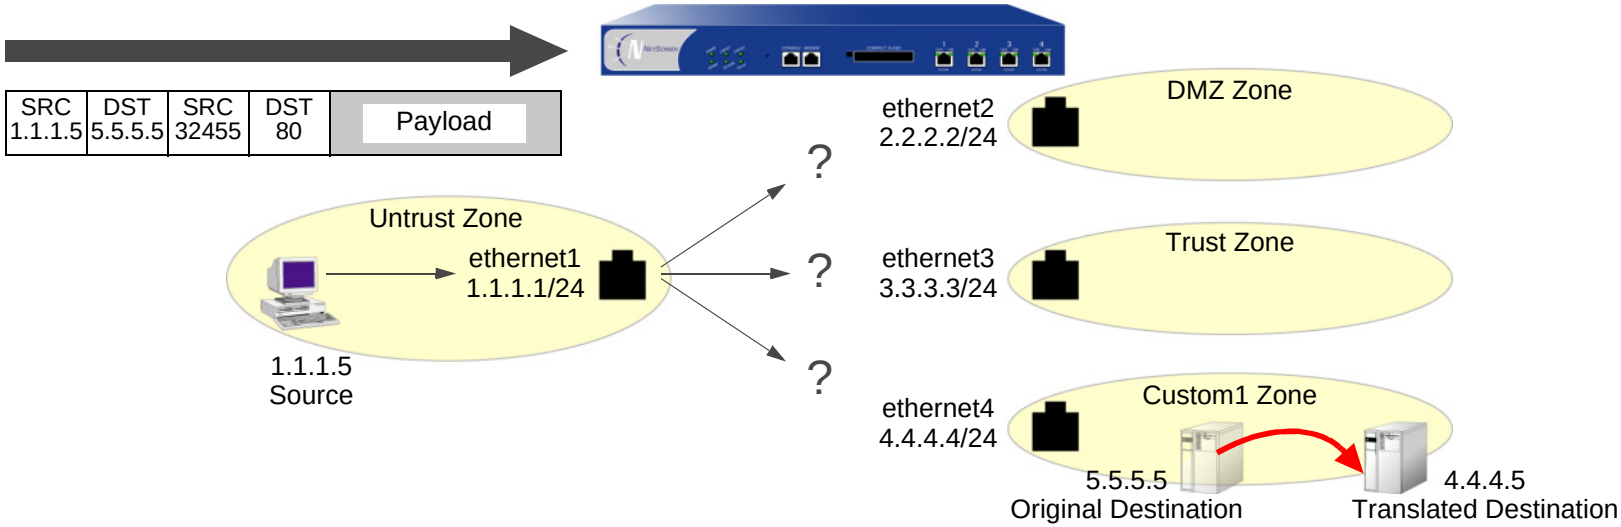

The NetScreen device has not yet performed the steps required to learn which interface it must use to forward the packet. This is indicated in the illustration by the three question marks.

2. If you have enabled SCREEN options for the Untrust zone, the NetScreen device activates the SCREEN module at this point. SCREEN checking can produce one of the following three results:
   - If a SCREEN mechanism detects anomalous behavior for which it is configured to block the packet, the NetScreen device drops the packet and makes an entry in the event log.
   - If a SCREEN mechanism detects anomalous behavior for which it is configured to record the event but not block the packet, the NetScreen device records the event in the SCREEN counters list for the ingress interface and proceeds to the next step.
   - If the SCREEN mechanisms detect no anomalous behavior, the NetScreen device proceeds to the next step.

If you have not enabled any SCREEN options for the Untrust zone, the NetScreen device immediately proceeds to the next step.

3. The session module performs a session lookup, attempting to match the packet with an existing session.

   If the packet does not match an existing session, the NetScreen device performs First Packet Processing, a procedure involving the remaining steps.

   If the packet matches an existing session, the NetScreen device performs Fast Processing, using the information available from the existing session entry to process the packet. Fast Processing bypasses all but the last step because the information generated by the bypassed steps has already been obtained during the processing of the first packet in the session.

4. The address-mapping module checks if a mapped IP (MIP) or virtual IP (VIP) configuration uses the destination IP address 5.5.5.5.

   If there is such a configuration, the NetScreen device resolves the MIP or VIP to the translated destination IP address and bases its route lookup on that. It then does a policy lookup between the Untrust and Global zones. If it finds a policy match that permits the traffic, the NetScreen device forwards the packet out the egress interface determined in the route lookup.

   If 5.5.5.5 is not used in a MIP or VIP configuration, the NetScreen device proceeds to the next step.

5. To determine the destination zone, the route module does a route lookup of the original destination IP address; that is, it uses the destination IP address that appears in the header of the packet that arrives at ethernet1. (The route module uses the ingress interface to determine which virtual router to use for the route lookup.) It discovers that 5.5.5.5/32 is accessed through ethernet4, which is bound to the Custom1 zone.

trust-vr Route Table			
To Reach:	**Use Interface:**	**In Zone:**	**Use Gateway:**
0.0.0.0/0	ethernet1	Untrust	1.1.1.250
1.1.1.0/24	ethernet1	Untrust	0.0.0.0
2.2.2.0/24	ethernet2	DMZ	0.0.0.0
3.3.3.0/24	ethernet3	Trust	0.0.0.0
4.4.4.0/24	etherne4	Custom1	0.0.0.0
5.5.5.5/32	**ethernet4**	**Custom1**	**0.0.0.0**

6. The policy engine does a policy lookup between the Untrust and Custom1 zones (as determined by the corresponding ingress and egress interfaces). The source and destination IP addresses and the service match a policy redirecting HTTP traffic from 5.5.5.5 to 4.4.4.5.

   set policy from untrust to custom1 any v-server1 http nat dst ip 4.4.4.5 permit

   (You have previously defined the address “v-server1” with IP address 5.5.5.5/32. It is in the Custom1 zone.)

   The NetScreen device translates the destination IP address from 5.5.5.5 to 4.4.4.5. The policy indicates that no source network address translation and no destination port address translation are required.

7. The NetScreen device does a second route lookup using the translated IP address and discovers that 4.4.4.5/32 is accessed through ethernet4.

8. The address-mapping module translates the destination IP address in the packet header to 4.4.4.5. The NetScreen device then forwards the packet out ethernet4 and makes an entry in its session table (unless this packet is part of an existing session and an entry already exists).

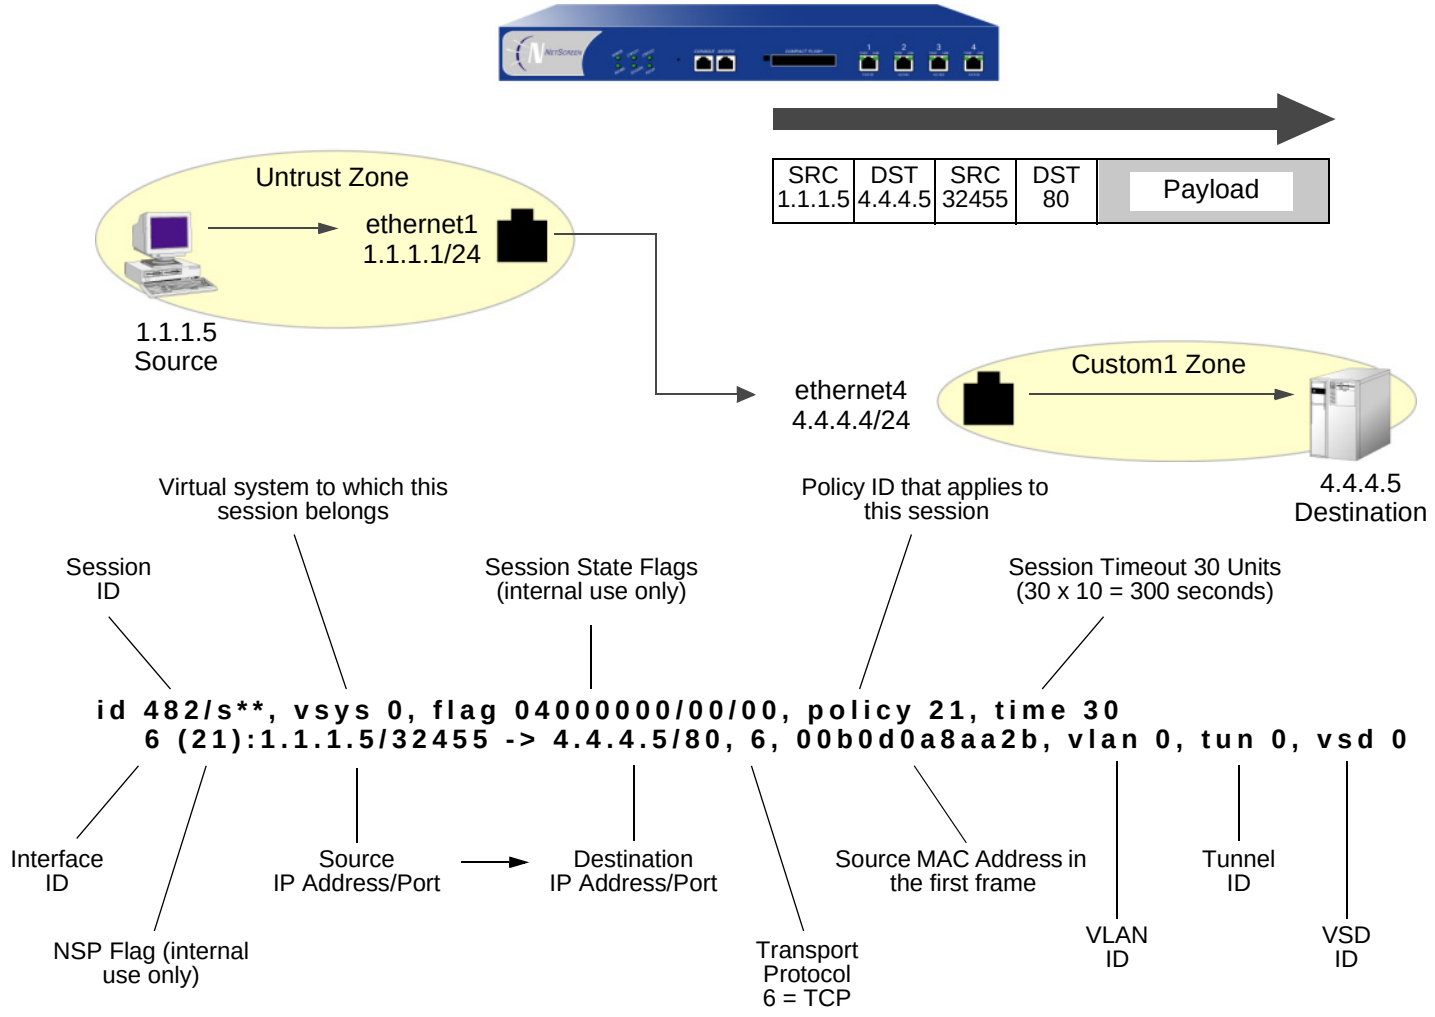

Note: *Because this session does not involve a virtual system, VLAN, VPN tunnel, or virtual security device (VSD), the setting for all these ID numbers is zero.*

Routing for Destination Translation

When you configure addresses for NAT-dst, the NetScreen device must have routes in its routing table to both the original destination address that appears in the packet header and the translated destination address (that is, the address to which the NetScreen device redirects the packet). As explained in "Packet Flow for Destination Translation" on page 278, the NetScreen device uses the original destination address to do a route lookup, and thereby determine the egress interface. The egress interface in turn provides the destination zone—the zone to which the interface is bound—so that the NetScreen device can do a policy lookup. When the NetScreen device finds a policy match, the policy defines the mapping of the original destination address to the translated destination address. The NetScreen device then performs a second route lookup to determine the interface through which it must forward the packet to reach the new destination address. In summary, the route to the original destination address provides a means to perform the policy lookup, and the route to the translated destination address specifies the egress interface through which the NetScreen device is to forward the packet.

In the following three scenarios, the need to enter static routes differs according to the network topology surrounding the destination addresses referenced in this policy:

set policy from untrust to trust any oda1 http nat dst ip 10.1.1.5 permit

in which "oda1" is the original destination address 10.2.1.5, and the translated destination address is 10.1.1.5.

Addresses Connected to the Same Interface

In this scenario, the routes to both the original and translated destination addresses direct traffic through the same interface, ethernet3. The NetScreen device automatically adds a route to 10.1.1.0/24 through ethernet3 when you configure the IP address of the ethernet3 interface as 10.1.1.1/24. To complete the routing requirements, you must add an additional route to 10.2.1.5/32 through ethernet3.

Note: *The Untrust zone is not shown.*

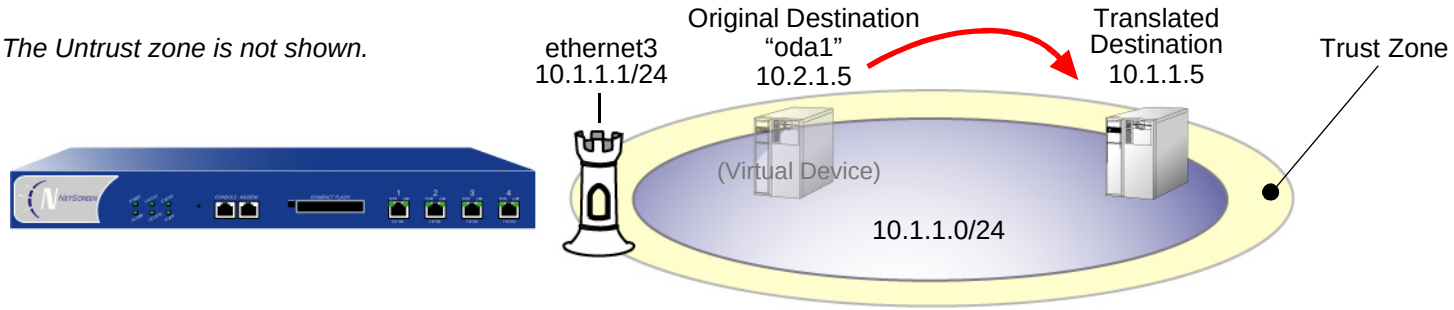

Note: *Although 10.2.1.5 is not in the 10.1.1.0/24 subnet, because its route does not specify a gateway, it is illustrated as if it is in the same connected subnet as the 10.1.1.0/24 address space.*

WebUI

Network > Routing > Routing Entries > (trust-vr) New: Enter the following, and then click **OK**:

Network Address / Netmask: 10.2.1.5/32
Gateway: (select)
Interface: ethernet3
Gateway IP Address: 0.0.0.0

CLI

```
set vrouter trust-vr route 10.2.1.5/32 interface ethernet3
save
```

Addresses Connected to the Same Interface but Separated by a Router

In this scenario, the routes to both the original and translated destination addresses direct traffic through ethernet3. The NetScreen device automatically adds a route to 10.1.1.0/24 through ethernet3 when you configure the IP address of the ethernet3 interface as 10.1.1.1/24. To complete the routing requirements, you must add a route to 10.2.1.0/24 through ethernet3 and the gateway connecting the 10.1.1.0/24 and the 10.2.1.0/24 subnets.

Note: *Because this route is required to reach any address in the 10.2.1.0/24 subnet, you have probably already configured it. If so, no extra route needs to be added just for the policy to apply NAT-dst to 10.2.1.5.*

Note: *The Untrust zone is not shown.*

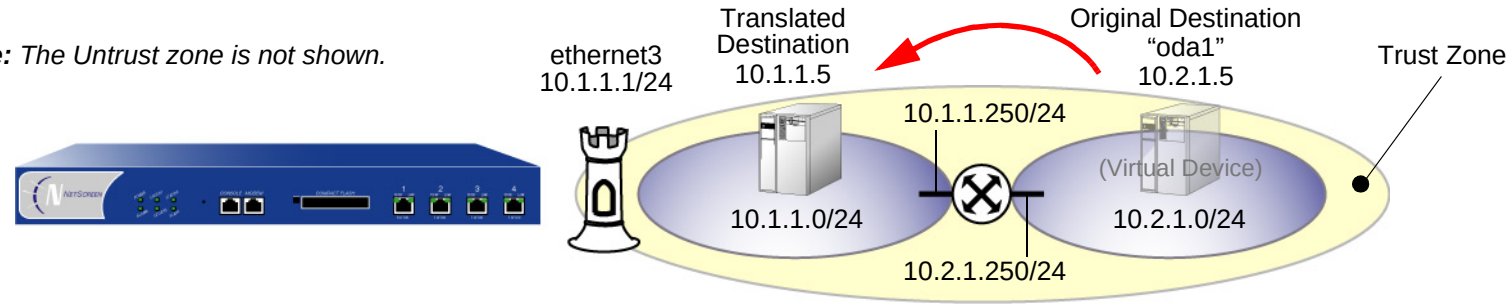

WebUI

Network > Routing > Routing Entries > (trust-vr) New: Enter the following, and then click **OK**:

Network Address / Netmask: 10.2.1.0/24

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 10.1.1.250

CLI

```
set vrouter trust-vr route 10.2.1.0/24 interface ethernet3 gateway 10.1.1.250
save
```

Addresses Separated by an Interface

In this scenario, two interfaces are bound to the Trust zone: ethernet3 with IP address 10.1.1.1/24 and ethernet4 with IP address 10.2.1.1/24. The NetScreen device automatically adds a route to 10.1.1.0/24 through ethernet3 and 10.2.1.0/24 through ethernet4 when you configure the IP addresses of these interfaces. By putting the original destination address in the 10.2.1.0/24 subnet and the translated destination address in the 10.1.1.0/24 subnet, you do not have to add any other routes for the NetScreen device to apply NAT-dst from 10.1.1.5 to 10.2.1.5.

Note: *The Untrust zone is not shown.*

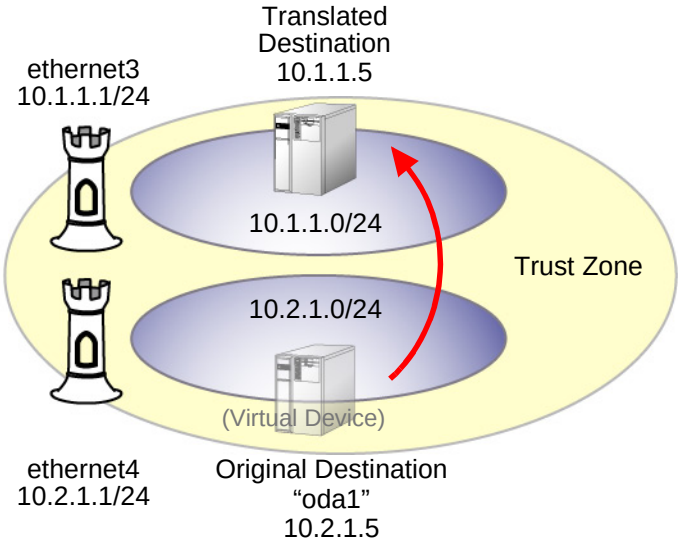

NAT-Dst: One-to-One Mapping

When applying destination network address translation (NAT-dst) without port address translation, the NetScreen device translates the destination IP address and performs stateful inspection as illustrated below (note that only the elements in the IP packet and TCP segment headers relevant to NAT-dst are shown):

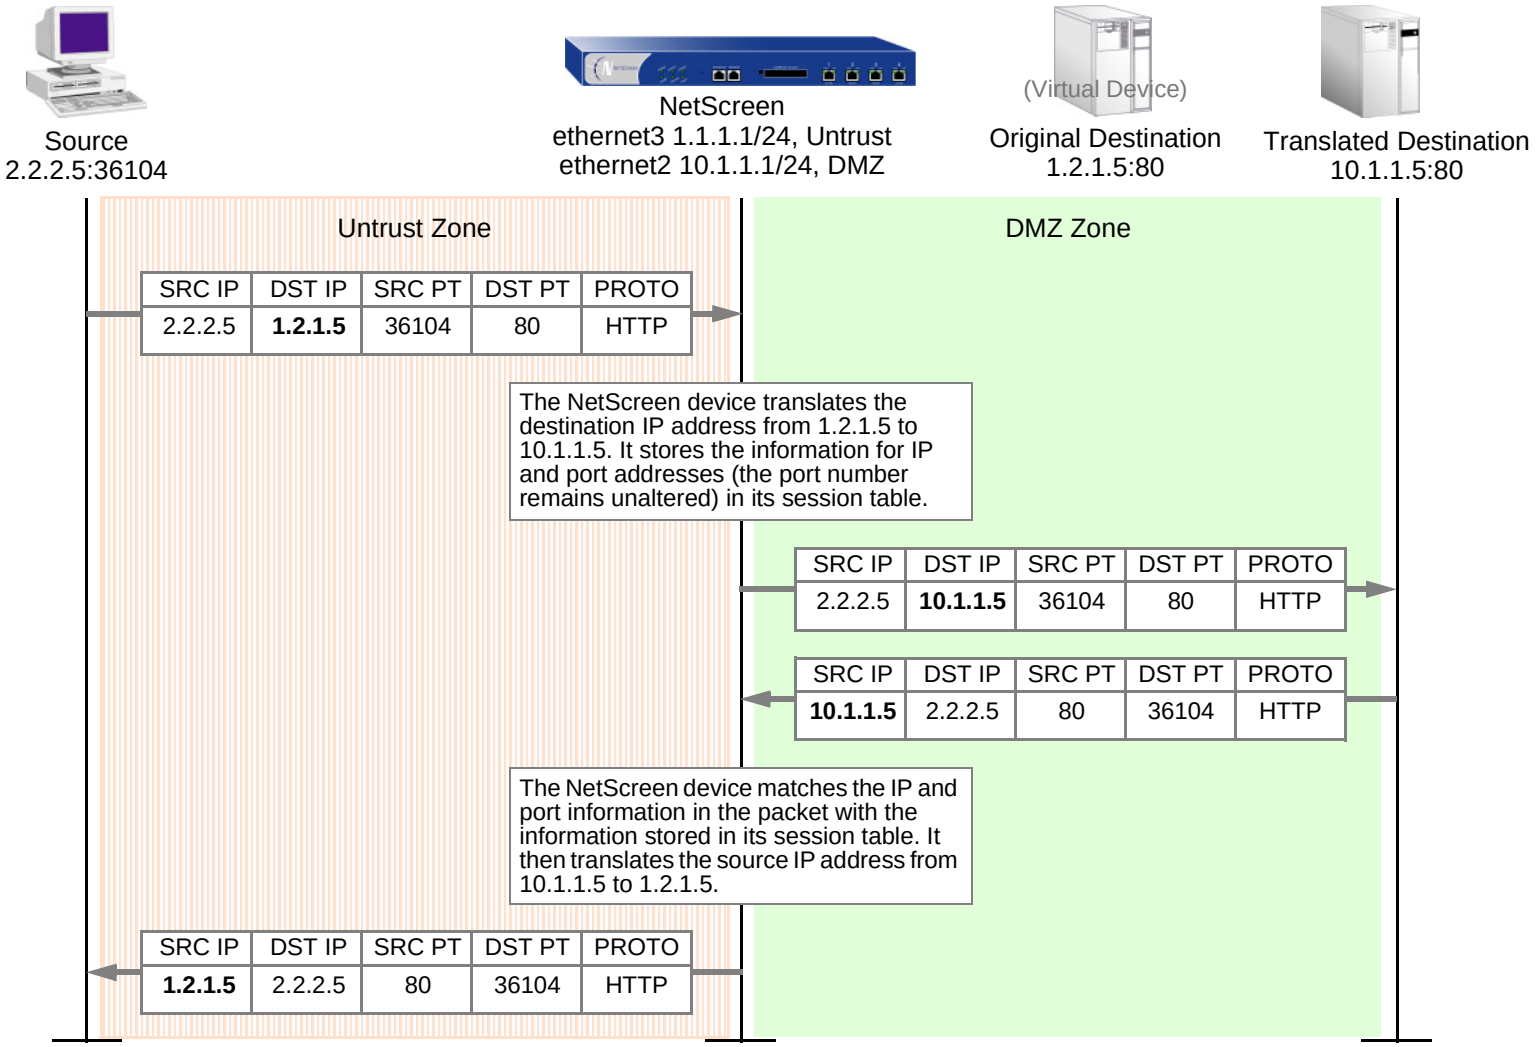

Example: One-to-One Destination Translation

In this example, you set a policy to provide one-to-one destination network address translation (NAT-dst) without changing the destination port addresses. The policy instructs the NetScreen device to perform the following tasks:

- Permit both FTP and HTTP traffic (defined as the service group “http-ftp”) from any address in the Untrust zone to a the original destination address named “oda2” with address 1.2.1.8 in the DMZ zone
- Translate the destination IP address in the IP packet header from 1.2.1.8 to 10.2.1.8
- Leave the original destination port number in the TCP segment header as is (80 for HTTP, and 21 for FTP)
- Forward HTTP and FTP traffic to 10.2.1.8 in the DMZ zone

You bind ethernet3 to the Untrust zone, and assign it IP address 1.1.1.1/24. You bind ethernet2 to the DMZ, and assign it IP address 10.2.1.1/24. You also define a route to the original destination address 1.2.1.8 through ethernet2. Both the Untrust zone and the DMZ zone are in the trust-vr routing domain.

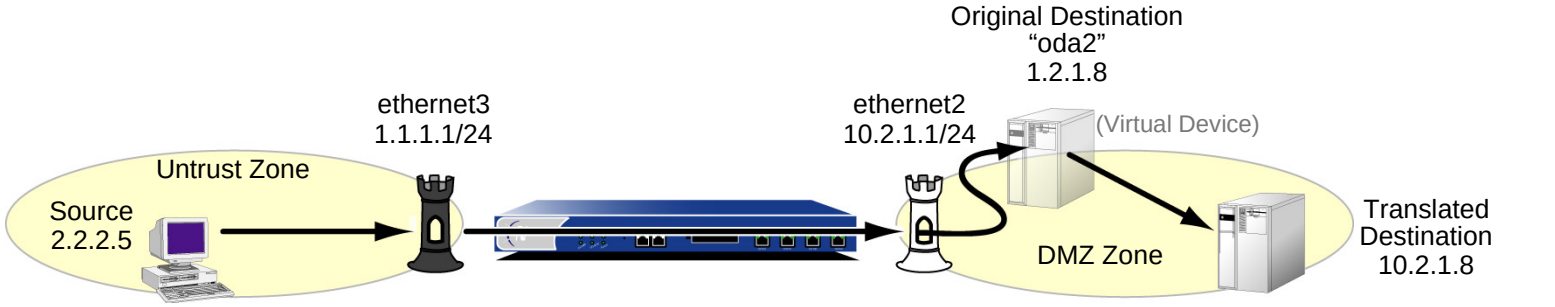

SRC IP	**DST IP**	SRC PT	DST PT	PROTO
2.2.2.5	**1.2.1.8**	25611	80	HTTP

SRC IP	**DST IP**	SRC PT	DST PT	PROTO
2.2.2.5	**10.2.1.8**	25611	80	HTTP

SRC IP	**DST IP**	SRC PT	DST PT	PROTO
2.2.2.5	**1.2.1.8**	40365	21	FTP

SRC IP	**DST IP**	SRC PT	DST PT	PROTO
2.2.2.5	**10.2.1.8**	40365	21	FTP

WebUI

1. Interfaces

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

Network > Interfaces > Edit (for ethernet2): Enter the following, and then click **OK**:

Zone Name: DMZ

Static IP: (select this option when present)

IP Address/Netmask: 10.2.1.1/24

2. Address

Objects > Addresses > List > New: Enter the following information, and then click **OK**:

Address Name: oda2

IP Address/Domain Name:

IP/Netmask: (select), 1.2.1.8/32

Zone: DMZ

3. Service Group

Objects > Services > Group: Enter the following group name, move the following services, and then click **OK**:

Group Name: HTTP-FTP

Select **HTTP** and use the **<<** button to move the service from the Available Members column to the Group Members column.

Select **FTP** and use the **<<** button to move the service from the Available Members column to the Group Members column.

4. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 1.2.1.8/32

Gateway: (select)

Interface: ethernet2

Gateway IP Address: 0.0.0.0

5. Policy

Policies > (From: Untrust, To: DMZ) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), oda2

Service: HTTP-FTP

Action: Permit

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

NAT:

Destination Translation: (select)

Translate to IP: (select), 10.2.1.8

Map to Port: (clear)

CLI

1. Interfaces

```
set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24

set interface ethernet2 zone dmz
set interface ethernet2 ip 10.2.1.1/24
```

2. Address

```
set address dmz oda2 1.2.1.8/32
```

3. Service Group

```
set group service http-ftp
set group service http-ftp add http
set group service http-ftp add ftp
```

4. Route

```
set vrouter trust-vr route 1.2.1.8/32 interface ethernet2
```

5. Policy

```
set policy from untrust to dmz any oda2 http-ftp nat dst ip 10.2.1.8 permit
save
```

Translating from One Address to Multiple Addresses

The NetScreen device can translate the same original destination address to different translated destination addresses specified in different policies, depending on the type of service or the source address specified in each policy. You might want the NetScreen device to redirect HTTP traffic from 1.2.1.8 to 10.2.1.8, and FTP traffic from 1.2.1.8 to 10.2.1.9 (see the following example). Perhaps you want the NetScreen device to redirect HTTP traffic sent from host1 to 1.2.1.8 over to 10.2.1.8, but HTTP traffic sent from host2 to 1.2.1.8 over to 10.2.1.37. In both cases, the NetScreen device redirects traffic sent to the same original destination address to different translated addresses.

Example: One-to-Many Destination Translation

In this example, you create two policies that use the same original destination address (1.2.1.8), but that direct traffic sent to that address to two different translated destination addresses based on the service type. These policies instruct the NetScreen device to perform the following tasks:

- Permit both FTP and HTTP traffic from any address in the Untrust zone to a user-defined address named “oda3” in the DMZ zone
- For HTTP traffic, translate the destination IP address in the IP packet header from 1.2.1.8 to 10.2.1.8
- For FTP traffic, translate the destination IP address from 1.2.1.8 to 10.2.1.9
- Leave the original destination port number in the TCP segment header as is (80 for HTTP, 21 for FTP)
- Forward HTTP traffic to 10.2.1.8 and FTP traffic to 10.2.1.9 in the DMZ zone

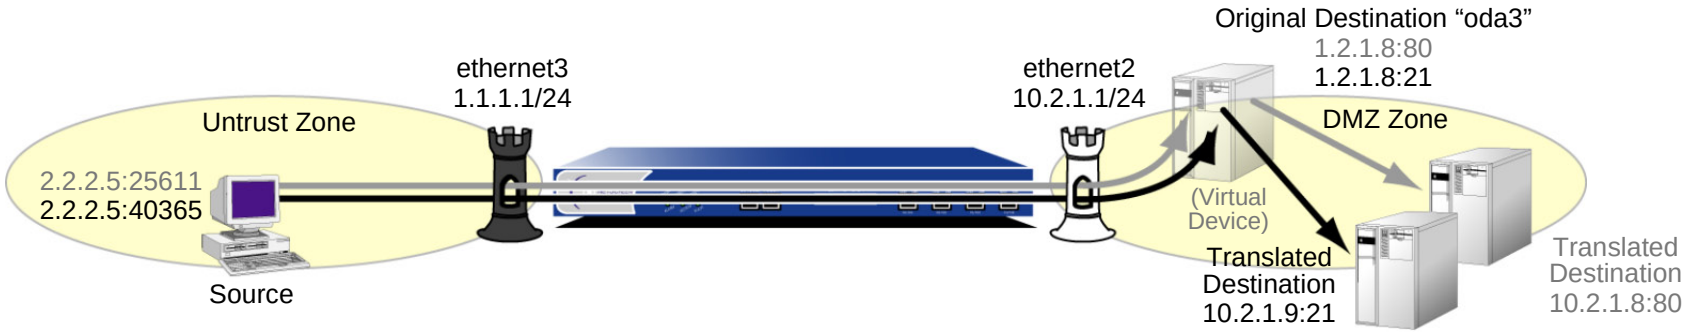

SRC IP	**DST IP**	SRC PT	DST PT	PROTO
2.2.2.5	**1.2.1.8**	25611	80	HTTP

SRC IP	**DST IP**	SRC PT	DST PT	PROTO
2.2.2.5	**10.2.1.8**	25611	80	HTTP

SRC IP	**DST IP**	SRC PT	DST PT	PROTO
2.2.2.5	**1.2.1.8**	40365	21	FTP

SRC IP	**DST IP**	SRC PT	DST PT	PROTO
2.2.2.5	**10.2.1.9**	40365	21	FTP

You bind ethernet3 to the Untrust zone, and assign it IP address 1.1.1.1/24. You bind ethernet2 to the DMZ, and assign it IP address 10.2.1.1/24. You also define a route to the original destination address 1.2.1.8 through ethernet2. Both the Untrust zone and the DMZ zone are in the trust-vr routing domain.

WebUI

1. Interfaces

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

Network > Interfaces > Edit (for ethernet2): Enter the following, and then click **OK**:

Zone Name: DMZ

Static IP: (select this option when present)

IP Address/Netmask: 10.2.1.1/24

2. Address

Objects > Addresses > List > New: Enter the following information, and then click **OK**:

Address Name: oda3

IP Address/Domain Name:

IP/Netmask: (select), 1.2.1.8/32

Zone: DMZ

3. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 1.2.1.8/32

Gateway: (select)

Interface: ethernet2

Gateway IP Address: 0.0.0.0

4. Policies

Policies > (From: Untrust, To: DMZ) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), oda3

Service: HTTP

Action: Permit

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

NAT:

Destination Translation: (select)

Translate to IP: (select), 10.2.1.8

Map to Port: (clear)

Policies > (From: Untrust, To: DMZ) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), oda3

Service: FTP

Action: Permit

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

NAT:

Destination Translation: (select)

Translate to IP: (select), 10.2.1.9

Map to Port: (clear)

CLI

1. Interfaces

```
set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24

set interface ethernet2 zone dmz
set interface ethernet2 ip 10.2.1.1/24
```

2. Address

```
set address dmz oda3 1.2.1.8/32
```

3. Route

```
set vrouter trust-vr route 1.2.1.8/32 interface ethernet2
```

4. Policies

```
set policy from untrust to dmz any oda3 http nat dst ip 10.2.1.8 permit
set policy from untrust to dmz any oda3 ftp nat dst ip 10.2.1.9 permit
save
```

NAT-Dst: Many-to-One Mapping

The relationship of the original destination address to the translated destination address can also be a many-to-one relationship. In this case, the NetScreen device forwards traffic sent to several original destination addresses to a single translated destination address. Optionally, you can also specify destination port mapping.

Example: Many-to-One Destination Translation

In this example, you create a policy that redirects traffic sent to different original destination addresses (1.2.1.10 and 1.2.1.20) to the same translated destination address. This policy instructs the NetScreen device to perform the following tasks:

- Permit HTTP traffic from any address in the Untrust zone to a user-defined address group named “oda45” with addresses “oda4” (1.2.1.10) and “oda5” (1.2.1.20) in the DMZ zone
- Translate the destination IP addresses in the IP packet header from 1.2.1.10 and 1.2.1.20 to 10.2.1.15
- Leave the original destination port number in the TCP segment header as is (80 for HTTP)
- Forward the HTTP traffic to 10.2.1.15 in the DMZ zone

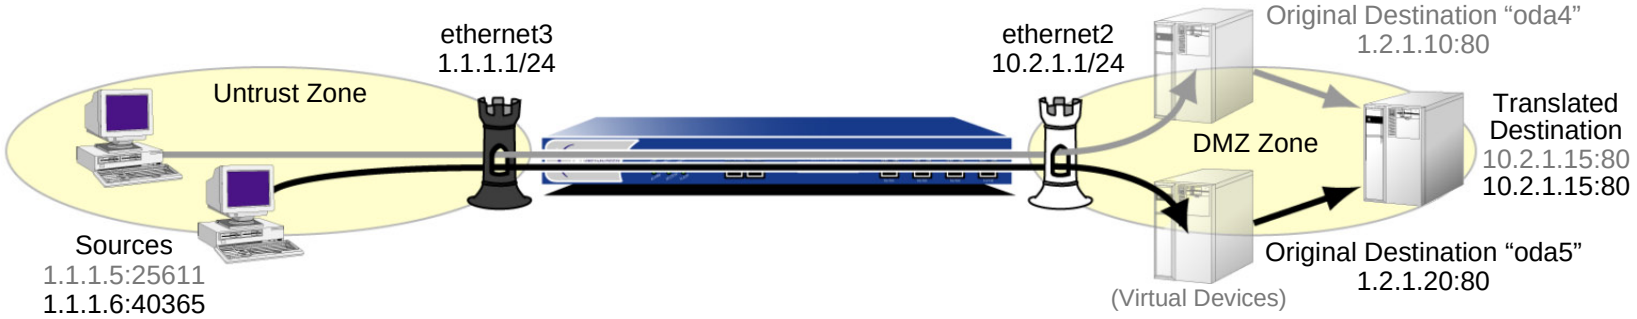

SRC IP	**DST IP**	SRC PT	DST PT	PROTO
1.1.1.5	**1.2.1.10**	25611	80	HTTP

SRC IP	**DST IP**	SRC PT	DST PT	PROTO
1.1.1.5	**10.2.1.15**	25611	80	HTTP

SRC IP	**DST IP**	SRC PT	DST PT	PROTO
1.1.1.6	**1.2.1.20**	40365	80	HTTP

SRC IP	**DST IP**	SRC PT	DST PT	PROTO
1.1.1.6	**10.2.1.15**	40365	80	HTTP

You bind ethernet3 to the Untrust zone, and assign it IP address 1.1.1.1/24. You bind ethernet2 to the DMZ, and assign it IP address 10.2.1.1/24. You also define a route to the original destination addresses 1.2.1.10 and 1.2.1.20 through ethernet2. Both the Untrust zone and the DMZ zone are in the trust-vr routing domain.

WebUI

1. Interfaces

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

Network > Interfaces > Edit (for ethernet2): Enter the following, and then click **OK**:

Zone Name: DMZ

Static IP: (select this option when present)

IP Address/Netmask: 10.2.1.1/24

2. Addresses

Objects > Addresses > List > New: Enter the following information, and then click **OK**:

Address Name: oda4

IP Address/Domain Name:

IP/Netmask: (select), 1.2.1.10/32

Zone: DMZ

Objects > Addresses > List > New: Enter the following information, and then click **OK**:

Address Name: oda5

IP Address/Domain Name:

IP/Netmask: (select), 1.2.1.20/32

Zone: DMZ

Objects > Addresses > Group > (for Zone: DMZ) New: Enter the following group name, move the following addresses, and then click **OK**:

Group Name: oda45

Select **oda4** and use the **<<** button to move the address from the Available Members column to the Group Members column.

Select **oda5** and use the **<<** button to move the address from the Available Members column to the Group Members column.

3. Routes

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 1.2.1.10/32

Gateway: (select)

Interface: ethernet2

Gateway IP Address: 0.0.0.0

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 1.2.1.20/32

Gateway: (select)

Interface: ethernet2

Gateway IP Address: 0.0.0.0

4. Policy

Policies > (From: Untrust, To: DMZ) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), oda45

Service: HTTP

Action: Permit

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

NAT:

Destination Translation: (select)

Translate to IP: (select), 10.2.1.15

Map to Port: (clear)

CLI

1. Interfaces

```
set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24

set interface ethernet2 zone dmz
set interface ethernet2 ip 10.2.1.1/24
```

2. Addresses

```
set address dmz oda4 1.2.1.10/32
set address dmz oda5 1.2.1.20/32
set group address dmz oda45 add oda4
set group address dmz oda45 add oda5
```

3. Routes

```
set vrouter trust-vr route 1.2.1.10/32 interface ethernet2
set vrouter trust-vr route 1.2.1.20/32 interface ethernet2
```

4. Policy

```
set policy from untrust to dmz any oda45 http nat dst ip 10.2.1.15 permit
save
```

NAT-Dst: Many-to-Many Mapping

You can use destination network address translation (NAT-dst) to translate one range of IP addresses to another range. The range of addresses can be a subnet or a smaller set of addresses within a subnet. NetScreen employs an address shifting mechanism to maintain the relationships among the original range of destination addresses after translating them to the new range of addresses. For example, if the range of original addresses is 10.1.1.1 – 10.1.1.50 and the starting address for the translated address range is 10.100.3.101, then the NetScreen device translates the addresses as follows:

- 10.1.1.1 – 10.100.3.101
- 10.1.1.2 – 10.100.3.102
- 10.1.1.3 – 10.100.3.103

  …
- 10.1.1.48 – 10.100.3.148
- 10.1.1.49 – 10.100.3.149
- 10.1.1.50 – 10.100.3.150

If, for example, you want to create a policy that applies the above translations to HTTP traffic from any address in zoneA to an address group named “addr1-50”, which contains all the addresses from 10.1.1.1 to 10.1.1.50, in zoneB, you can enter the following CLI command:

```
set policy id 1 from zoneA to zoneB any addr1-50 http nat dst ip 10.100.3.101
    10.100.3.150 permit
```

If any host in zoneA initiates HTTP traffic to an address within the defined range in zoneB, such as 10.1.1.37, then the NetScreen device applies this policy and translates the destination address to 10.100.3.137.

The NetScreen device only performs NAT-dst if the source and destination zones, the source and destination addresses, and the service specified in the policy all match these components in the packet. For example, you might create another policy that permits traffic from any host in zoneA to any host in zoneB and position it after policy 1 in the policy list:

```
set policy id 1 from zoneA to zoneB any addr1-50 http nat dst ip 10.100.3.101
    10.100.3.150 permit
set policy id 2 from zoneA to zoneB any any any permit
```

If you have these two policies configured, the following kinds of traffic sent from a host in zoneA to a host in zoneB bypass the NAT-dst mechanism:

- A zoneA host initiates non-HTTP traffic to 10.1.1.37 in zone B. The NetScreen device applies policy 2 because the service is not HTTP and passes the traffic without translating the destination address.
- A zoneA host initiates HTTP traffic to 10.1.1.51 in zone B. The NetScreen device also applies policy 2 because the destination address is not in the addr1-50 address group, and passes the traffic without translating the destination address.

Example: Many-to-Many Destination Translation

In this example, you configure a policy that applies NAT-dst when any kind of traffic is sent to any host in a subnet, instructing the NetScreen device to perform the following tasks:

- Permit all traffic types from any address in the Untrust zone to any address in the DMZ zone
- Translate the original destination address named “oda6” from the 1.2.1.0/24 subnet to a corresponding address in the 10.2.1.0/24 subnet
- Leave the original destination port number in the TCP segment header as is
- Forward HTTP traffic to the translated address in the DMZ zone

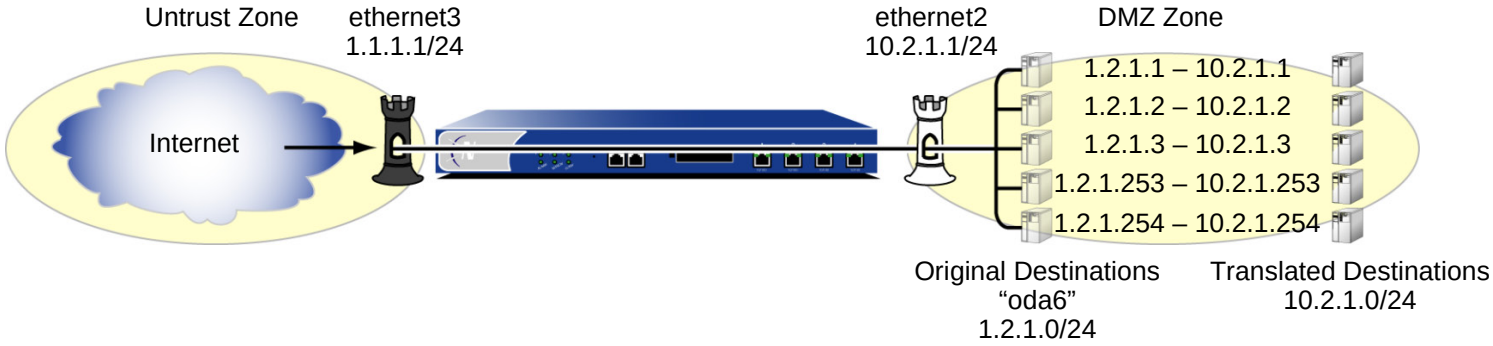

You bind ethernet3 to the Untrust zone, and assign it IP address 1.1.1.1/24. You bind ethernet2 to the DMZ, and assign it IP address 10.2.1.1/24. You also define a route to the original destination address subnet (1.2.1.0/24) through ethernet2. Both the Untrust zone and the DMZ zone are in the trust-vr routing domain.

WebUI

1. Interfaces

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

Network > Interfaces > Edit (for ethernet2): Enter the following, and then click **OK**:

Zone Name: DMZ

Static IP: (select this option when present)

IP Address/Netmask: 10.2.1.1/24

2. Address

Objects > Addresses > List > New: Enter the following information, and then click **OK**:

Address Name: oda6

IP Address/Domain Name:

IP/Netmask: (select), 1.2.1.0/24

Zone: DMZ

3. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 1.2.1.0/24

Gateway: (select)

Interface: ethernet2

Gateway IP Address: 0.0.0.0

4. Policy

Policies > (From: Untrust, To: DMZ) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), oda6

Service: Any

Action: Permit

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

NAT:

Destination Translation: (select)

Translate to IP Range: (select), 10.2.1.0 – 10.2.1.254

CLI

1. Interfaces

```
set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24

set interface ethernet2 zone dmz
set interface ethernet2 ip 10.2.1.1/24
```

2. Address

```
set address dmz oda6 1.2.1.0/24
```

3. Route

```
set vrouter trust-vr route 1.2.1.0/24 interface ethernet2
```

4. Policy

```
set policy from untrust to dmz any oda6 any nat dst ip 10.2.1.1 10.2.1.254
    permit
save
```

NAT-Dst with Port Mapping

When you configure the NetScreen device to perform destination network address translation (NAT-dst), you can optionally enable port mapping. One reason to enable port mapping is to support multiple server processes for a single service on a single host. For example, one host can run two Web servers—one at port 80 and another at port 8081. For HTTP service 1, the NetScreen device performs NAT-dst without port mapping (dst port 80 -> 80). For HTTP service 2, the NetScreen device performs NAT-dst to the same destination IP address with port mapping (dst port 80 -> 8081). The host can sort HTTP traffic to the two Web servers by the two distinct destination port numbers.

Note: *NetScreen does not support port mapping for NAT-dst with address shifting. See* *"NAT-Dst: Many-to-Many Mapping" on page 300.*

Example: NAT-Dst with Port Mapping

In this example, you create two policies that perform NAT-dst and port mapping on Telnet traffic from the Trust and Untrust zones to a Telnet server in the DMZ zone. These policies instruct the NetScreen device to perform the following tasks:

- Permit Telnet from any address in the Untrust and Trust zones to 1.2.1.15 in the DMZ zone
- Translate the original destination IP address named "oda7" from 1.2.1.15 to 10.2.1.15
- Translate the original destination port number in the TCP segment header from 23 to 2200
- Forward Telnet traffic to the translated address in the DMZ zone

You configure the following interface-to-zone bindings and address assignments:

- ethernet1: Trust zone, 10.1.1.1/24
- ethernet2: DMZ zone, 10.2.1.1/24.
- ethernet3: Untrust zone, 1.1.1.1/24.

You define an address entry "oda7" with IP address 1.2.1.15/32 in the DMZ zone. You also define a route to the original destination address 1.2.1.15 through ethernet2. The Trust, Untrust, and DMZ zones are all in the trust-vr routing domain.

WebUI

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet2): Enter the following, and then click **OK**:

Zone Name: DMZ

Static IP: (select this option when present)

IP Address/Netmask: 10.2.1.1/24

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

2. Address

Objects > Addresses > List > New: Enter the following information, and then click **OK**:

Address Name: oda7

IP Address/Domain Name:

IP/Netmask: (select), 1.2.1.15/32

Zone: DMZ

3. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 1.2.1.15/32

Gateway: (select)

Interface: ethernet2

Gateway IP Address: 0.0.0.0

4. Policies

Policies > (From: Trust, To: DMZ) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), oda7

Service: Telnet

Action: Permit

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

NAT:

Destination Translation: (select)

Translate to IP: (select), 10.2.1.15

Map to Port: (select), 2200

Policies > (From: Untrust, To: DMZ) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), oda7

Service: Telnet

Action: Permit

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

NAT:

Destination Translation: (select)

Translate to IP: (select), 10.2.1.15

Map to Port: (select), 2200

CLI

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet2 zone dmz
set interface ethernet2 ip 10.2.1.1/24

set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
```

2. Address

```
set address dmz oda7 1.2.1.15/32
```

3. Route

```
set vrouter trust-vr route 1.2.1.15/32 interface ethernet2
```

4. Policies

```
set policy from trust to dmz any oda7 telnet nat dst ip 10.2.1.15 port 2200
    permit
set policy from untrust to dmz any oda7 telnet nat dst ip 10.2.1.15 port 2200
    permit
save
```

NAT-Src and NAT-Dst in the Same Policy

You can combine source and destination network address translation (NAT-src and NAT-dst) in the same policy. This combination provides you with a means to change both the source and destination IP address at a single point in the data path.

Example: NAT-Src and NAT-Dst Combined

In this example, you configure a NetScreen device (NetScreen-1) that is between a service provider's customers and server farms. The customers connect to NetScreen-1 through ethernet1, which has IP address 10.1.1.1/24 and is bound to the Trust zone. NetScreen-1 then forwards their traffic through one of two route-based VPN tunnels to reach the servers they want to target[11]. The tunnel interfaces that are bound to these tunnels are in the Untrust zone. Both the Trust and Untrust zones are in the trust-vr routing domain.

Because the customers might have the same addresses as those of the servers to which they want to connect, NetScreen-1 must perform both source and destination address translation (NAT-src and NAT-dst). To retain addressing independence and flexibility, the NetScreen devices protecting the server farms—NetScreen-A and NetScreen-B—perform NAT-dst. The service provider instructs the customers and the server farm admins to reserve addresses 10.173.10.1–10.173.10.7, 10.173.20.0/24, 10.173.30.0/24, 10.173.40.0/24, and 10.173.50.0/24 for this purpose. These addresses are used as follows:

- The two tunnel interfaces have the following address assignments:
  - tunnel.1, 10.173.10.1/30
  - tunnel.2, 10.173.10.5/30
- Each tunnel interface supports the following DIP pools with PAT enabled:
  - tunnel.1, DIP ID 5: 10.173.10.2–10.173.10.2
  - tunnel.2, DIP ID 6: 10.173.10.6–10.173.10.6
- When NetScreen-1 performs NAT-dst, it translates original destination addresses with address shifting as follows[12]:
  - 10.173.20.0/24 to 10.173.30.0/24
  - 10.173.40.0/24 to 10.173.50.0/24

11. Policy-based VPNs do not support NAT-dst. You must use a route-based VPN configuration with NAT-dst.

12. For information about address shifting when performing NAT-dst, see "NAT-Dst: Many-to-Many Mapping" on page 300.

The configurations for both tunnels—vpn1 and vpn2—use the following parameters: AutoKey IKE, preshared key ("netscreen1" for vpn1, and "netscreen2" for vpn2), and the security level predefined as "Compatible" for both Phase 1 and Phase 2 proposals. (For details about these proposals, see "Tunnel Negotiation" on page **5**-11.) The proxy ID for both vpn1 and vpn2 is 0.0.0.0/0 - 0.0.0.0/0 - any.

Note: *The configuration for NetScreen-1 is provided first. The VPN configurations for NetScreen-A and NetScreen-B follow, and are included for completeness.*

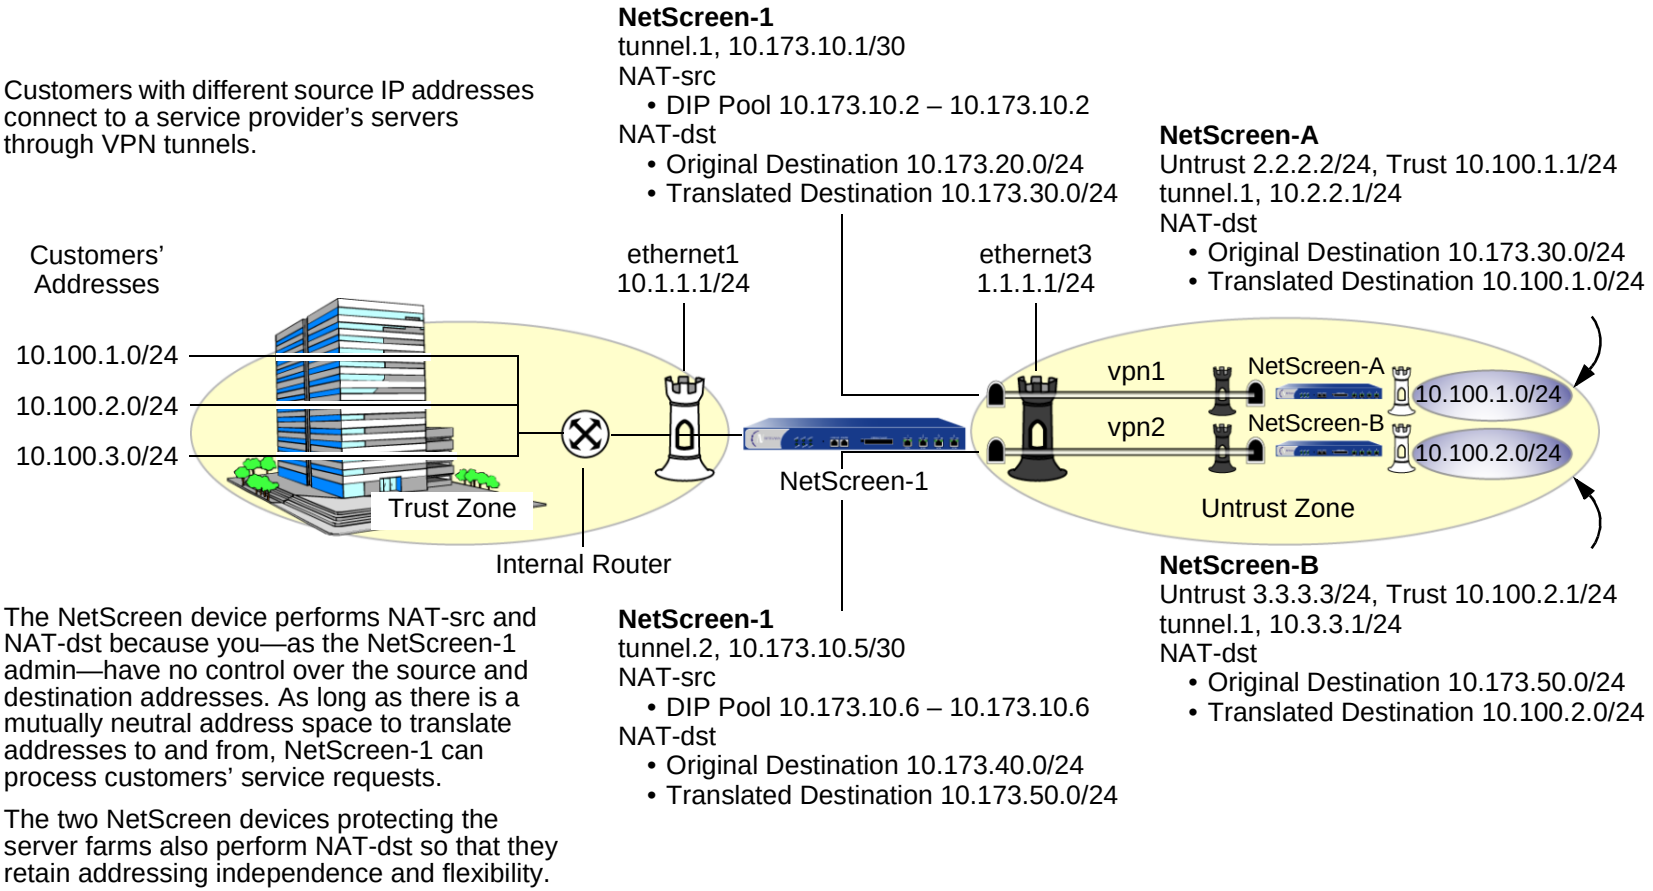

WebUI (NetScreen-1)

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

Network > Interfaces > New Tunnel IF: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.1

Zone (VR): Untrust (trust-vr)

Fixed IP: (select)

IP Address / Netmask: 10.173.10.1/30

Network > Interfaces > New Tunnel IF: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.1

Zone (VR): Untrust (trust-vr)

Fixed IP: (select)

IP Address / Netmask: 10.173.10.5/30

2. DIP Pools

Network > Interfaces > Edit (for tunnel.1) > DIP > New: Enter the following, and then click **OK**:

ID: 5

IP Address Range: (select), 10.173.10.2 ~ 10.173.10.2

Port Translation: (select)

In the same subnet as the interface IP or its secondary IPs: (select)

Network > Interfaces > Edit (for tunnel.2) > DIP > New: Enter the following, and then click **OK**:

ID: 6

IP Address Range: (select), 10.173.10.6 ~ 10.173.10.6

Port Translation: (select)

In the same subnet as the interface IP or its secondary IPs: (select)

3. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: serverfarm-A

IP Address/Domain Name:

IP/Netmask: (select), 10.173.20.0/24

Zone: Untrust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: serverfarm-B

IP Address/Domain Name:

IP/Netmask: (select), 10.173.40.0/24

Zone: Untrust

4. VPNs

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: vpn1

Security Level: Compatible

Remote Gateway: Create a Simple Gateway: (select)

Gateway Name: gw-A

Type: Static IP: (select), Address/Hostname: 2.2.2.2

Preshared Key: netscreen1

Security Level: Compatible

Outgoing Interface: ethernet3[13]

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Bind to Tunnel Interface: (select), tunnel.1

Proxy-ID: (select)

Local IP / Netmask: 0.0.0.0/0

Remote IP / Netmask: 0.0.0.0/0

Service: ANY

13. The outgoing interface does not have to be in the same zone to which the tunnel interface is bound, although in this case they are in the same zone.

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: vpn2

Security Level: Compatible

Remote Gateway: Create a Simple Gateway: (select)

Gateway Name: gw-B

Type: Static IP: (select), Address/Hostname: 3.3.3.3

Preshared Key: netscreen2

Security Level: Compatible

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Bind to Tunnel Interface: (select), tunnel.2

Proxy-ID: (select)

Local IP / Netmask: 0.0.0.0/0

Remote IP / Netmask: 0.0.0.0/0

Service: ANY

5. Routes

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 1.1.1.250

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 10.173.20.0/24

Gateway: (select)

Interface: tunnel.1

Gateway IP Address: 0.0.0.0

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 10.173.30.0/24

Gateway: (select)

Interface: tunnel.1

Gateway IP Address: 0.0.0.0

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 10.173.40.0/24

Gateway: (select)

Interface: tunnel.2

Gateway IP Address: 0.0.0.0

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 10.173.50.0/24

Gateway: (select)

Interface: tunnel.2

Gateway IP Address: 0.0.0.0

6. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), serverfarm-A

Service: ANY

Action: Permit

Position at Top: (select)

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Policy configuration page:

NAT:

Source Translation: (select)

(DIP on): 5 (10.173.10.2–10.173.10.2)/X-late

Destination Translation: (select)

Translate to IP Range: (select), 10.173.30.0 – 10.173.30.255

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), serverfarm-B

Service: ANY

Action: Permit

Position at Top: (select)

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Policy configuration page:

NAT:

Source Translation: (select)

(DIP on): 6 (10.173.10.6–10.173.10.6)/X-late

Destination Translation: (select)

Translate to IP Range: (select), 10.173.50.0 – 10.173.50.255

CLI (NetScreen-1)

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24

set interface tunnel.1 zone untrust
set interface tunnel.1 ip 10.173.10.1/30

set interface tunnel.2 zone untrust
set interface tunnel.2 ip 10.173.10.5/30
```

2. DIP Pools

```
set interface tunnel.1 dip-id 5 10.173.10.2 10.173.10.2
set interface tunnel.2 dip-id 6 10.173.10.6 10.173.10.6
```

3. Addresses

```
set address untrust serverfarm-A 10.173.20.0/24
set address untrust serverfarm-B 10.173.40.0/24
```

4. VPNs

```
set ike gateway gw-A ip 2.2.2.2 main outgoing-interface ethernet3 preshare
    netscreen1 sec-level compatible
set vpn vpn1 gateway gw-A sec-level compatible
set vpn vpn1 bind interface tunnel.1
set vpn vpn1 proxy-id local-ip 0.0.0.0/0 remote-ip 0.0.0.0/0 any
```

```
set ike gateway gw-B ip 3.3.3.3 main outgoing-interface ethernet3 preshare
    netscreen2 sec-level compatible
set vpn vpn2 gateway gw-B sec-level compatible
set vpn vpn2 bind interface tunnel.2
set vpn vpn2 proxy-id local-ip 0.0.0.0/0 remote-ip 0.0.0.0/0 any
```

5. Routes

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
set vrouter trust-vr route 10.173.20.0/24 interface tunnel.1
set vrouter trust-vr route 10.173.30.0/24 interface tunnel.1
set vrouter trust-vr route 10.173.40.0/24 interface tunnel.2
set vrouter trust-vr route 10.173.50.0/24 interface tunnel.2
```

6. Policies

```
set policy top from trust to untrust any serverfarm-A any nat src dip-id 5 dst
    ip 10.173.30.0 10.173.30.255 permit
set policy top from trust to untrust any serverfarm-B any nat src dip-id 6 dst
    ip 10.173.50.0 10.173.50.255 permit
save
```

WebUI (NetScreen-A)

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.100.1.1/24

Select the following, and then click **OK**:
Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 2.2.2.2/24

Network > Interfaces > New Tunnel IF: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.1

Zone (VR): Untrust (trust-vr)

Fixed IP: (select)

IP Address / Netmask: 10.2.2.1/24

2. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: serverfarm-A

IP Address/Domain Name:

IP/Netmask: (select), 10.173.30.0/24

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: customer1

IP Address/Domain Name:

IP/Netmask: (select), 10.173.10.2/32

Zone: Untrust

3. VPN

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: vpn1

Security Level: Compatible

Remote Gateway: Create a Simple Gateway: (select)

Gateway Name: gw-1

Type: Static IP: (select), Address/Hostname: 1.1.1.1

Preshared Key: netscreen1

Security Level: Compatible

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Bind to Tunnel Interface: (select), tunnel.1

Proxy-ID: (select)

Local IP / Netmask: 0.0.0.0/0

Remote IP / Netmask: 0.0.0.0/0

Service: ANY

4. Routes

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 2.2.2.250

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 10.173.10.2/32

Gateway: (select)

Interface: tunnel.1

Gateway IP Address: 0.0.0.0

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 10.173.30.0/24

Gateway: (select)

Interface: ethernet1

Gateway IP Address: 0.0.0.0

5. Policy

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), customer1

Destination Address:

Address Book Entry: (select), serverfarm-A

Service: ANY

Action: Permit

Position at Top: (select)

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Policy configuration page:

NAT:

Destination Translation: (select)

Translate to IP Range: (select), 10.100.1.0 – 10.100.1.255

CLI (NetScreen-A)

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.100.1.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 2.2.2.2/24

set interface tunnel.1 zone untrust
set interface tunnel.1 ip 10.2.2.1/24
```

2. Addresses

```
set address trust serverfarm-A 10.173.30.0/24
set address untrust customer1 10.173.10.2/32
```

3. VPN

```
set ike gateway gw-1 ip 1.1.1.1 main outgoing-interface ethernet3 preshare
    netscreen1 sec-level compatible
set vpn vpn1 gateway gw-1 sec-level compatible
set vpn vpn1 bind interface tunnel.1
set vpn vpn1 proxy-id local-ip 0.0.0.0/0 remote-ip 0.0.0.0/0 any
```

4. Routes

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 2.2.2.250
set vrouter trust-vr route 10.173.10.2/32 interface tunnel.1
set vrouter trust-vr route 10.173.30.0/24 interface ethernet1
```

5. Policy

```
set policy top from untrust to trust customer1 serverfarm-A any nat dst ip
    10.100.1.0 10.100.1.255 permit
save
```

WebUI (NetScreen-B)

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.100.2.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 3.3.3.3/24

Network > Interfaces > New Tunnel IF: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.1

Zone (VR): Untrust (trust-vr)

Fixed IP: (select)

IP Address / Netmask: 10.3.3.1/24

2. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: serverfarm-B

IP Address/Domain Name:

IP/Netmask: (select), 10.173.50.0/24

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: customer1

IP Address/Domain Name:

IP/Netmask: (select), 10.173.10.6/32

Zone: Untrust

3. VPN

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: vpn1

Security Level: Compatible

Remote Gateway: Create a Simple Gateway: (select)

Gateway Name: gw-1

Type: Static IP: (select), Address/Hostname: 1.1.1.1

Preshared Key: netscreen2

Security Level: Compatible

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Bind to Tunnel Interface: (select), tunnel.1

Proxy-ID: (select)

Local IP / Netmask: 0.0.0.0/0

Remote IP / Netmask: 0.0.0.0/0

Service: ANY

4. Routes

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 3.3.3.250

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 10.173.10.6/32

Gateway: (select)

Interface: tunnel.1

Gateway IP Address: 0.0.0.0

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 10.173.50.0/24

Gateway: (select)

Interface: ethernet1

Gateway IP Address: 0.0.0.0

5. Policy

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), customer1

Destination Address:

Address Book Entry: (select), serverfarm-B

Service: ANY

Action: Permit

Position at Top: (select)

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Policy configuration page:

NAT:

Destination Translation: (select)

Translate to IP Range: (select), 10.100.2.0 – 10.100.2.255

CLI (NetScreen-B)

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.100.2.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 3.3.3.3/24

set interface tunnel.1 zone untrust
set interface tunnel.1 ip 10.3.3.1/24
```

2. Addresses

```
set address trust serverfarm-B 10.173.50.0/24
set address untrust customer1 10.173.10.6/32
```

3. VPN

```
set ike gateway gw-1 ip 1.1.1.1 main outgoing-interface ethernet3 preshare
    netscreen2 sec-level compatible
set vpn vpn2 gateway gw-1 sec-level compatible
set vpn vpn2 bind interface tunnel.1
set vpn vpn2 proxy-id local-ip 0.0.0.0/0 remote-ip 0.0.0.0/0 any
```

4. Routes

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 3.3.3.250
set vrouter trust-vr route 10.173.10.6/32 interface tunnel.1
set vrouter trust-vr route 10.173.50.0/24 interface ethernet1
```

5. Policy

```
set policy top from untrust to trust customer1 serverfarm-B any nat dst ip
    10.100.2.0 10.100.2.255 permit
save
```

MAPPED IP ADDRESSES

Mapped IP (MIP) is a direct one-to-one mapping of one IP address to another. The NetScreen device forwards incoming traffic destined for a MIP to the host with the address to which the MIP points. Essentially, a MIP is static destination address translation, mapping the destination IP address in an IP packet header to another static IP address. When a MIP host initiates outbound traffic, the NetScreen device translates the source IP address of the host to that of the MIP address. This bidirectional translation symmetry differs from the behavior of source and destination address translation (see "Directional Nature of NAT-Src and NAT-Dst" on page 257).

MIPs allow inbound traffic to reach private addresses in a zone whose interface is in NAT mode. MIPs also provide part of the solution to the problem of overlapping address spaces[14] at two sites connected by a VPN tunnel. (For the complete solution to this problem, see "VPN Sites with Overlapping Addresses" on page **5**-168.)

You can create a MIP in the same subnet as a tunnel interface with an IP address/netmask, or in the same subnet as the IP address/netmask of an interface bound to a Layer 3 (L3) security zone[15]. Although you configure MIPs for interfaces bound to tunnel zones and security zones, the MIP that you define is stored in the Global zone.

Mapped IP: Inbound traffic from the Untrust zone is mapped from 210.1.1.5 to 10.1.1.5 in the Trust zone.

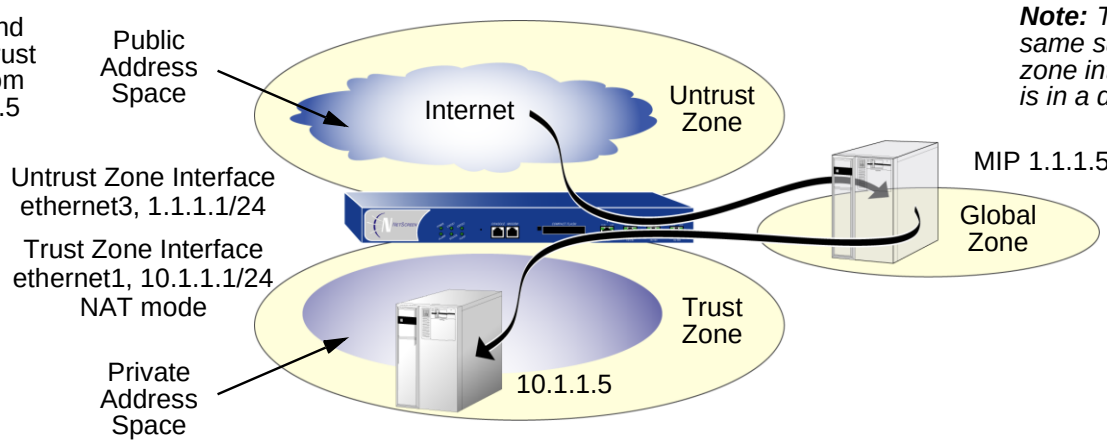

Note: *The MIP (1.1.1.5) is in the same subnet as the Untrust zone interface (1.1.1.1/24) but it is in a different zone.*

14. An overlapping address space is when the IP address range in two networks are partially or completely the same.

15. An exception is a MIP defined for an interface in the Untrust zone. That MIP can be in a different subnet from an Untrust zone interface IP address. However, if that is the case, you must add a route on the external router pointing to an Untrust zone interface so that incoming traffic can reach the MIP. Also, you must define a static route on the NetScreen device associating the MIP with the interface that hosts it.

Note: *On some NetScreen devices, a MIP can use the same address as an interface, but a MIP address cannot be in a DIP pool.*

You can map an address-to-address or subnet-to-subnet relationship. When a subnet-to-subnet mapped IP configuration is defined, the netmask is applied to both the mapped IP subnet and the original IP subnet.

MIP and the Global Zone

Setting a MIP for an interface in any zone generates an entry for the MIP in the Global zone address book. The Global zone address book stores all MIPs, regardless of the zone to which their interfaces belong. You can use these MIP addresses as the destination addresses in policies between any two zones, and as the destination addresses when defining a Global policy. (For information about Global policies, see "Global Policies" on page 201.) Although the NetScreen device stores MIPs in the Global zone, you can use either the Global zone or the zone with the address to which the MIP points as the destination zone in a policy referencing a MIP.

Example: Adding a MIP to an Untrust Zone Interface

In this example, you bind ethernet1 to the Trust zone and assign it IP address 10.1.1.1/24. You bind ethernet2 to the Untrust zone and assign it IP address 1.1.1.1/24. Then you configure a MIP to direct incoming HTTP traffic destined for 1.1.1.5 in the Untrust zone to a Web server at 10.1.1.5 in the Trust zone. Finally, you create a policy permitting HTTP traffic from the any address in the Untrust zone to the MIP—and consequently to the host with the address to which the MIP points—in the Trust zone. All security zones are in the trust-vr routing domain.

Note: *No address book entry is required for a Mapped IP or the host to which it points.*

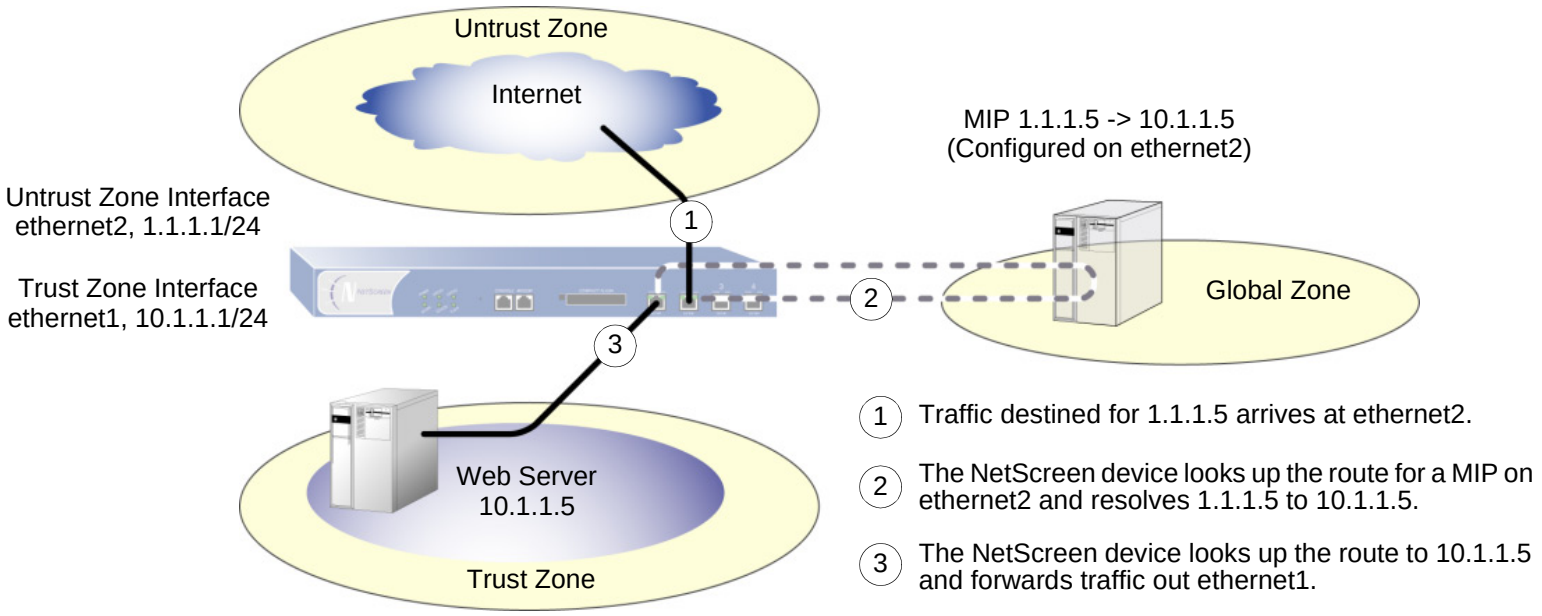

WebUI

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet2): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

2. MIP

Network > Interfaces > Edit (for ethernet2) > MIP > New: Enter the following, and then click **OK**:

Mapped IP: 1.1.1.5

Netmask: 255.255.255.255

Host IP Address: 10.1.1.5

Host Virtual Router Name: trust-vr

3. Policy

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), MIP(1.1.1.5)

Service: HTTP

Action: Permit

CLI

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet2 zone untrust
set interface ethernet2 ip 1.1.1.1/24
```

2. MIP

```
set interface ethernet2 mip 1.1.1.5 host 10.1.1.5 netmask 255.255.255.255[16]
    vrouter trust-vr[17]
```

3. Policy

```
set policy from untrust to trust any mip(1.1.1.5) http permit
save
```

16. By default, the netmask for a MIP is 32 bits (255.255.255.255), mapping the address to a single host. You can also define a MIP for a range of addresses. For example, to define 1.1.1.5 as a MIP for the addresses 10.1.10.129–10.1.10.254 within a class C subnet through the CLI, use the following syntax: **set interface** *interface* **mip 1.1.1.5 host 10.1.10.128 netmask 255.255.255.128**. Be careful not to use a range of addresses that includes the interface or router addresses.

17. The default virtual router is the trust-vr. You do not have to specify that the virtual router is the trust-vr or that the MIP has a 32-bit netmask . These arguments are included in this command to provide symmetry with the WebUI configuration.

Example: Reaching a MIP from Different Zones

Traffic from different zones can still reach a MIP through other interfaces than the one on which you configured the MIP. To accomplish this, you must set a route on the routers in each of the other zones that points inbound traffic to the IP address of their respective interfaces to reach the MIP[18].

In this example, you configure a MIP (1.1.1.5) on the interface in the Untrust zone (ethernet2, 1.1.1.1/24) to map to a Web server in the Trust zone (10.1.1.5). The interface bound to the Trust zone is ethernet1 with IP address 10.1.1.1/24.

You create a security zone named X-Net, bind ethernet3 to it, and assign the interface the IP address 1.3.3.1/24. You define an address for 1.1.1.5 for use in a policy to allow HTTP traffic from any address in the X-Net zone to the MIP in the Untrust zone. You also configure a policy to allow the HTTP traffic to pass from the Untrust zone to the Trust zone. All security zones are in the trust-vr routing domain.

Note: *You must enter a route on the router in the X-Net zone directing traffic destined for 1.1.1.5 (MIP) to 1.3.3.1 (IP address of ethernet3).*

18. If the MIP is in the same subnet as the interface on which you configured it, you do not have to add a route to the NetScreen device for traffic to reach the MIP via a different interface. However, if the MIP is in a different subnet than the IP address of its interface (which is possible only for a MIP on an interface in the Untrust zone), you must add a static route to the NetScreen routing table. Use the **set vrouter** *name_str* **route** *ip_addr* **interface** *interface* command (or its equivalent in the WebUI) , where *name_str* is the virtual router to which the specified interface belongs, and *interface* is interface on which you configured the MIP.

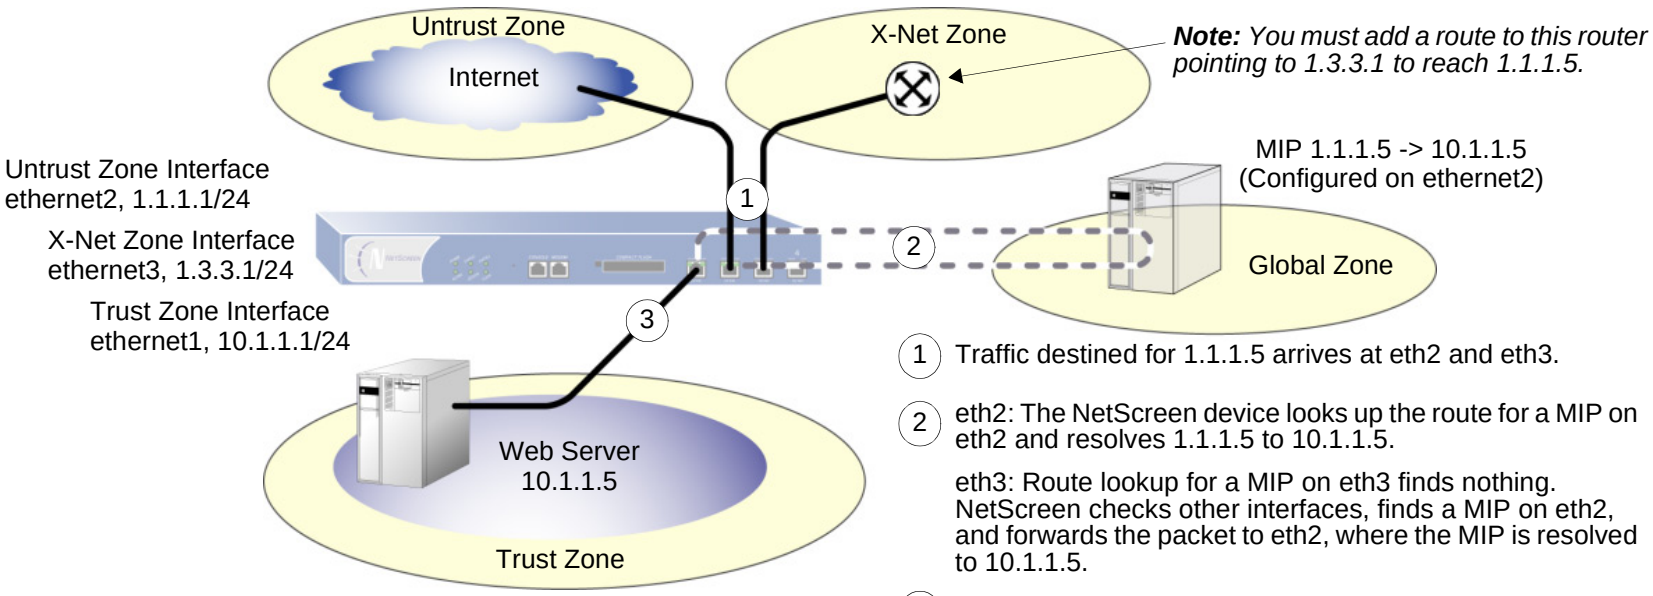

WebUI

1. Interfaces and Zones

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **OK**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet2): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

Network > Zones > New: Enter the following, and then click **OK**:

Zone Name: X-Net

Virtual Router Name: untrust-vr

Zone Type: Layer 3

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: X-Net

IP Address/Netmask: 1.3.3.1/24

2. Address

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: 1.1.1.5

IP Address/Domain Name:

IP/Netmask: (select), 1.1.1.5/32

Zone: Untrust

3. MIP

Network > Interfaces > Edit (for ethernet2) > MIP > New: Enter the following, and then click **OK**:

Mapped IP: 1.1.1.5

Netmask: 255.255.255.255

Host IP Address: 10.1.1.5

Host Virtual Router Name: trust-vr

4. Policies

Policies > (From: X-Net, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), 1.1.1.5

Service: HTTP

Action: Permit

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), MIP(1.1.1.5)

Service: HTTP

Action: Permit

CLI

1. Interfaces and Zones

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet2 zone untrust
set interface ethernet2 ip 1.1.1.1/24

set zone name X-Net
set interface ethernet3 zone X-Net
set interface ethernet3 ip 1.3.3.1/24
```

2. Address

```
set address untrust “1.1.1.5” 1.1.1.5/32
```

3. MIP

```
set interface ethernet2 mip 1.1.1.5 host 10.1.1.5 netmask 255.255.255.255
    vrouter trust-vr[19]
```

4. Policies

```
set policy from X-Net to untrust any “1.1.1.5” http permit
set policy from untrust to trust any mip(1.1.1.5) http permit
save
```

19. By default, the netmask for a MIP is 32 bits (255.255.255.255) and the default virtual router is the trust-vr. You do not have to specify them in the command. These arguments are included here to provide symmetry with the WebUI configuration.

Example: Adding a MIP to a Tunnel Interface

In this example, the IP address space for the network in the Trust zone is 10.1.1.0/24 and the IP address for the tunnel interface “tunnel.8” is 10.20.3.1. The physical IP address for a server on the network in the Trust zone is 10.1.1.25. To allow a remote site whose network in the Trust zone uses an overlapping address space to access the local server through a VPN tunnel, you create a MIP in the same subnet as the tunnel.8 interface. The MIP address is 10.20.3.25/32. (For a more complete example of a MIP with a tunnel interface, see “VPN Sites with Overlapping Addresses” on page **5**-168.)

WebUI

Network > Interfaces > Edit (for tunnel.8) > MIP > New: Enter the following, and then click **OK**:

Mapped IP: 10.20.3.25

Netmask: 255.255.255.255

Host IP Address: 10.1.1.25

Host Virtual Router Name: trust-vr

CLI

```
set interface tunnel.8 mip 10.20.3.25 host 10.1.1.25 netmask 255.255.255.255
    vrouter trust-vr[20]
save
```

Note: *When the remote administrator adds the address for the server to his Untrust zone address book, he must enter the MIP (10.20.3.25), not the physical IP address (10.1.1.25) of the server.*

The remote administrator also needs to apply policy-based NAT-src (using DIP) on the outgoing packets bound for the server through the VPN so that the local administrator can add an Untrust zone address that does not conflict with the local Trust zone addresses. Otherwise, the source address in the incoming policy would seem to be in the Trust zone.

20. By default, the netmask for a MIP is 32 bits (255.255.255.255) and the default virtual router is the trust-vr. You do not have to specify them in the command. These arguments are included here to provide symmetry with the WebUI configuration.

MIP-Same-as-Untrust

As IPv4 addresses become increasingly scarce, ISPs are becoming increasingly reluctant to give their customers more than one or two IP addresses. If you only have one IP address for the interface bound to the Untrust zone—the interface bound to the Trust zone is in Network Address Translation (NAT) mode—you can use the Untrust zone interface IP address as a mapped IP (MIP) to provide inbound access to an internal server or host, or to a VPN or L2TP tunnel endpoint.

A MIP maps traffic arriving at the one address to another address; so by using the Untrust zone interface IP address as a MIP, the NetScreen device maps all inbound traffic using the Untrust zone interface to a specified internal address. If the MIP on the Untrust interface maps to a VPN or L2TP tunnel endpoint, the device automatically forwards the IKE or L2TP packets that it receives to the tunnel endpoint, as long as there is no VPN or L2TP tunnel configured on the Untrust interface.

If you create a policy in which the destination address is a MIP using the Untrust zone interface IP address and you specify HTTP as the service in the policy, you lose Web management of the NetScreen device via that interface (because all inbound HTTP traffic to that address is mapped to an internal server or host). You can still manage the device via the Untrust zone interface using the WebUI by changing the port number for Web management. To change the Web management port number, do the following:

1. Admin > Web: Enter a registered port number (from 1024 to 65,535) in the HTTP Port field. Then click **Apply**.
2. When you next connect to the Untrust zone interface to manage the device, append the port number to the IP address—for example, http://209.157.66.170:5000.

Example: MIP on the Untrust Interface

In this example, you select the IP address of the Untrust zone interface (ethernet3, 1.1.1.1/24) as the MIP for a Web server whose actual IP address is 10.1.1.5 in the Trust zone. Because you want to retain Web management access to the ethernet3 interface, you change the web management port number to 8080. You then create a policy permitting HTTP service (on the HTTP default port number—80) from the Untrust zone to the MIP—and consequently to the host with the address to which the MIP points—in the Trust zone.

WebUI

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Enter the following and then click **OK**:

NAT:[21] (select)

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

2. HTTP Port

Configuration > Admin > Management: Type **8080** in the HTTP Port field, and then click **Apply**.

(The HTTP connection is lost.)

21. By default, any interface that you bind to the Trust zone is in NAT mode. Consequently, this option is already enabled for interfaces bound to the Trust zone.

3. Reconnection

Reconnect to the NetScreen device, appending 8080 to the IP address in the URL address field in your web browser. (If you are currently managing the device via the untrust interface, type **http://1.1.1.1:8080**.)

4. MIP

Network > Interface > Edit (for ethernet3) > MIP > New: Enter the following, and then click **OK**:

Mapped IP: 1.1.1.1
Netmask: 255.255.255.255[22]
Host IP Address: 10.1.1.5
Host Virtual Router Name: trust-vr

5. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 0.0.0.0/0
Gateway: (select)
  Interface: ethernet3
  Gateway IP Address: 1.1.1.250

6. Policy

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:
  Address Book Entry: (select), Any
Destination Address:
  Address Book Entry: (select), MIP(1.1.1.1)
Service: HTTP
Action: Permit

22. The netmask for a MIP using an Untrust zone interface IP address must be 32 bits.

CLI

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
```

2. HTTP Port

```
set admin port 8080
```

3. MIP

```
set interface ethernet3 mip 1.1.1.1 host 10.1.1.5 netmask 255.255.255.255
    vrouter trust-vr[23]
```

4. Route

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
```

5. Policy

```
set policy from untrust to trust any mip(1.1.1.1) http permit
save
```

23. By default, the netmask for a MIP is 32 bits (255.255.255.255) and the default virtual router is the trust-vr. You do not have to specify them in the command. These arguments are included here to provide symmetry with the WebUI configuration.

MIP and the Loopback Interface

Defining a MIP on the loopback interface allows a MIP to be accessed by a group of interfaces. The primary application for this is to allow access to a host through one of several VPN tunnels using a single MIP address. The MIP host can also initiate traffic to a remote site through the appropriate tunnel.

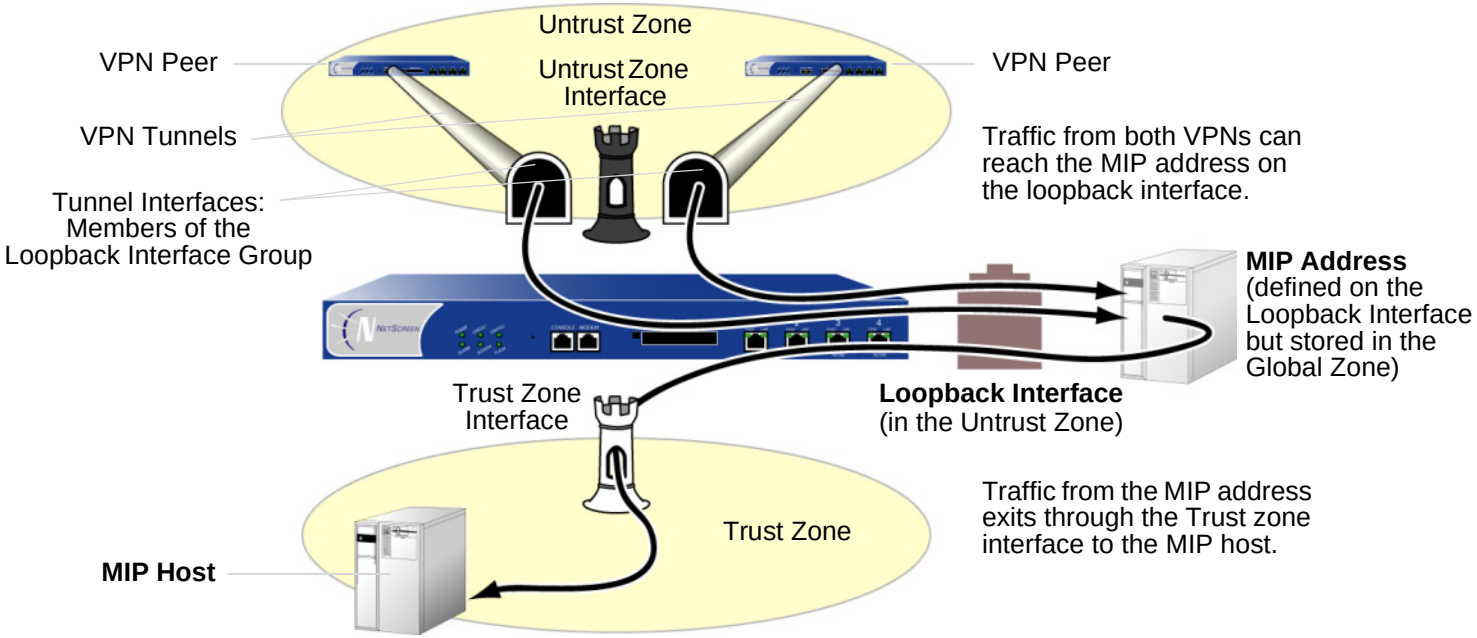

You can think of the loopback interface as a resource holder that contains a MIP address. You configure a loopback interface with the name loopback.*id_num* (where *id_num* is an index number that uniquely identifies the interface in the device), and assign an IP address to the interface (see "Loopback Interfaces" on page 86). To allow other interfaces to use a MIP on the loopback interface, you then add the interfaces as members of the loopback group.

The loopback interface and its member interfaces must be in different IP subnets in the same zone. Any type of interface can be a member of a loopback group as long as the interface has an IP address. If you configure a MIP on both a loopback interface and one of its member interfaces, the loopback interface configuration takes precedence. A loopback interface cannot be a member of another loopback group.

Example: MIP for Two Tunnel Interfaces

In this example, you configure the following interfaces:

- ethernet1, Trust zone, 10.1.1.1/24
- ethernet3, Untrust zone, 1.1.1.1/24
- tunnel.1, Untrust zone, 10.10.1.1/24, bound to vpn1
- tunnel.2, Untrust zone, 10.20.1.1/24, bound to vpn2
- loopback.3, Untrust zone, 10.100.1.1/24

The tunnel interfaces are members of the loopback.3 interface group. The loopback.3 interface contains MIP address 10.100.1.5, which maps to a host at 10.1.1.5 in the Trust zone.

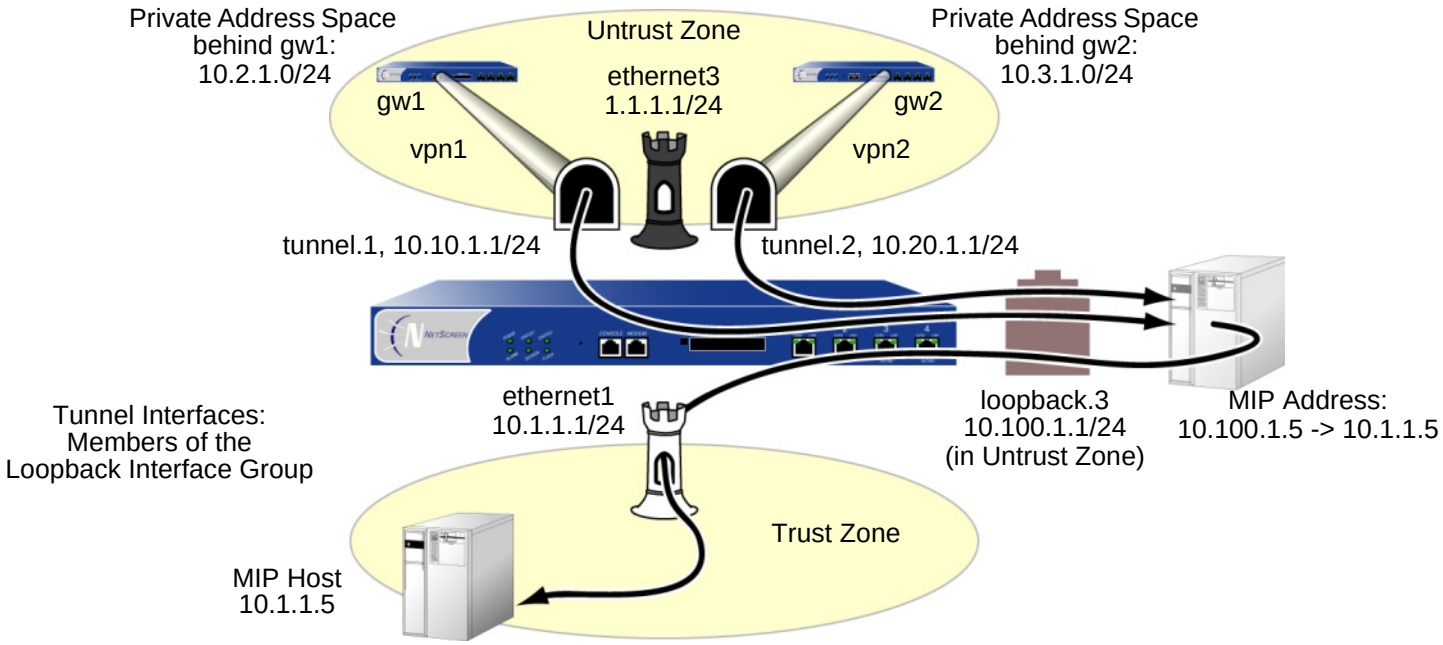

When a packet destined for 10.100.1.5 arrives at through a VPN tunnel to tunnel.1, the NetScreen device searches for the MIP on the loopback interface loopback.3. When it finds a match on loopback.3, the NetScreen device

translates the original destination IP (10.100.1.5) to the host IP address (10.1.1.5) and forwards the packet through ethernet1 to the MIP host. Traffic destined for 10.100.1.5 can also arrive through a VPN tunnel bound to tunnel.2. Again, the NetScreen device finds a match on loopback.3 and translates the original destination IP 10.100.1.5 to 10.1.1.5 and forwards the packet to the MIP host.

You also define addresses, VPN tunnels, routes, and policies as needed to complete the configuration. All security zones are in the trust-vr routing domain.

WebUI

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Enter the following, and then click **OK**:

NAT:[24] (select)

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

Network > Interfaces > New Loopback IF: Enter the following, and then click **OK**:

Interface Name: loopback.3

Zone: Untrust (trust-vr)

IP Address / Netmask: 10.100.1.1/24

24. By default, any interface that you bind to the Trust zone is in NAT mode. Consequently, this option is already enabled for interfaces bound to the Trust zone.

Network > Interfaces > New Tunnel IF: Enter the following, and then click **Apply**:

Tunnel Interface Name: tunnel.1

Zone (VR): Untrust (trust-vr)

Fixed IP: (select)

IP Address / Netmask: 10.10.1.1/24

Select **loopback.3** in the Member of Loopback Group drop-down list, and then click **OK**.

Network > Interfaces > New Tunnel IF: Enter the following, and then click **Apply**:

Tunnel Interface Name: tunnel.2

Zone (VR): Untrust (trust-vr)

Fixed IP: (select)

IP Address / Netmask: 10.20.1.1/24

Select **loopback.3** in the Member of Loopback Group drop-down list, and then click **OK**.

2. MIP

Network > Interfaces > Edit (for loopback.3) > MIP > New: Enter the following, and then click **OK**:

Mapped IP: 10.100.1.5

Netmask: 255.255.255.255

Host IP Address: 10.1.1.5

Host Virtual Router Name: trust-vr

3. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: local_lan

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.0/24

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: peer-1

IP Address/Domain Name:

IP/Netmask: (select), 10.2.1.0/24

Zone: Untrust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: peer-2

IP Address/Domain Name:

IP/Netmask: (select), 10.3.1.0/24

Zone: Untrust

4. VPNs

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: vpn1

Security Level: Compatible

Remote Gateway: Create a Simple Gateway: (select)

Gateway Name: gw1

Type: Static IP: (select), Address/Hostname: 2.2.2.2

Preshared Key: netscreen1

Security Level: Compatible

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Bind to Tunnel Interface: (select), tunnel.1

Proxy-ID: (select)

Local IP / Netmask: 0.0.0.0/0

Remote IP / Netmask: 0.0.0.0/0

Service: ANY

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: vpn2

Security Level: Compatible

Remote Gateway: Create a Simple Gateway: (select)

Gateway Name: gw2

Type: Static IP: (select), Address/Hostname: 3.3.3.3

Preshared Key: netscreen2

Security Level: Compatible

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Bind to Tunnel Interface: (select), tunnel.2

Proxy-ID: (select)

Local IP / Netmask: 0.0.0.0/0

Remote IP / Netmask: 0.0.0.0/0

Service: ANY

5. Routes

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 10.2.1.0/24

Gateway: (select)

Interface: tunnel.1

Gateway IP Address: 0.0.0.0

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 10.3.1.0/24

Gateway: (select)

Interface: tunnel.2

Gateway IP Address: 0.0.0.0

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 1.1.1.250

6. Policies

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), peer-1

Destination Address:

Address Book Entry: (select), MIP(10.100.1.5)

Service: ANY

Action: Permit

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), peer-2

Destination Address:

Address Book Entry: (select), MIP(10.100.1.5)

Service: ANY

Action: Permit

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), local_lan

Destination Address:

Address Book Entry: (select), Any

Service: ANY

Action: Permit

CLI

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24

set interface loopback.3 zone trust
set interface loopback.3 ip 10.100.1.1/24

set interface tunnel.1 zone untrust
set interface tunnel.1 ip 10.10.1.1/24
set interface tunnel.1 loopback-group loopback.3

set interface tunnel.2 zone untrust
set interface tunnel.2 ip 10.20.1.1/24
set interface tunnel.2 loopback-group loopback.3
```

2. MIP

```
set interface loopback.3 mip 10.100.1.5 host 10.1.1.5 netmask 255.255.255.255
    vrouter trust-vr[25]
```

3. Addresses

```
set address trust local_lan 10.1.1.0/24
set address untrust peer-1 10.2.1.0/24
set address untrust peer-2 10.3.1.0/24
```

25. By default, the netmask for a MIP is 32 bits (255.255.255.255) and the default virtual router is the trust-vr. You do not have to specify them in the command. These arguments are included here to provide symmetry with the WebUI configuration.

4. VPNs

```
set ike gateway gw1 address 2.2.2.2 outgoing-interface ethernet3 preshare
    netscreen1 sec-level compatible
set vpn vpn1 gateway gw1 sec-level compatible
set vpn vpn1 bind interface tunnel.1
set vpn vpn1 proxy-id local-ip 0.0.0.0/0 remote-ip 0.0.0.0/0 any

set ike gateway gw2 address 3.3.3.3 outgoing-interface ethernet3 preshare
    netscreen2 sec-level compatible
set vpn vpn2 gateway gw2 sec-level compatible
set vpn vpn2 bind interface tunnel.2
set vpn vpn2 proxy-id local-ip 0.0.0.0/0 remote-ip 0.0.0.0/0 any
```

5. Routes

```
set vrouter trust-vr route 10.2.1.0/24 interface tunnel.1
set vrouter trust-vr route 10.3.1.0/24 interface tunnel.2
set vrouter untrust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
```

6. Policies

```
set policy top from untrust to trust peer-1 mip(10.100.1.5) any permit
set policy top from untrust to trust peer-2 mip(10.100.1.5) any permit
set policy from trust to untrust local_lan any any permit
save
```

VIRTUAL IP ADDRESSES

A virtual IP (VIP) address maps traffic received at one IP address to another address based on the destination port number in the TCP or UDP segment header. For example,

- An HTTP packet destined for 1.1.1.3:80 (that is, IP address 1.1.1.3 and port 80) might get mapped to a web server at 10.1.1.10.
- An FTP packet destined for 1.1.1.3:21 might get mapped to an FTP server at 10.1.1.20.
- An SMTP packet destined for 1.1.1.3:25 might get mapped to a mail server at 10.1.1.30.

The destination IP addresses are the same. The destination port numbers determine the host to which the NetScreen device forwards traffic.

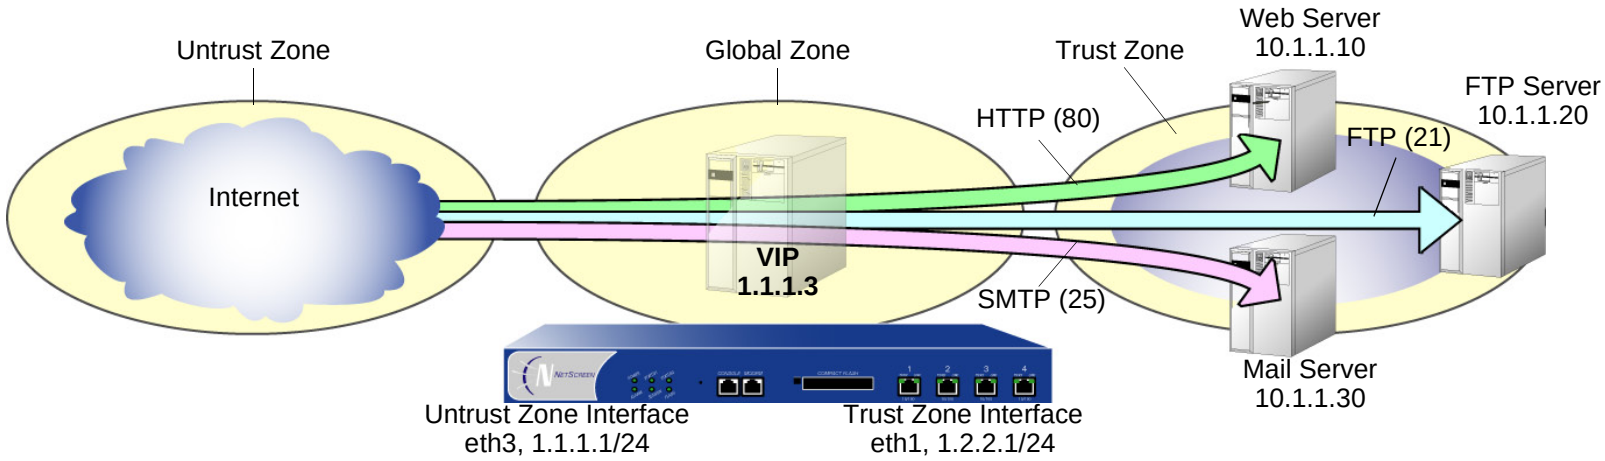

Virtual IP Forwarding Table

Interface IP in Untrust Zone	VIP in Global Zone	Port	Forward to	Host IP in Trust Zone
1.1.1.1/24	*1.1.1.3*	*80 (HTTP)*	→	*10.1.1.10*
1.1.1.1/24	*1.1.1.3*	*21 (FTP)*	→	*10.1.1.20*
1.1.1.1/24	*1.1.1.3*	*25 (SMTP)*	→	*10.1.1.30*

The NetScreen device forwards incoming traffic destined for a VIP to the host with the address to which the VIP points. When a VIP host initiates outbound traffic, the NetScreen device translates the source IP address of the host to that of the VIP address. This bidirectional translation symmetry differs from the behavior of source and destination address translation (see "Directional Nature of NAT-Src and NAT-Dst" on page 257).

You need the following information to define a Virtual IP:

- The IP address for the VIP must be in the same subnet as an interface in the Untrust zone or—on some NetScreen devices—can even be the same address as that interface[26]
- The IP addresses for the servers that process the requests
- The type of service you want the NetScreen device to forward from the VIP to the IP address of the host

> ***Note:*** *You can only set a VIP on an interface in the Untrust zone.*

Some notes about NetScreen VIPs:

- You can use virtual port numbers for well-known services when running multiple server processes on a single machine. For example, if you have two FTP servers on the same machine, you can run one server on port 21 and the other on port 2121. Only those who know the virtual port number in advance and append it to the IP address in the packet header can gain access to the second FTP server.
- You can map predefined services and user-defined services.
- A single VIP can distinguish custom services with the same source and destination port numbers but different transports.
- Custom services can use any destination port number or number range from 1 to 65,535, not just from 1024 to 65,535.

26. On some NetScreen devices, an interface in the Untrust zone can receive its IP address dynamically via DHCP or PPPoE. If you want to use a VIP in such a situation, do either of the following: In the WebUI (Network > Interfaces > Edit (for an interface in the Untrust zone) > VIP: Select the **Same as the untrusted interface IP address** option when setting up the VIP. In the CLI, use the **set interface** *name* **vip untrust-ip** command.

    If you configure a VIP to use the same IP address as an Untrust zone interface on a NetScreen device that supports multiple VIPs, the other "regular" VIPs become unusable. If a regular VIP is configured, you cannot create a VIP using an Untrust zone interface until you delete the regular VIP first.

- A single VIP can support custom services with multiple port entries by creating multiple service entries under that VIP—one service entry in the VIP for each port entry in the service. By default, you can use single-port services in a VIP. To be able to use multiple-port services in a VIP, you must first issue the CLI command **set vip multi-port**, and then reset the NetScreen device. (See "Example: VIP with Custom and Multiple-Port Services" on page 363.)
- The host to which the NetScreen device maps VIP traffic must be reachable from the trust-vr. If the host is in a routing domain other than that of the trust-vr, you must define a route to reach it.
- Custom services can use any destination port number or number range from 1 to 32,767, not just from 1024 to 32,767.

VIP and the Global Zone

Setting a VIP for an interface in the Untrust zone generates an entry in the Global zone address book. The Global zone address book keeps all the VIPs of all interfaces, regardless of the zone to which the interface belongs. You can use these VIP addresses as the destination address in policies between any two zones, and as the destination address in Global policies.

Example: Configuring Virtual IP Servers

In this example, you bind interface ethernet1 to the Trust zone and assign it IP address 10.1.1.1/24. You bind interface ethernet3 to the Untrust zone and assign it IP address 1.1.1.1/24.

Then, you configure a VIP at 1.1.1.10 to forward inbound HTTP traffic to a Web server at 10.1.1.10, and you create a policy permitting traffic from the Untrust zone to reach the VIP—and consequently to the host with the address to which the VIP points—in the Trust zone.

Because the VIP is in the same subnet as the Untrust zone interface (1.1.1.0/24), you do not need to define a route for traffic from the Untrust zone to reach it[27]. Also, no address book entry is required for the host to which a VIP forwards traffic. All security zones are in the trust-vr routing domain.

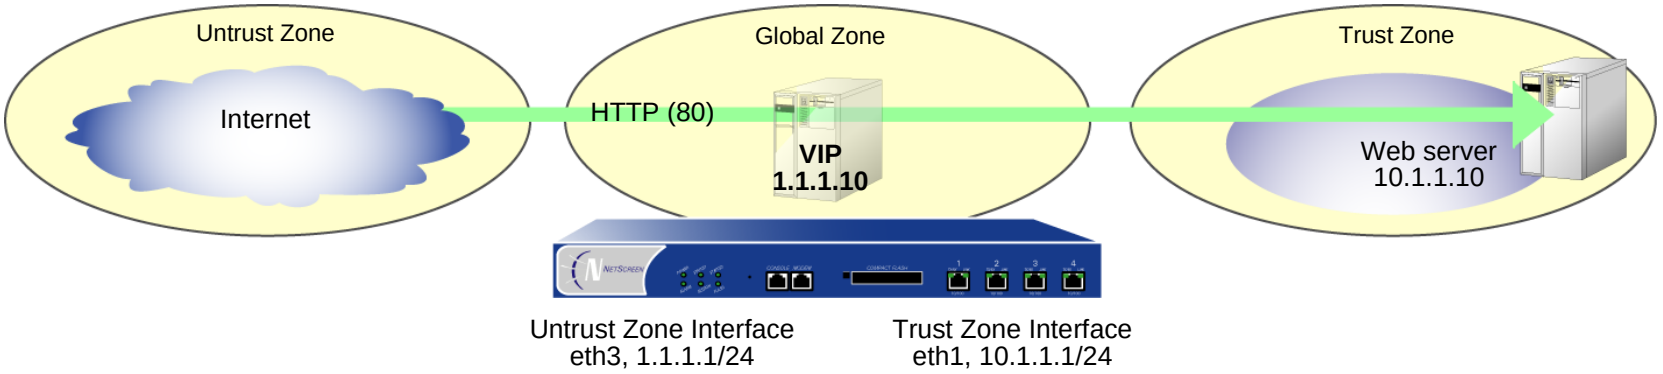

27. If you want HTTP traffic from a security zone other than the Untrust zone to reach the VIP, you must set a route for 1.1.1.10 on the router in the other zone to point to an interface bound to that zone. For example, imagine that ethernet2 is bound to a user-defined zone, and you have configured a router in that zone to send traffic destined for 1.1.1.10 to ethernet2. After the router sends traffic to ethernet2, the forwarding mechanism in the NetScreen device locates the VIP on ethernet3, which maps it to 10.1.1.10, and sends it out ethernet1 to the Trust zone. This process is similar to that described in "Example: Reaching a MIP from Different Zones" on page 336. You must also set a policy permitting HTTP traffic from the source zone to the VIP in the Trust zone.

WebUI

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

2. VIP

Network > Interfaces > Edit (for ethernet3) > VIP: Enter the following address, and then click **Add**:

Virtual IP Address: 1.1.1.10

Network > Interfaces > Edit (for ethernet3) > VIP > New VIP Service: Enter the following, and then click **OK**:

Virtual IP: 1.1.1.10

Virtual Port: 80

Map to Service: HTTP (80)

Map to IP: 10.1.1.10

3. Policy

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), ANY

Destination Address:

Address Book Entry: (select), VIP(1.1.1.10)

Service: HTTP

Action: Permit

CLI

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet2 ip 1.1.1.1/24
```

2. VIP

```
set interface ethernet3 vip 1.1.1.10 80 http 10.1.1.10
```

3. Policy

```
set policy from untrust to trust any vip(1.1.1.10) http permit
save
```

Example: Editing a VIP Configuration

In this example, you modify the Virtual IP server configuration you created in the previous example. To restrict access to the Web server, you change the virtual port number for HTTP traffic from 80 (the default) to 2211. Now, only those that know to use port number 2211 when connecting to the Web server can access it.

WebUI

Network > Interfaces > Edit (for ethernet3) > VIP > Edit (in the VIP Services Configure section for 1.1.1.10): Enter the following, and then click **OK**:

Virtual Port: 2211

CLI

```
unset interface ethernet3 vip 1.1.1.10 port 80
set interface ethernet3 vip 1.1.1.10 2211 http 10.1.1.10
save
```

Example: Removing a VIP Configuration

In this example, you delete the VIP configuration that you previously created and modified. Before you can remove a VIP, you must first remove any existing policies associated with it. The ID number for the policy that you created in “Example: Configuring Virtual IP Servers” on page 359 is 5.

WebUI

Policies > (From: Untrust, To: Trust) > Go: Click **Remove** for policy ID 5.

Network > Interfaces > Edit (for ethernet3) > VIP: Click **Remove** in the VIP Configure section for 1.1.1.10.

CLI

```
unset policy id 5
unset interface ethernet3 vip 1.1.1.10
save
```

Example: VIP with Custom and Multiple-Port Services

In the following example, you configure a VIP at 1.1.1.3 to route the following services to the following internal addresses:

Service	Transport	Virtual Port Number	Actual Port Number	Host IP Address
DNS	TCP, UDP	53	53	10.1.1.3
HTTP	TCP	80	80	10.1.1.4
PCAnywhere	TCP, UDP	5631, 5632	5631, 5632	10.1.1.4
LDAP	TCP, UDP	5983	389	10.1.1.5

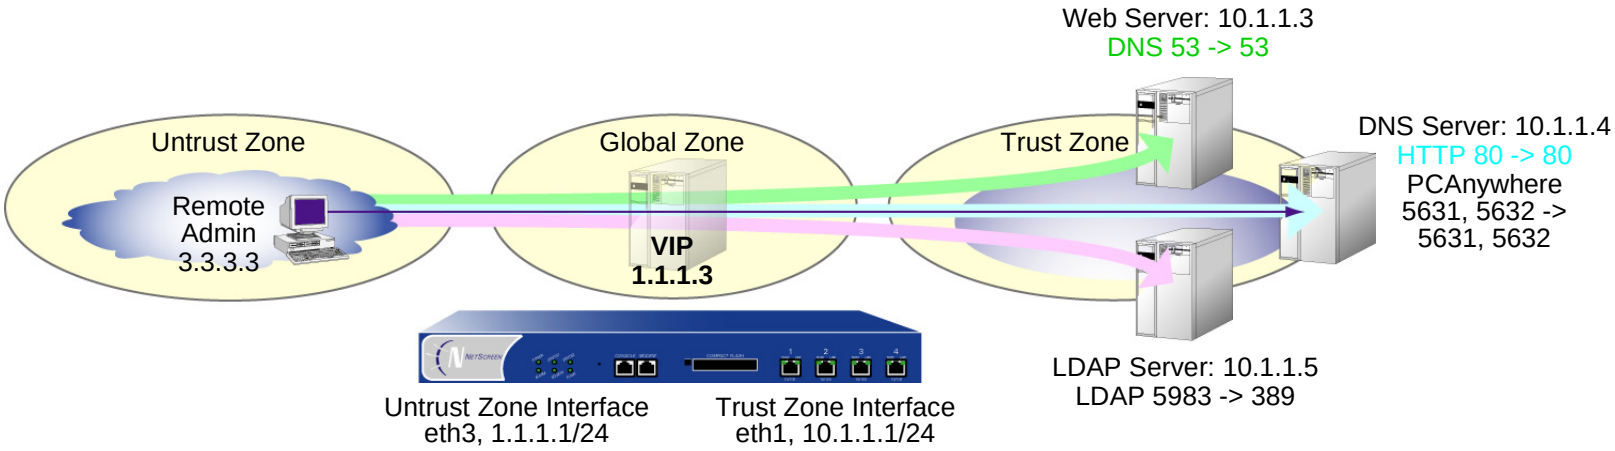

The VIP routes DNS lookups to the DNS server at 10.1.1.3, HTTP traffic to the web server at 10.1.1.4, and authentication checks to the database on the LDAP server at 10.1.1.5. For HTTP, DNS, and PCAnywhere, the virtual port numbers remain the same as the actual port numbers. For LDAP, a virtual port number (5983) is used to add an extra level of security to the LDAP authentication traffic.

For managing the HTTP server remotely, you define a custom service and name it PCAnywhere. PCAnywhere is a multiple-port service that sends and listens for data on TCP port 5631 and status checks on UDP port 5632.

You also enter the address of the Remote Admin at 3.3.3.3 in the Untrust zone address book, and configure policies from the Untrust zone to the Trust zone for all the traffic that you want to use the VIPs. All security zones are in the trust-vr routing domain.

WebUI

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

2. Address

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Remote Admin

IP Address/Domain Name:

IP/Netmask: (select), 3.3.3.3/32

Zone: Untrust

3. Custom Service

Object > Services > Custom > New: Enter the following, and then click **OK**:

Service Name: PCAnywhere
No 1:
Transport protocol: TCP
Source Port Low: 0
Source Port High: 65535
Destination Port Low: 5631
Destination Port High: 5631
No 2:
Transport protocol: UDP
Source Port Low: 0
Source Port High: 65535
Destination Port Low: 5632
Destination Port High: 5632

4. VIP Address and Services[28]

Network > Interfaces > Edit (for ethernet3) > VIP: click here to configure: Type **1.1.1.3** in the Virtual IP Address field, and then click **Add**.

> New VIP Service: Enter the following, and then click **OK**:

Virtual IP: 1.1.1.3
Virtual Port: 53
Map to Service: DNS
Map to IP: 10.1.1.3

28. To enable the VIP to support multiple-port services, you must use enter the CLI command **set vip multi-port**, save the configuration, and then reboot the device.

> New VIP Service: Enter the following, and then click **OK**:

Virtual IP: 1.1.1.3

Virtual Port: 80

Map to Service: HTTP

Map to IP: 10.1.1.4

> New VIP Service: Enter the following, and then click **OK**:

Virtual IP: 1.1.1.3

Virtual Port: 5631[29]

Map to Service: PCAnywhere

Map to IP: 10.1.1.4

> New VIP Service: Enter the following, and then click **OK**:

Virtual IP: 1.1.1.3

Virtual Port: 5983[30]

Map to Service: LDAP

Map to IP: 10.1.1.5

29. For multiple-port services, enter the lowest port number of the service as the virtual port number.

30. Using non-standard port numbers adds another layer of security, thwarting common attacks that check for services at standard port numbers.

5. Policies

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), VIP(1.1.1.3)

Service: DNS

Action: Permit

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), VIP(1.1.1.3)

Service: HTTP

Action: Permit

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), VIP(1.1.1.3)

Service: LDAP

Action: Permit

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Remote Admin

Destination Address:

Address Book Entry: (select), VIP(1.1.1.3)

Service: PCAnywhere

Action: Permit

CLI

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
```

2. Address

```
set address untrust “Remote Admin” 3.3.3.3/32
```

3. Custom Service

```
set service pcanywhere protocol udp src-port 0-65535 dst-port 5631-5631
set service pcanywhere + tcp src-port 0-65535 dst-port 5632-5632
```

4. VIP Address and Services

```
set vip multi-port
save
reset
System reset, are you sure? y/[n] y
```

```
set interface ethernet3 vip 1.1.1.3 53 dns 10.1.1.3
set vip 1.1.1.3 + 80 http 10.1.1.4
set vip 1.1.1.3 + 5631 pcanywhere 10.1.1.4[31]
set vip 1.1.1.3 + 5983 ldap 10.1.1.5
```

5. Policies

```
set policy from untrust to trust any vip(1.1.1.3) dns permit
set policy from untrust to trust any vip(1.1.1.3) http permit
set policy from untrust to trust any vip(1.1.1.3) ldap permit
set policy from untrust to trust “Remote Admin” vip(1.1.1.3) pcanywhere permit
save
```

31. For multiple-port services, enter the lowest port number of the service as the virtual port number.

9

User Authentication

This chapter focuses on the methods available for authenticating users. It begins by examining different kinds of authentication servers—the local database built into every NetScreen device, and external RADIUS, SecurID, and LDAP authentication servers. It then describes how to define different user accounts (or "profiles"), how to create user groups, and how to reference those users and user groups in policies, AutoKey IKE gateways, and L2TP tunnels. The chapter concludes with a section about customizing the banners that appear on HTTP, FTP, L2TP, Telnet, and XAuth login prompts. This material is presented in the following sections:

- "Authentication Servers" on page 372
- "Local Database" on page 374
- "External Auth Servers" on page 376
  - "Auth Server Types" on page 379
  - "Defining Auth Server Objects" on page 388
  - "Defining Default Auth Servers" on page 395
- "Authentication Types and Applications" on page 397
  - "Auth Users and User Groups" on page 398
  - "IKE Users and User Groups" on page 431
  - "XAuth Users and User Groups" on page 436
  - "L2TP Users and User Groups" on page 460
  - "Admin Users" on page 465
- "Multiple-Type Users" on page 467
- "Group Expressions" on page 468
- "Banner Customization" on page 476

AUTHENTICATION SERVERS

You can configure the NetScreen device to use the local database or one or more external authentication servers to verify the identities of the following types of users:

- Auth users
- IKE users
- L2TP users
- XAuth users
- Admin users

> ***Note:*** *IKE user accounts must be stored on the local database. The only external server to support L2TP and XAuth remote setting assignments and admin privilege assignments is RADIUS.*

In addition to its local database, a NetScreen device supports external RADIUS, SecurID and LDAP servers. You can use each kind of authentication server to authenticate auth users, L2TP users, XAuth users, and admin users. NetScreen also supports WebAuth, an alternative authentication scheme for auth users. (For a WebAuth example, see "Example: WebAuth + SSL (External User Group)" on page 427.) Any auth server that contains auth user account types is eligible to be the default WebAuth auth server. The following table summarizes which servers support which user types and authentication features:

Server Type	Supported User Types and Features									
	Auth Users	IKE Users	L2TP Users		XAuth Users		Admin Users		User Groups	Group Expressions
			Auth	Remote Settings	Auth	Remote Settings	Auth	Privileges		
Local	✓	✓	✓	✓	✓	✓	✓	✓	✓	
RADIUS	✓		✓	✓	✓	✓	✓	✓	✓	✓
SecurID	✓		✓		✓		✓			
LDAP	✓		✓		✓		✓			

On most NetScreen devices, you can simultaneously employ up to 10 primary authentication servers per system—root system and virtual system—in any combination of types. This total includes the local database and excludes backup authentication servers. A RADIUS or LDAP server supports two backup servers, and a SecurID server supports one backup server; so, for example, you might use the local database and 9 different primary RADIUS servers, with each RADIUS server having two backup servers assigned to it.

Multiple authentication servers operating simultaneously

The color of each connection request relates to the matching color of the authentication check:

IKE/XAuth User (orange) –> Local Database

IKE User (red) –> Local Database

IKE/L2TP User (green) –> SecurID server

Admin User (purple) –> RADIUS server

Auth User (blue) –> LDAP server

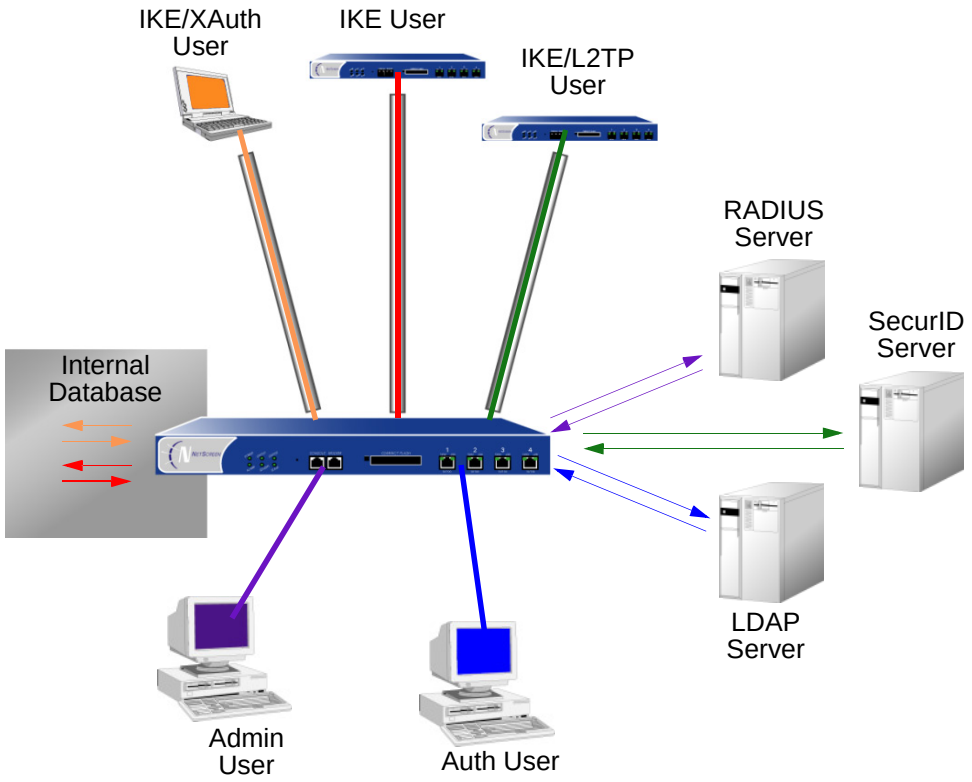

Note: *You can use a single authentication server for multiple types of user authentication. For example, a RADIUS server can simultaneously store admin, auth, L2TP, and XAuth users.*

The following sections explore the local database and each authentication server in greater detail.

LOCAL DATABASE

All NetScreen devices support a built-in user database for authentication. When you define a user on the NetScreen device, the NetScreen device enters the user name and password in its local database.

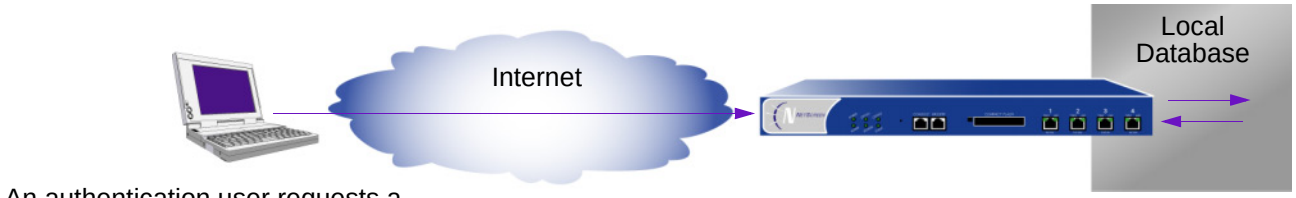

An authentication user requests a connection to the protected network.

The NetScreen device consults its internal database to authenticate the authentication user before approving the connection request.

Supported User Types and Features

The local database supports the following types of users and authentication features:

- Auth users
- IKE users
- L2TP users
- XAuth users
- Admin users
- Admin privileges
- WebAuth
- User groups
- Group expressions*

* You define the group expressions on the NetScreen device, but the users and user groups must be stored on an external RADIUS auth server. For more information about group expressions, see "Group Expressions" on page 468.

The local database is the default authentication server (auth server) for all types of authentication. For instructions on how to add users and user groups to the local database via the WebUI and CLI, see "Authentication Types and Applications" on page 397.

Example: Setting the Local Database Timeout

By default, the local database authentication timeout for both admins and auth users is 10 minutes. In this example, you change it to never time out for admins and to time out after 30 minutes for auth users.

WebUI

Configuration > Admin > Management: Clear the Enable Web Management Idle Timeout check box, and then click **Apply**.

Configuration > Auth > Servers > Edit (for Local): Enter **30** in the Timeout field, and then click **Apply**.

CLI

```
set admin auth timeout 0
set auth-server Local timeout 30
save
```

EXTERNAL AUTH SERVERS

A NetScreen device can connect to one or more external authentication servers, or “auth servers”, on which you store user accounts. When the NetScreen device receives a connection request that requires authentication verification, the NetScreen device requests an authentication check from the external auth server specified in the policy, L2TP tunnel configuration, or IKE gateway configuration. The NetScreen then acts as a relay between the user requesting authentication and the auth server granting authentication. A successful authentication check by an external auth server proceeds as follows:

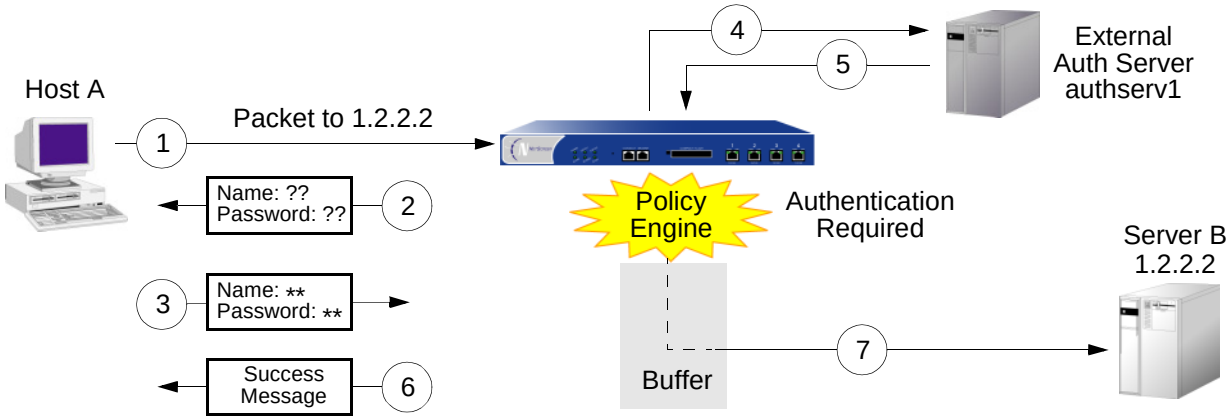

1. Host A sends an FTP, HTTP, or Telnet TCP SYN packet to 1.2.2.2.
2. The NetScreen device intercepts the packet, notes that its corresponding policy requires authentication from authserv1, buffers the packet, and prompts the user for a user name and password.
3. The user replies with a user name and password.
4. The NetScreen device relays the login information to authserv1.
5. Authserv1 sends back a notification of success to the NetScreen device.
6. The NetScreen device informs the auth user of his or her login success.
7. The NetScreen device then forwards the packet from its buffer to its destination of 1.2.2.2.

Auth Server Object Properties

A NetScreen device treats each auth server as an object that it can reference in policies, IKE gateways, and L2TP tunnels. The following properties define and uniquely identify an auth server object:

- Object name: A name string, such as “authserv1” (The only predefined auth server is “Local”.)
- ID number: You can set the ID number or allow the NetScreen device to set it automatically. If you set an ID number, you must choose one that is not already in use.
- Type: RADIUS, SecurID, LDAP.
- Server name: The IP address or domain name of the server
- Backup1: The IP address or domain name of a primary backup server
- Backup2: (RADIUS and LDAP) The IP address or domain name of a secondary backup server
- Account Type: One or more of the following types of users: Auth, L2TP, XAuth; or Admin by itself.
- Timeout value: The timeout value takes on a different meaning if it is for an auth user or if it is for an admin user.
  - Auth user: The timeout countdown begins after the first authenticated session completes. If the user initiates a new session before the countdown reaches the timeout threshold, he does not have to reauthenticate himself and the timeout countdown resets. The default timeout value is 10 minutes, and the maximum is 255 minutes. You can also set the timeout value at 0 so that the authentication period never times out.

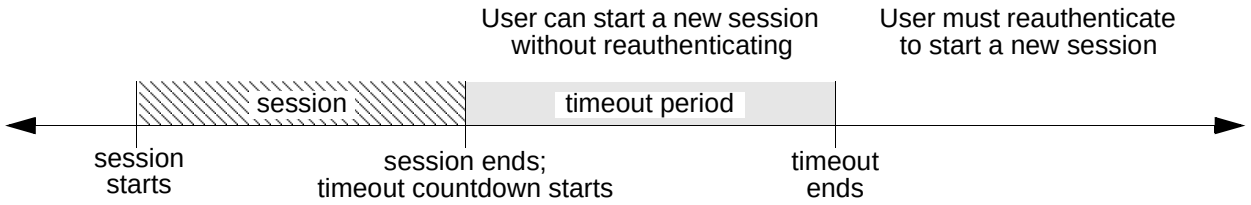

Note: *User authentication timeout is not the same as session idle timeout. If no activity occurs in a session for a predefined length of time, the NetScreen device automatically removes the session from its session table.*

– Admin user: If the the length of idle time reaches the timeout threshold, the NetScreen device terminates the admin session. To continue managing the NetScreen device, the admin must reconnect to the device and reauthenticate himself. The default timeout value is 10 minutes, and the maximum is 1000 minutes. You can also set the timeout value at 0 so that an admin session never times out.

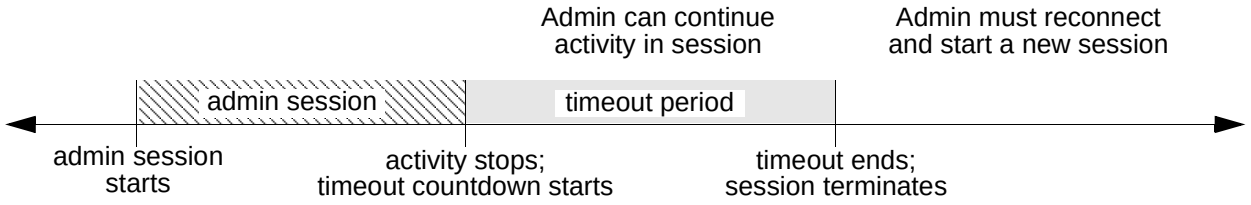

In addition to the above properties that apply to all auth server objects, each server has a few others specific to itself. These are explained in the RADIUS, SecurID, and LDAP auth server property sections that follow.

Auth Server Types

In addition to the internal database, NetScreen devices support three types of external auth servers: RADIUS, SecurID, and LDAP.

RADIUS

The Remote Authentication Dial-In User Service (RADIUS) is a protocol for an authentication server that can support up to tens of thousands of users.

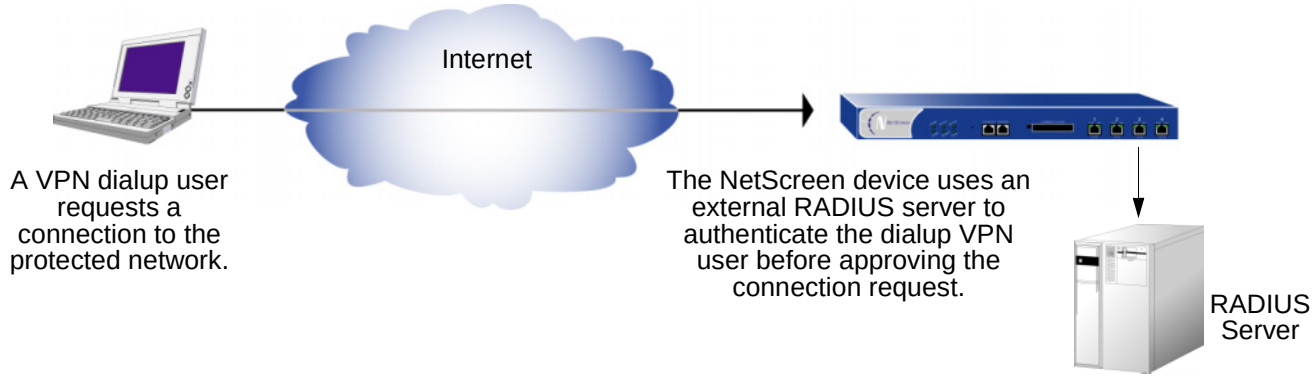

The RADIUS client (that is, the NetScreen device) authenticates users through a series of communications between the client and the server. Basically, RADIUS asks the person logging on to enter his or her user name and password. It then compares these values to those in its database, and once a user is authenticated, the client provides the user with access to the appropriate network services.

To configure the NetScreen device for RADIUS, you must specify the IP address of the RADIUS server and define a shared secret—the same as that defined on the RADIUS server. The shared secret is a password the RADIUS server uses to generate a key to encrypt traffic between the NetScreen and RADIUS devices.

RADIUS Auth Server Object Properties

In addition to the generic auth server properties listed in “Auth Server Object Properties” on page 377, a RADIUS server also makes use of the following properties:

- **Shared Secret:** The secret (password) shared between the NetScreen device and the RADIUS server. The devices use this secret to encrypt the user’s password that it sends to the RADIUS server.
- **RADIUS Port:** The port number on the RADIUS server to which the NetScreen device sends authentication requests. The default port number is 1645.
- **RADIUS Retry Timeout:** The interval (in seconds) that the NetScreen device waits before sending another authentication request to the RADIUS server if the previous request does not elicit a response. The default is three seconds.

Supported User Types and Features

A RADIUS server supports the following types of users and authentication features:

- Auth users
- L2TP users (authentication and remote settings)
- XAuth users (authentication and remote settings)
- Admin users (authentication and privilege assignments)
- User groups

A RADIUS server can support all of the user types and features that the local database supports except IKE users. Among the three types of external auth servers, RADIUS is the only one at this time with such broad support. For a RADIUS server to support such NetScreen-specific attributes as admin privileges, user groups, and remote L2TP and XAuth IP address[1], and DNS and WINS server address assignments, you must load a NetScreen dictionary file that defines these attributes onto the RADIUS server.

1. NetScreen uses the standard RADIUS attribute for IP address assignments. If you only want to use RADIUS for IP address assignments, you do not have to load the NetScreen vendor-specific attributes (VSAs).

NetScreen Dictionary File

A dictionary file defines vendor-specific attributes (VSAs) that you can load onto a RADIUS server. After defining values for these VSAs, NetScreen can then query them when a user logs on to the NetScreen device. NetScreen VSAs include admin privileges, user groups, and remote L2TP and XAuth IP address, and DNS and WINS server address assignments. There are two NetScreen dictionary files, one for Cisco RADIUS servers and one for Funk Software RADIUS servers. If using a Microsoft RADIUS server, there is no dictionary file. You must configure it as outlined in *Using a Windows NT Domain / Active Directory for User Authentication NetScreen Devices*, which you can download from www.netscreen.com/resources/application_notes/technical.jsp.

Each NetScreen dictionary file contains the following specific information:

- **Vendor ID:** The NetScreen vendor ID (VID; also called an "IETF number") is 3224. The VID identifies a specific vendor for a particular attribute. Some types of RADIUS server require you to enter the VID for each attribute entry, while other types only require you to enter it once and then apply it globally. Refer to your RADIUS server documentation for further information.
- **Attribute Name:** The attribute names describe individual NetScreen-specific attributes, such as NS-Admin-Privilege, NS-User-Group, NS-Primary-DNS-Server, and so forth.
- **Attribute Number:** The attribute number identifies an individual vendor-specific attribute. NetScreen-specific attribute numbers fall into two ranges:
  - NetScreen ScreenOS: 1 – 199
  - NetScreen-Global PRO: 200 and above

  For example, the ScreenOS attribute number for user groups is 3. The NetScreen-Global PRO attribute number for user groups is 200.
- **Attribute Type:** The attribute type identifies the form in which attribute data (or "value") appears—a string, an IP address, or an integer.

The RADIUS server automatically receives the above information when you load the NetScreen dictionary file onto it. To make new data entries, you must manually enter a value in the form indicated by the attribute type. For example, an entry for a read-write admin appears as follows:

VID	Attribute Name	Attribute Number	Attribute Type	Value
3224	NS-Admin-Privileges	1	data=int4 (ie, integer)	2 (2 = all privileges)

To download a dictionary file, go to www.netscreen.com/services/tac_online/index.jsp, select a NetScreen product, and then select a RADIUS dictionary file.

RADIUS Access-Challenge

NetScreen devices can now process access-challenge packets from an external RADIUS server when an authentication user attempts to log on via Telnet. Access-challenge presents an additional condition to the login process after the approval of a user name and password. After an authentication user responds to a login prompt with the correct user name and password, the RADIUS server sends an access-challenge to the NetScreen device, which then forwards it to the user. When the user replies, the NetScreen device sends a new access-request with the user's response to the RADIUS server. If the user's response is correct, the authentication process concludes successfully. Consider the following scenario in which an authentication user wants to telnet to a server:

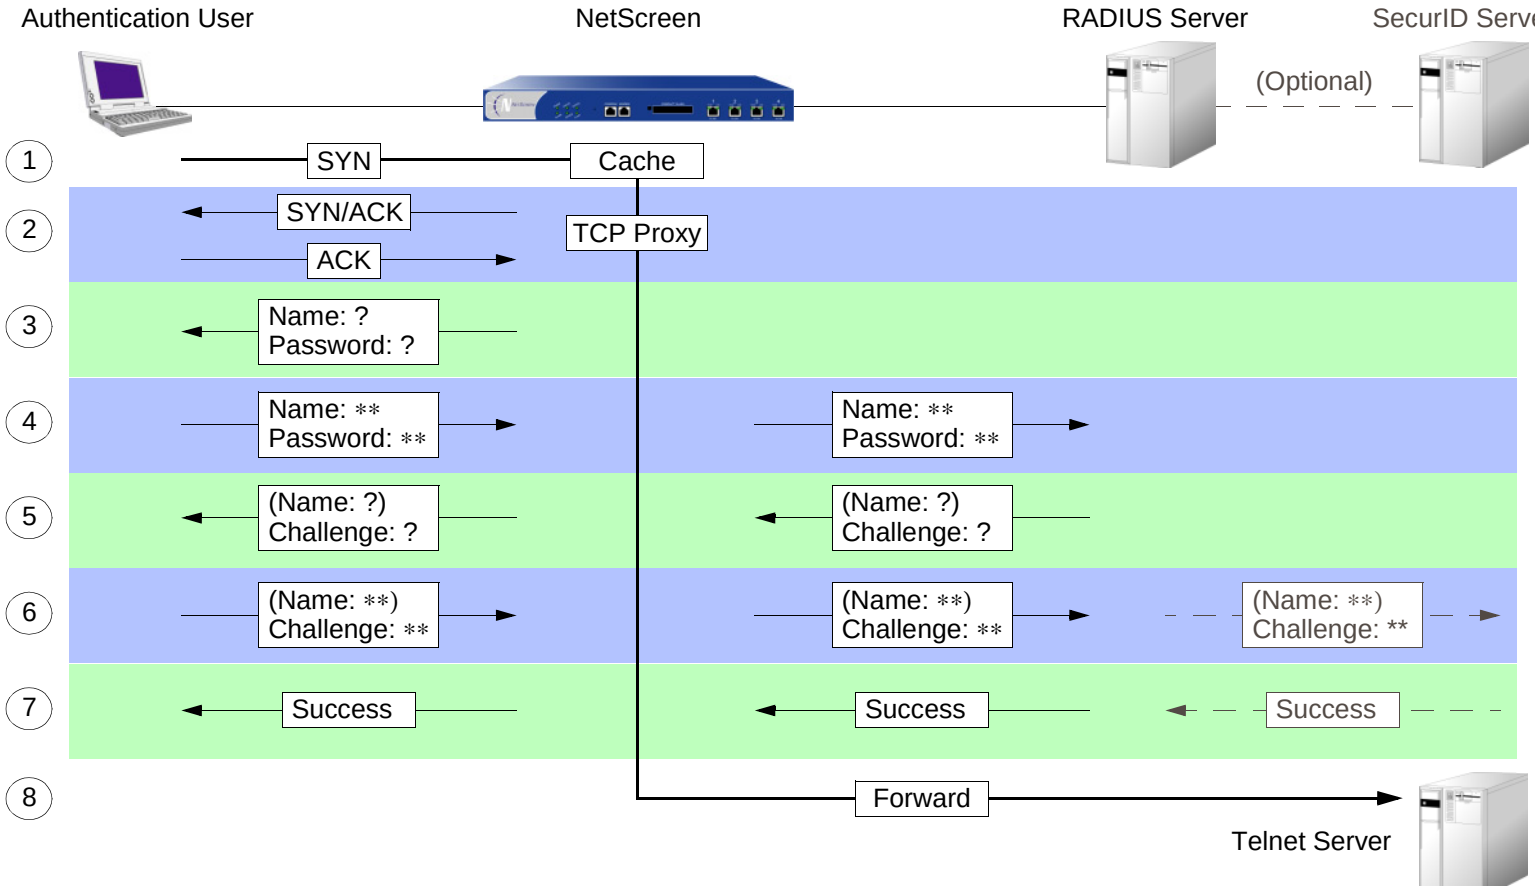

1. An authentication user sends a SYN packet to initiate a TCP connection for a Telnet session to a Telnet server.
2. A NetScreen device intercepts the packet, checks its policy list, and determines that this session requires user authentication. The NetScreen device caches the SYN packet and proxies the TCP 3-way handshake with the user.
3. The NetScreen device prompts the user to log in with a user name and password.
4. The authentication user enters his or her user name and password and sends it to the NetScreen device. The NetScreen device then sends an access-request with the login information to a RADIUS server.
5. If the information is correct, the RADIUS server sends the NetScreen device an access-challenge with a reply-message attribute that prompts the user to provide a response to a challenge. (The access-challenge can optionally prompt the authentication to provide a user name again. The second user name can be the same as the first or a different one.) The NetScreen device then sends the user another login prompt that contains the content of the reply-message attribute.
6. The authentication user enters his or her challenge response (and optionally a user name) and sends it to the NetScreen device. The NetScreen device then sends a second access-request, with the user's challenge-response, to the RADIUS server.

   If the RADIUS server needs to authenticate the challenge-response via another auth server—for example, if a SecurID server must authenticate a token code—the RADIUS server sends the access-request to the other auth server.
7. If the RADIUS server forwarded the challenge-response to another auth server and that server sends an access-accept, or if the RADIUS server itself approves the challenge-response, the RADIUS server sends an access-accept message to the NetScreen device. The NetScreen device then notifies the authentication user that his or her login is successful.
8. The NetScreen device forwards the initial SYN packet to its original destination: the Telnet server.

Note: *NetScreen does not support access-challenge with L2TP at the time of this release.*

SecurID

Instead of a fixed password, SecurID combines two factors to create a dynamically changing password. SecurID issues a credit card sized device called an authenticator that has an LCD window that displays a randomly generated string of numbers called a token code that changes every minute. The user also has a personal identification number (PIN). When the user logs on, he enters a user name and his PIN plus the current token code.

SecurID Authentication Device (Authenticator)

The token code changes to a different pseudo-random number every 60 seconds.

The authenticator performs an algorithm known only by RSA to create the values that appear in the LCD window. When the user to be authenticated enters his PIN and the number on his card, the ACE server, which also performs the same algorithm, compares the values received with those in its database. If they match, the authentication is successful.

The relationship of NetScreen device and a RSA SecurID ACE server is similar to that of a NetScreen device and a RADIUS server. That is, the NetScreen device acts as a client, forwarding authentication requests to the external server for approval and relaying login information between the user and the server. SecurID differs from RADIUS in that the user's "password" involves a continually changing token code.

SecurID Auth Server Object Properties

In addition to the generic auth server properties listed in “Auth Server Object Properties” on page 377, a SecurID server also makes use of the following properties:

- **Authentication Port:** The port number on the SecurID ACE server to which the NetScreen device sends authentication requests. The default port number is 5500.
- **Encryption Type:** The algorithm used for encrypting communication between the NetScreen device and the SecurID ACE server—either SDI or DES.
- **Client Retries:** The number of times that the SecurID client (that is, the NetScreen device) tries to establish communication with the SecurID ACE server before aborting the attempt.
- **Client Timeout:** The length of time in seconds that the NetScreen device waits between authentication retry attempts.
- **Use Duress:** An option that prevents or allows use of a different PIN number. When this option is enabled, and a user enters a previously determined duress PIN number, the NetScreen device sends a signal to the SecurID ACE server, indicating that the user is performing the login against his or her will; that is, while under duress. The SecurID ACE server permits access that one time, and then it denies any further login attempts by that user until he or she contacts the SecurID administrator. Duress mode is available only if the SecurID ACE server supports this option.

Supported User Types and Features

A SecurID Ace server supports the following types of users and authentication features:

- Auth users
- L2TP users (user authentication; L2TP user receives default L2TP settings from the NetScreen device)
- XAuth users (user authentication; no support for remote setting assignments)
- Admin users (user authentication; admin user receives default privilege assignment of read-only)

At present, a SecurID ACE server cannot assign L2TP or XAuth remote settings or NetScreen admin privileges, although you can use a SecurID server to store L2TP, XAuth, and admin user accounts for authentication purposes. Also, NetScreen does not provide user group support when used with SecurID.

LDAP

Lightweight Directory Access Protocol (LDAP) is a directory server standard developed at the University of Michigan in 1996. LDAP is a protocol for organizing and accessing information in a hierarchical structure resembling a branching tree. Its purpose is twofold:

- To locate resources, such as organizations, individuals, and files on a network
- To help authenticate users attempting to connect to networks controlled by directory servers

The basic LDAP structure branches from countries to organizations to organizational units to individuals. There can also be other, intermediary levels of branching, such as “states” and “counties”. The following illustration shows an example of the branching organizational structure of LDAP.

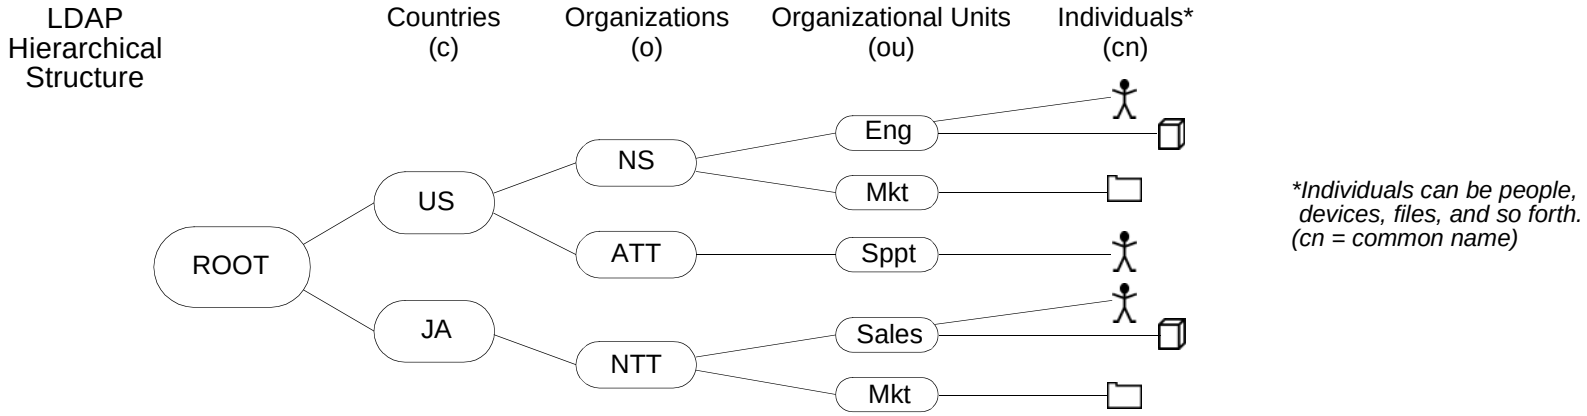

Note: *For information about LDAP, see RFC-1777 “Lightweight Directory Access Protocol”.*

You can configure the NetScreen device to link to a Lightweight Directory Access Protocol (LDAP) server. This server uses the LDAP hierarchical syntax to identify each user uniquely.

LDAP Auth Server Object Properties

In addition to the generic auth server properties listed in "Auth Server Object Properties" on page 377, an LDAP server also makes use of the following properties:

- **LDAP Server Port:** The port number on the LDAP server to which the NetScreen device sends authentication requests. The default port number is 389.

  Note: *If you change the LDAP port number on the NetScreen device, also change it on the LDAP server.*

- **Common Name Identifier:** The identifier used by the LDAP server to identify the individual entered in a LDAP server. For example, an entry of "uid" means "user ID" and "cn" for "common name".
- **Distinguished Name (dn):** The path used by the LDAP server before using the common name identifier to search for a specific entry. (For example, c=us;o=netscreen, where "c" stands for "country", and "o" for "organization".)

Supported User Types and Features

An LDAP server supports the following types of users and authentication features:

- Auth users
- L2TP users (user authentication; L2TP user receives default L2TP settings from the NetScreen device)
- XAuth users (user authentication; no support for remote setting assignments)
- Admin users (user authentication; admin user receives default privilege assignment of read-only)

At present, an LDAP server cannot assign L2TP or XAuth remote settings or NetScreen admin privileges, although you can use an LDAP server to store L2TP, XAuth, and admin user accounts for authentication purposes. Also, NetScreen does not provide user group support when used with LDAP.

Defining Auth Server Objects

Before you can refer to external authentication servers (auth servers) in policies, IKE gateways, and L2TP tunnels, you must first define the auth server objects. The following examples illustrate how to define auth server objects for a RADIUS server, a SecurID server, and an LDAP server.

Example: Defining an Auth Server Object for RADIUS

In the following example, you define an auth server object for a RADIUS server. You specify its user account types as auth, L2TP, and XAuth. You name the RADIUS server “radius1” and accept the ID number that the NetScreen device automatically assigns it. You enter its IP address, which is 10.20.1.100; and change its port number from the default (1645) to 4500. You define its shared secret as “A56htYY97kl”. You change the authentication timeout value from the default (10 minutes) to 30 minutes and the RADIUS retry timeout from 3 seconds to 4 seconds. You also assign its two backup servers the IP addresses 10.20.1.110 and 10.20.1.120.

In addition, you load the NetScreen dictionary file on the RADIUS server so that it can support queries for the following vendor-specific attributes (VSAs): user groups, admin privileges, remote L2TP and XAuth settings.

The NetScreen device sends auth, L2TP, and XAuth authentication requests to the primary RADIUS server, “radius1”, at 10.20.1.100.

If the NetScreen device loses network connectivity with the primary RADIUS server, it then redirects authentication requests to backup1 at 10.20.1.110.

If the NetScreen device cannot reach backup1, it then redirects authentication requests to backup2 at 10.20.1.120.

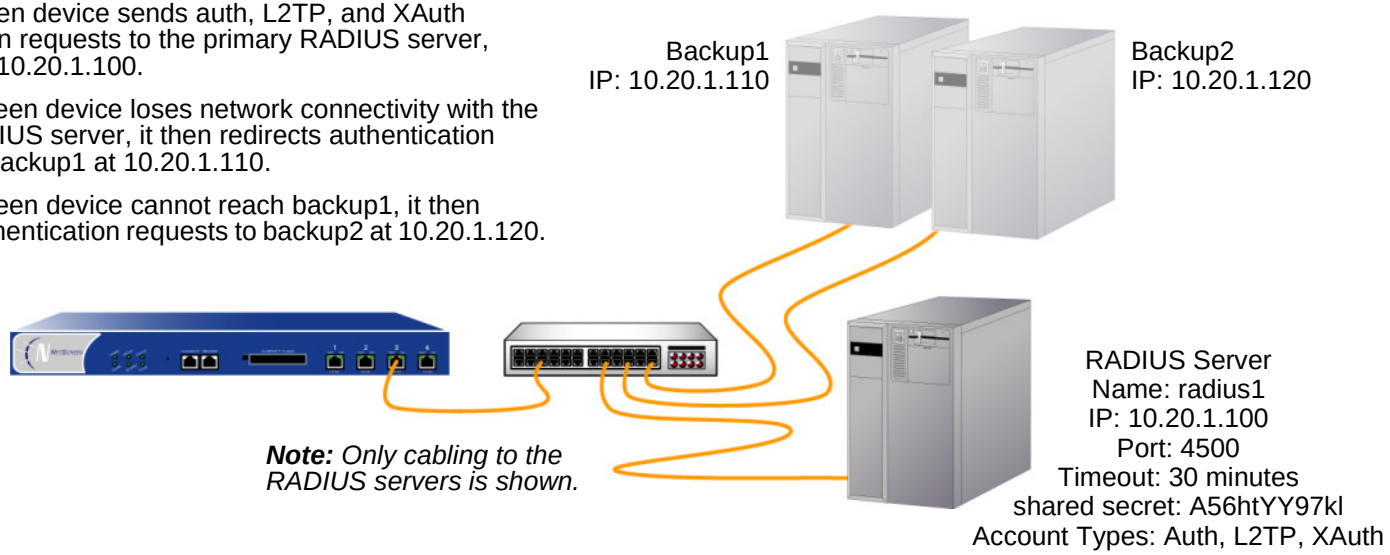

WebUI

Configuration > Auth > Servers > New: Enter the following, and then click **OK**:

Name: radius1

IP/Domain Name: 10.20.1.100

Backup1: 10.20.1.110

Backup2: 10.20.1.120

Timeout: 30

Account Type: Auth, L2TP, Xauth

RADIUS: (select)

RADIUS Port: 4500

Retry Timeout: 4 (seconds)

Shared Secret: A56htYY97kl

Load the NetScreen dictionary file on the RADIUS server.

Note: *For more information on the NetScreen dictionary file, see* *"NetScreen Dictionary File" on page 381**. For instructions on how to load the dictionary file onto a RADIUS server, refer to the documentation for your specific server.*

CLI

```
set auth-server radius1 type radius
set auth-server radius1 account-type auth l2tp xauth[2]
set auth-server radius1 server-name 10.20.1.100
set auth-server radius1 backup1 10.20.1.110
set auth-server radius1 backup2 10.20.1.120
set auth-server radius1 timeout 30
set auth-server radius1 radius port 4500[3]
set auth-server radius1 radius timeout 4
set auth-server radius1 radius secret A56htYY97kl
save
```

Load the NetScreen dictionary file on the RADIUS server.

Note: *For more information on the NetScreen dictionary file, see* *"NetScreen Dictionary File" on page 381**. For instructions on how to load the dictionary file onto a RADIUS server, refer to the documentation for your specific server.*

2. The order in which you enter the account types is important. For example, if you first type **set auth-server radius1 account-type l2tp**, then your only subsequent choice is **xauth**; you cannot type **auth** after **l2tp**. The correct order is easily remembered because it is alphabetical.
3. Changing the port number helps deter potential attacks targeted at the default RADIUS port number (1645).

Example: Defining an Auth Server Object for SecurID

In the following example, you configure an auth server object for a SecurID ACE server. You specify its user account type as admin. You name the server “securid1” and accept the ID number that the NetScreen device automatically assigns it. You enter the IP address of the primary server, which is 10.20.2.100, and the IP address of a backup server: 10.20.2.110. You change its port number from the default (5500) to 15000. The NetScreen device and the SecurID ACE server protect the authentication information using DES encryption. There are three allowable retries, and a client timeout value of 10 seconds[4]. You change the idle timeout value from the default (10 minutes) to 60 minutes[5]. The **Use Duress** setting is disabled.

The NetScreen device sends admin authentication requests to the primary SecurID server, “securid1”, at 10.20.2.100.

If the NetScreen device loses network connectivity with the primary SecurID server, it then redirects authentication requests to backup1 at 10.20.2.110.

Note: *NetScreen supports only one backup server for SecurID.*

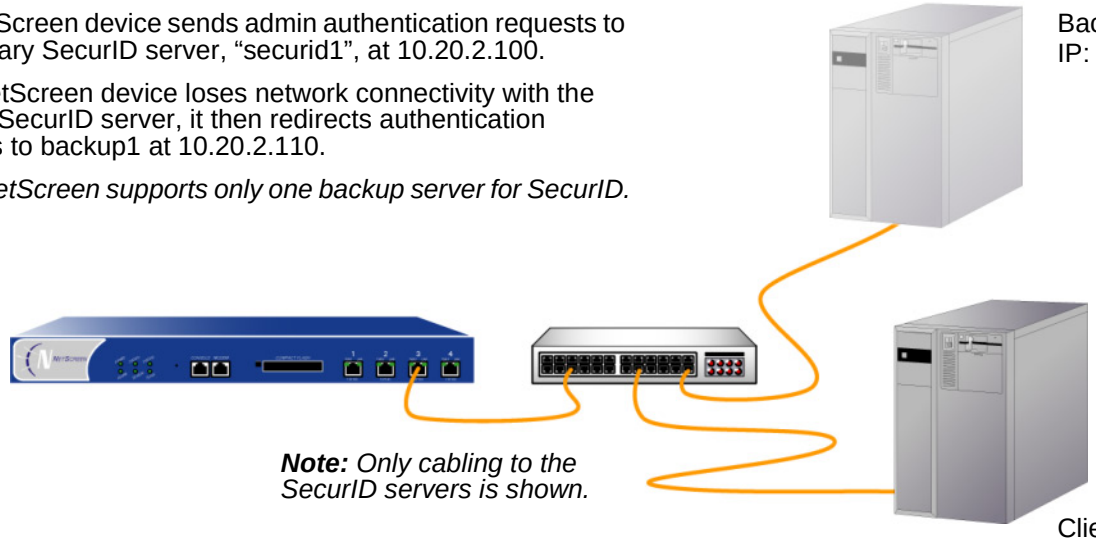

4. The client timeout value is the length of time in seconds that the SecurID client (that is, the NetScreen device) waits between authentication retry attempts.
5. The idle timeout value is the length of idle time in minutes that can elapse before the NetScreen device automatically terminates an inactive admin session. (For information comparing the timeout value as applied to admin users and other user types, see “Auth Server Object Properties” on page 377.)

WebUI

Configuration > Auth > Servers > New: Enter the following, and then click **OK**:

Name: securid1

IP/Domain Name: 10.20.2.100

Backup1: 10.20.2.110

Timeout: 60

Account Type: Admin

SecurID: (select)

Client Retries: 3

Client Timeout: 10 seconds

Authentication Port: 15000

Encryption Type: DES

User Duress: No

CLI

```
set auth-server securid1 type securid
set auth-server securid1 server-name 10.20.2.100
set auth-server securid1 backup1 10.20.2.110
set auth-server securid1 timeout 60
set auth-server securid1 account-type admin
set auth-server securid1 securid retries 3
set auth-server securid1 securid timeout 10
set auth-server securid1 securid auth-port 15000
set auth-server securid1 securid encr 1
set auth-server securid1 securid duress 0
save
```

Example: Defining an Auth Server Object for LDAP

In the following example, you configure an auth server object for an LDAP server. You specify its user account type as auth. You name the LDAP server “ldap1” and accept the ID number that the NetScreen device automatically assigns it. You enter its IP address, which is 10.20.3.100; and change its port number from the default (389) to 19830. You change the timeout value from the default (10 minutes) to 40 minutes. You also assign its two backup servers the IP addresses 10.20.3.110 and 10.20.3.120. The LDAP common name identifier is cn, and the Distinguished Name is c=us;o=netscreen;ou=marketing.

The NetScreen device sends authentication requests from auth users to the primary LDAP server, “ldap1”, at 10.20.3.100.

If the NetScreen device loses network connectivity with the primary LDAP server, it then redirects authentication requests to backup1 at 10.20.3.110.

If the NetScreen device cannot reach backup1, it then redirects authentication requests to backup2 at 10.20.3.120.

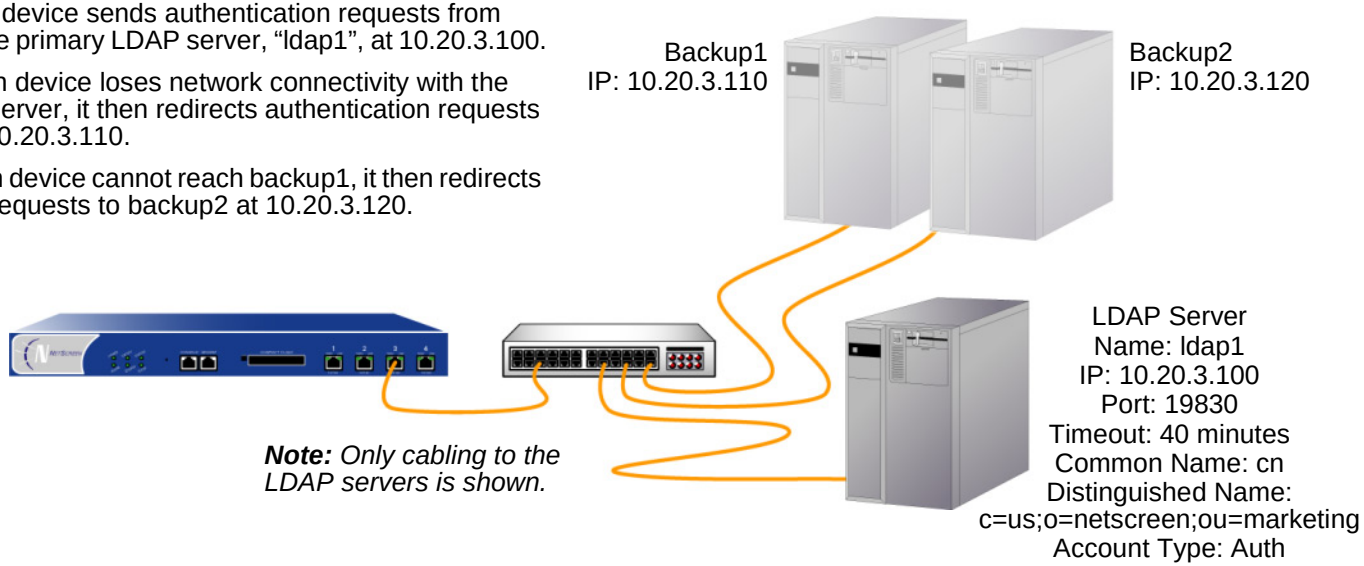

WebUI

Configuration > Auth > Servers > New: Enter the following, and then click **OK**:

Name: ldap1

IP/Domain Name: 10.20.3.100

Backup1: 10.20.3.110

Backup2: 10.20.3.120

Timeout: 40

Account Type: Auth

LDAP: (select)

LDAP Port: 4500

Common Name Identifier: cn

Distinguished Name (dn): c=us;o=netscreen;ou=marketing

CLI

```
set auth-server ldap1 type ldap
set auth-server ldap1 account-type auth
set auth-server ldap1 server-name 10.20.3.100
set auth-server ldap1 backup1 10.20.3.110
set auth-server ldap1 backup2 10.20.3.120
set auth-server ldap1 timeout 40
set auth-server ldap1 ldap port 15000
set auth-server ldap1 ldap cn cn
set auth-server ldap1 ldap dn c=us;o=netscreen;ou=marketing
save
```

Defining Default Auth Servers

By default, the local database is the default auth server for all user types. You can specify external auth servers to be the default auth servers for one or more of the following user types:

- Admin
- Auth
- L2TP
- XAuth

Then, if you want to use the default auth server for a specific user type when configuring authentication in policies, L2TP tunnels, or IKE gateways, you do not have to specify an auth server in every configuration. The NetScreen device refers to the appropriate auth servers that you previously appointed to be the defaults.

Example: Changing the Default Auth Servers

In this example, you use the RADIUS, SecurID, and LDAP auth server objects that you created in the previous examples:

- radius1 ("Example: Defining an Auth Server Object for RADIUS" on page 388)
- securid1 ("Example: Defining an Auth Server Object for SecurID" on page 391)
- ldap1 ("Example: Defining an Auth Server Object for LDAP" on page 393)

You then assign the local database, radius1, securid1, and ldap1 as the default servers for the following user types:

- radius1: Default auth server for admin users
- securid1: Default auth server for L2TP users
- ldap1: Default auth server for auth users
- Local: Default auth server for XAuth users[6]

6. By default, the local database is the default auth server for all user types. Therefore, unless you have previously assigned an external auth server as the default server for XAuth users, you do not need to configure it as such.

WebUI

Configuration > Admin > Administrators: Select **Local/radius1** from the Admin Auth Server drop-down list, and then click **Apply**.

VPNs > AutoKey Advanced > XAuth Settings: Select **Local** from the Default Authentication Server drop-down list and then click **Apply**[7].

Note: *You cannot set and defer to a default auth server in the WebUI for auth user authentication in a policy or XAuth user authentication in an IKE gateway. You must select an auth server from a drop-down list in each policy and in each IKE gateway configuration to which you want to apply user authentication.*

CLI

```
set admin auth server radius1
set auth default auth server ldap1
set l2tp default auth server securid1
set xauth default auth server Local[7]
save
```

7. By default, the local database is the default auth server for all user types. Therefore, unless you have previously assigned an external auth server as the default server for XAuth users, you do not need to configure it as such.

AUTHENTICATION TYPES AND APPLICATIONS

The following sections describe the different types of users and user groups that you can create, and how to use them when configuring policies, IKE gateways, and L2TP tunnels:

- "Auth Users and User Groups" on page 398
- "IKE Users and User Groups" on page 431
- "XAuth Users and User Groups" on page 436
- "L2TP Users and User Groups" on page 460
- "Admin Users" on page 465

The NetScreen device authenticates the different types of users at different stages in the connection process. To get a sense of when IKE, XAuth, L2TP, and auth authentication techniques occur during the creation of an L2TP-over-IPSec VPN tunnel, refer to the following illustration:

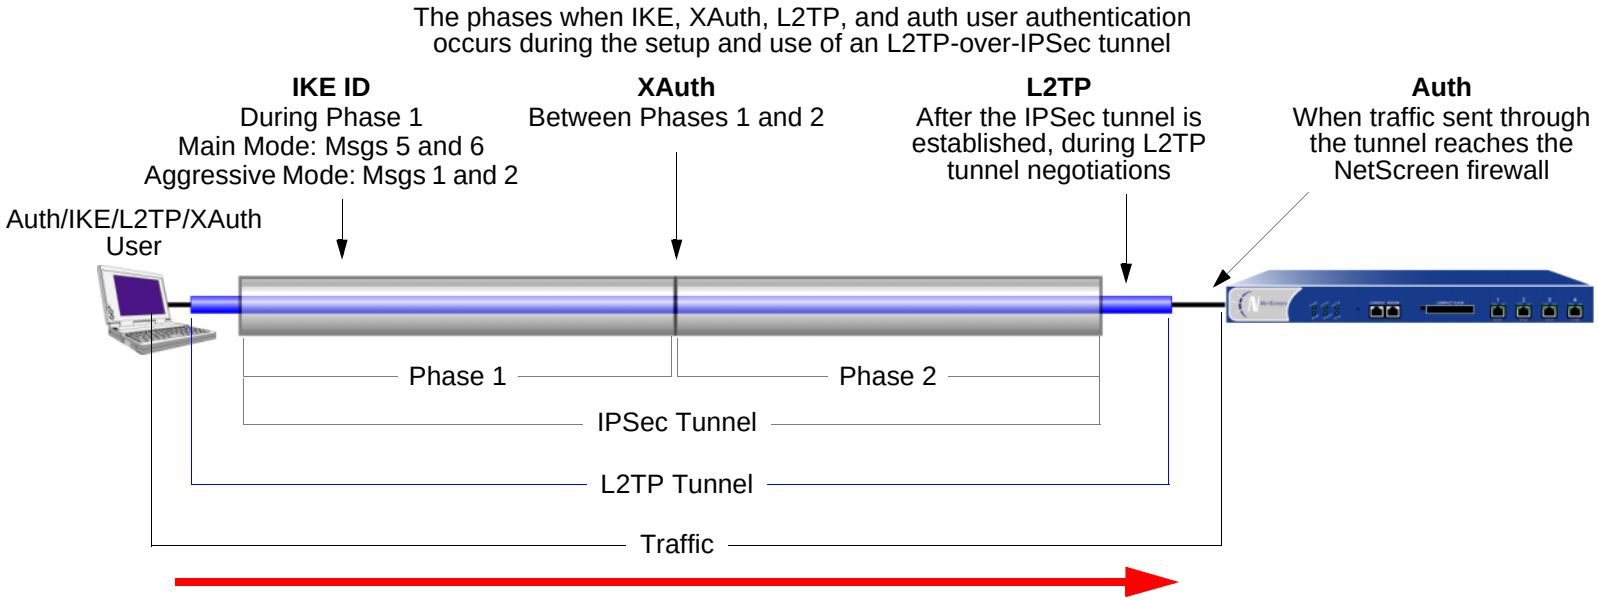

Note: Because XAuth and L2TP both provide user authentication and address assignments, they are seldom used together. They are shown together here solely to illustrate when each type of authentication occurs during the creation of a VPN tunnel.

Auth Users and User Groups

An auth user is a network user that must provide a user name and password for authentication when initiating a connection across the firewall. You can store an auth user account on the local database or on an external RADIUS, SecurID, or LDAP server.

You can put several auth user accounts together to form an auth user group, which you can store on the local database or on a RADIUS server. A single auth user account can be in up to four user groups on the local database or on a RADIUS server. If you create an external user group on a RADIUS server, you must also create an identical—but unpopulated—user group on the NetScreen device. For example, if you define an auth user group named “au_grp1” on a RADIUS server named “rs1” and add 10 members to the group, then on the NetScreen device you need to define an auth user group also named “au_grp1”, identify it as an external user group, but add no members to it. When you reference the external auth user group “au_grp1” and auth server “rs1” in a policy, the NetScreen device can properly query the specified RADIUS server when traffic matching the policy provokes an authentication check.

Referencing Auth Users in Policies

After you define an auth user, you can then create a policy that requires the user to authenticate himself or herself through one of two authentication schemes. The first scheme authenticates users when FTP, HTTP, or Telnet traffic matching a policy requiring authentication reaches the NetScreen device. In the second scheme, users authenticate themselves before sending traffic (of any kind—not just FTP, HTTP, or Telnet) to which a policy requiring user authentication applies.

Run-Time Authentication

When a user attempts to initiate an HTTP, FTP, or Telnet connection request to which a policy requiring authentication applies, the NetScreen device intercepts the request and prompts the user to enter a name and password (see “User Authentication” on page 210). Before granting permission, the NetScreen device validates the user name and password by checking them against those stored in the local database or on an external auth server.

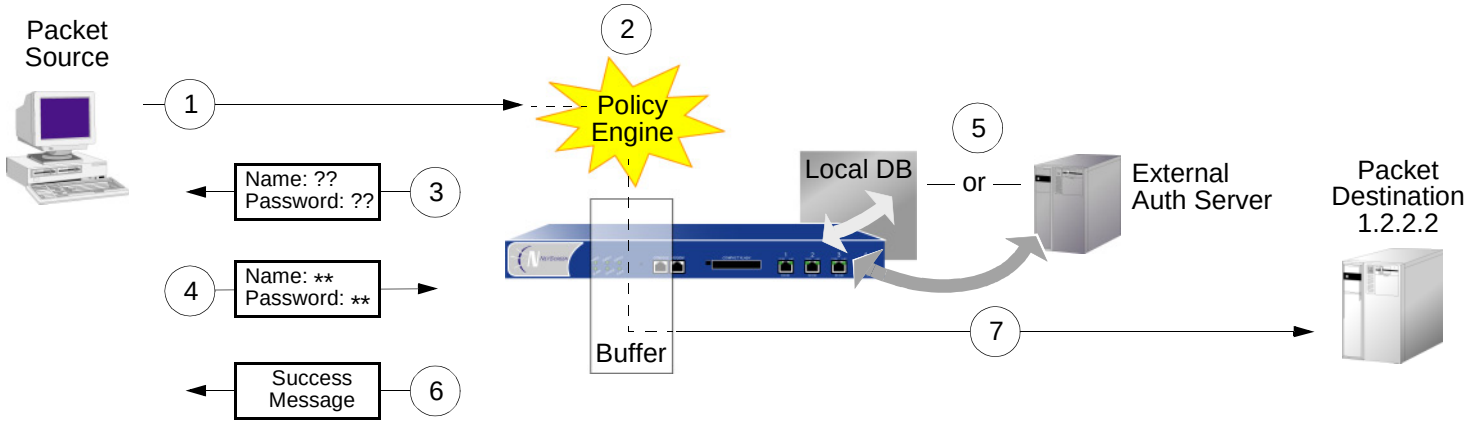

1. An auth user sends an FTP, HTTP, or Telnet packet to 1.2.2.2.
2. The NetScreen device intercepts the packet, notes that its policy requires authentication from either the local database or an external auth server and buffers the packet.
3. The NetScreen device prompts the user for login information via FTP, HTTP, or Telnet.
4. The user replies with a user name and password.
5. The NetScreen device either checks for an auth user account on its local database or it sends the login information to the external auth server as specified in the policy.
6. Finding a valid match (or receiving notice of such a match from the external auth server), the NetScreen device informs the user that the login has been successful.
7. The NetScreen device forwards the packet from its buffer to its destination of 1.2.2.2.

Pre-Policy Check Authentication (WebAuth)

Before sending traffic to an intended destination, an auth user initiates an HTTP session to the IP address hosting the WebAuth feature on the NetScreen device and authenticates himself or herself. After the NetScreen device authenticates the user, he or she can then send traffic to the destination as permitted by a policy requiring authentication via WebAuth. (For more information, see "Auth Users and User Groups" on page 398.)

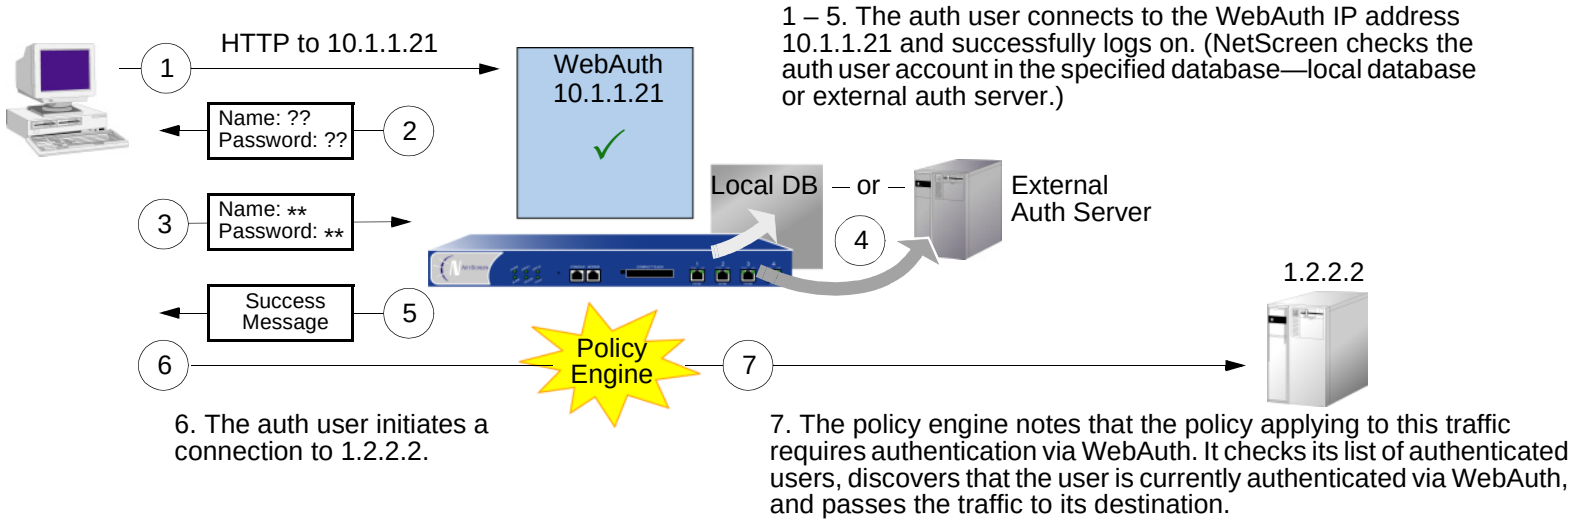

Some details about WebAuth:

- You can leave the default WebAuth auth server as the local database or you can choose an external auth server for the role. The main requirement for a WebAuth auth server is that the auth server must have auth user account-types.
- The WebAuth address must be in the same subnet as the interface that you want to use to host it. For example, if you want auth users to connect to WebAuth via ethernet3, which has IP address 1.1.1.1/24, then you can assign WebAuth an IP address in the 1.1.1.0/24 subnet.
- You can put a WebAuth address in the same subnet as the IP address of any physical interface, subinterface, or virtual security interface (VSI). (For information about different types of interfaces, see "Interfaces" on page 65.)

- If you want to use WebAuth while in Transparent mode, you can put a WebAuth address in the same subnet as the VLAN1 IP address.
- You can put WebAuth addresses on multiple interfaces.
- If you have multiple interfaces bound to the same security zone, you can put a WebAuth address in a subnet on one interface, and traffic from the same zone but using a different interface can still reach it.
- Be aware that after a NetScreen device authenticates a user at a particular source IP address, it subsequently permits traffic—as specified in the policy requiring authentication via WebAuth—from any other user at that same address. This might be the case if the users originate traffic from behind a NAT device that changes all original source addresses to a single translated address.

Referencing Auth User Groups in Policies

To manage a number of auth users, you can create auth user groups and store them either on the local NetScreen device or on an external RADIUS server.

Note: *If you store users in groups on a RADIUS server, you must create unpopulated external user groups on the NetScreen device with names that correspond with those of the user groups you create on the RADIUS server.*

Rather than manage each user individually, you can gather users into a group, so that any changes made to the group propagate to each group member. An auth user can be a member of up to four user groups on the local database or on a RADIUS server. An auth user who belongs to more than one group is required to supply a username and password only once, before being granted access to the resources defined for each group in which the user is a member.

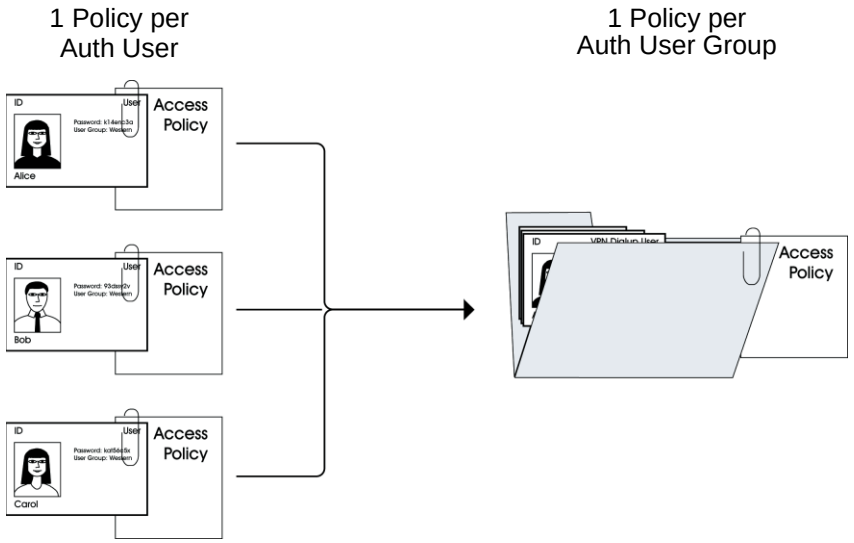

Example: Run-Time Authentication (Local User)

In this example, you define a local auth user named louis with password iDa84rNk, and an address named “host1” in the Trust zone address book. You then configure two outgoing policies: one that denies all outbound traffic, and another from host1 requiring louis to authenticate himself. (Louis must initiate all outbound traffic from host1.) The NetScreen device denies outbound access from any other address, as well as unauthenticated traffic from “host1”.

WebUI

1. Local Auth User and Address

Objects > Users > Local > New: Enter the following, and then click **OK**:

User Name: louis

Status: Enable

Authentication User: (select)

User Password: iDa84rNk

Confirm Password: iDa84rNk

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: host1

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.4/32

Zone: Trust

2. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Any

Service: ANY

Action: Deny

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), host1

Destination Address:

Address Book Entry: (select), Any

Service: ANY

Action: Permit

Position at Top: (select)

> Advanced**:** Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

Authentication: (select)

Auth Server: (select)

Use: Local

User: (select), Local Auth User - louis

CLI

1. Local User and Address

```
set user louis password iDa84rNk[8]
set address trust host1 10.1.1.4/32
```

2. Policies

```
set policy from trust to untrust any any any deny
set policy top from trust to untrust host1 any any permit auth user louis
save
```

8. By default, a user to whom you assign a password is classified as an auth user.

Example: Run-Time Authentication (Local User Group)

In this example, you define a local user group named auth_grp1. You add previously created auth-users louis and lara to the group[9]. Then you configure a policy referencing auth_grp1. The policy provides FTP-GET and FTP-PUT privileges for auth_grp1, with address name "auth_grp1" (IP address 10.1.8.0/24) in the Trust zone to access an FTP server named "ftp1" (IP address 1.2.2.3/32) in the DMZ zone.

WebUI

1. Local User Group and Members

Objects > Users > Local Groups > New: Enter **auth_grp1** in the Group Name field, do the following, and then click **OK**:

Select **louis** and use the **<<** button to move him from the Available Members column to the Group Members column.

Select **lara** and use the **<<** button to move her from the Available Members column to the Group Members column.

2. Address

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: auth_grp1

IP Address/Domain Name:

IP/Netmask: (select), 10.1.8.0/24

Zone: Trust

9. When you create a user group in the local database, its user type remains undefined until you add a user to it. At that point, the user group takes the type or types of users that you add to it. You can create a multiple-type user group by adding auth, IKE, L2TP, and XAuth user types. You cannot combine Admin users with any other user type.

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: ftp1

IP Address/Domain Name:

IP/Netmask: (select), 1.2.2.3/32

Zone: DMZ

3. Policy

Policies > (From: Trust; To: DMZ) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), auth_grp1

Destination Address:

Address Book Entry: (select), ftp1

Service: FTP

Action: Permit

Position at Top: (select)

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

Authentication: (select)

Auth Server: (select)

Use: Local

User Group: (select), Local Auth Group - auth_grp1

CLI

1. Local User Group and Members

```
set user-group auth_grp1 location local
set user-group auth_grp1 user louis
set user-group auth_grp1 user lara
```

2. Address

```
set address trust auth_grp1 10.1.8.0/24
set address dmz ftp1 1.2.2.3/32
```

3. Policy

```
set policy top from trust to dmz auth_grp1 ftp1 ftp permit auth user-group
    auth_grp1
save
```

Example: Run-Time Authentication (External User)

In this example, you define an external LDAP auth server named “x_srv1” with the following attributes:

- Account type: auth
- IP address: 10.1.1.100
- Backup1 IP address: 10.1.1.110
- Backup2 IP address: 10.1.1.120
- Authentication timeout: 60 minutes
- LDAP port number: 14500
- Common name identifier: cn
- Distinguished name: c=us;o=netscreen

You load the auth user “euclid” with password eTcS114u on the external auth server. You then configure an outgoing policy that requires authentication on auth server x_srv1 for external user euclid.

WebUI

1. Auth Server

Configuration > Auth > Servers > New: Enter the following, and then click **OK**:

Name: x_srv1

IP/Domain Name: 10.1.1.100

Backup1: 10.1.1.110

Backup2: 10.1.1.120

Timeout: 60

Account Type: Auth

LDAP: (select)

LDAP Port: 14500

Common Name Identifier: cn

Distinguished Name (dn): c=us;o=netscreen

2. External User

Define the auth user “euclid” with password eTcS114u on the external LDAP auth server x_serv1.

3. Address

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: euc_host

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.20/32

Zone: Trust

4. Policy

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Any

Service: ANY

Action: Permit

Position at Top: (select)

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

Authentication: (select)

Auth Server: (select)

Use: x_srv1

User: (select), External User

External User: euclid

CLI

1. **Auth Server**

```
set auth-server x_srv1
set auth-server x-srv1 type ldap
set auth-server xsrv1 account-type auth
set auth-server x_srv1 server-name 10.1.1.100
set auth-server lx_srv1 backup1 10.1.1.110
set auth-server x_srv1 backup2 10.1.1.120
set auth-server x_srv1 timeout 60
set auth-server x_srv1 ldap port 14500
set auth-server x_srv1 ldap cn cn
set auth-server x_srv1 ldap dn c=us;o=netscreen
```

2. **External User**

Define the auth user “euclid” with password eTcS114u on the external LDAP auth server x_serv1.

3. **Address**

```
set address trust euc_host 10.1.1.20/32
```

4. **Policy**

```
set policy top from trust to untrust euc_host any any auth server x_srv1 user
    euclid
save
```

Example: Run-Time Authentication (External User Group)

In this example, you configure an external RADIUS auth server named “radius1”[10] and define an external auth user group named “auth_grp2”. You define the external auth user group auth_grp2 in two places:

1. External RADIUS auth server “radius1”
2. NetScreen device

You populate the auth user group “auth_grp2” with auth users on the RADIUS server only, leaving the group unpopulated on the NetScreen device. The members in this group are accountants who require exclusive access to a server at IP address 10.1.1.80. You create an address book entry for the server and name the address “midas.” You then configure an intrazone policy permitting only authenticated traffic from auth_grp2 to midas, both of which are in the Trust zone. (For more information on intrazone policies, see Chapter 7, “Policies”.)

RADIUS Server

1. Load the NetScreen dictionary file on the RADIUS server[11].

   Note: *For information on the NetScreen dictionary file, see “NetScreen Dictionary File” on page 381. For instructions on loading the dictionary file onto a RADIUS server, refer to the RADIUS server documentation.*

2. After you define auth user accounts on the RADIUS server, use the NetScreen user group VSA to create the user group “auth_grp2” and apply it to the auth user accounts that you want to add to that group.

10. The RADIUS auth server configuration is nearly identical to that in “Example: Defining an Auth Server Object for RADIUS” on page 388, except that in this example you only specify “auth” as the user account type.

11. If you are using a Microsoft IAS RADIUS server, there is no dictionary file to load. Instead, define the correct vendor-specific attributes (VSAs) on the server.

WebUI

1. Auth Server

Configuration > Auth > Servers > New: Enter the following, and then click **OK**:

Name: radius1

IP/Domain Name: 10.20.1.100

Backup1: 10.20.1.110

Backup2: 10.20.1.120

Timeout: 30

Account Type: Auth

RADIUS: (select)

RADIUS Port: 4500

Shared Secret: A56htYY97kl

2. External User Group

Objects > Users > External Groups > New: Enter the following, and then click **OK**:

Group Name: auth_grp2

Group Type: Auth

3. Address

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: midas

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.80/32

Zone: Trust

4. Policy

Policies > (From: Trust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), midas

Service: ANY

Action: Permit

Position at Top: (select)

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

Authentication: (select)

Auth Server: (select)

Use: radius1

User Group: (select), External Auth Group - auth_grp2

CLI

1. Auth-Server

```
set auth-server radius1 type radius
set auth-server radius1 account-type auth
set auth-server radius1 server-name 10.20.1.100
set auth-server radius1 backup1 10.20.1.110
set auth-server radius1 backup2 10.20.1.120
set auth-server radius1 timeout 30
set auth-server radius1 radius port 4500
set auth-server radius1 radius secret A56htYY97kl
```

2. External User Group

```
set user-group auth_grp2 location external
set user-group auth_grp2 type auth
```

3. Address

```
set address trust midas 10.1.1.80/32
```

4. Policy

```
set policy top from trust to trust any midas any permit auth server radius1
    user-group auth_grp2
save
```

Example: Local Auth User in Multiple Groups

In this example, you define a local auth user named Mary. Mary is a sales manager who needs access to two servers: server A, which is for the salespeople (sales_reps group), and server B, which is for the managers (sales_mgrs group). To provide access to both, you add Mary to the two user groups. You then create two policies—one for each group.

Note: *This example does not show the configuration for the other group members.*

WebUI

1. Local User

Objects > Users > Local > New: Enter the following, and then click **OK**:

User Name: mary

Status: Enable

Authentication User: (select)

User Password: iFa8rBd

Confirm Password: iFa8rBd

2. Local User Groups and Member

Objects > Users > Local Groups > New: Enter **sales_mgrs** in the Group Name field, do the following, and then click **OK**:

Select **mary** and use the **<<** button to move her from the Available Members column to the Group Members column.

Objects > Users > Local Groups > New: Enter **sales_reps** in the Group Name field, do the following, and then click **OK**:

Select **mary** and use the **<<** button to move her from the Available Members column to the Group Members column.

3. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: sales
IP Address/Domain Name:
IP/Netmask: (select), 10.1.8.0/24
Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: server_a
IP Address/Domain Name:
IP/Netmask: (select), 1.1.1.5/32
Zone: Untrust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: server_b
IP Address/Domain Name:
IP/Netmask: (select), 1.1.1.6/32
Zone: Untrust

4. Policies

Policies > (From: Trust; To: Untrust) > New: Enter the following, and then click **OK**:

Source Address:
Address Book Entry: (select), sales
Destination Address:
Address Book Entry: (select), server_a
Service: FTP
Action: Permit

Position at Top: (select)

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

Authentication: (select)

Auth Server: (select)

Use: Local

User Group: (select), Local Auth Group - sales_reps

Policies > (From: Trust; To: Untrust) > New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), sales

Destination Address:

Address Book Entry: (select), server_b

Service: FTP

Action: Permit

Position at Top: (select)

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

Authentication: (select)

Auth Server: (select)

Use: Local

User Group: (select), Local Auth Group - sales_mgrs

CLI

1. Local User

```
set user mary password iFa8rBd
```

2. Local User Groups and Member

```
set user-group sales_mgrs location local
set user-group sales_mgrs user mary
set user-group sales_reps location local
set user-group sales_reps user mary
```

3. Addresses

```
set address trust sales 10.1.8.0/24
set address untrust server_a 1.1.1.5/32
set address untrust server_b 1.1.1.6/32
```

4. Policy

```
set policy top from trust to untrust sales server_a ftp permit auth user-group
    sales_reps
set policy top from trust to untrust sales server_b ftp permit auth user-group
    sales_mgrs
save
```

Example: WebAuth (Local User Group)

In this example, you require users to preauthenticate themselves via the WebAuth method before initiating outbound traffic to the Internet. You create a user group named “auth_grp3” in the local database on the NetScreen device. You then create auth user accounts for everyone in the Trust zone and add them to “auth_grp3”.

The Trust zone interface uses ethernet1 and has IP address 10.1.1.1/24. You assign 10.1.1.50 as the WebAuth IP address, and you use keep the local database as the default WebAuth server. Consequently, before a user can initiate traffic to the Internet, he or she must first make an HTTP connection to 10.1.1.50 and log on with a user name and password. The NetScreen device then checks the user name and password against those in its database and either approves or rejects the authentication request. If it approves the request, the authenticated user has 30 minutes to initiate traffic to the Internet. After terminating that initial session, the user has another 30 minutes to initiate another session before the NetScreen device requires him or her to reauthenticate himself or herself.

WebUI

1. WebAuth

Configuration > Auth > WebAuth: Select **Local** from the WebAuth Server drop-down list, and then click **Apply**.

Network > Interfaces > Edit (for ethernet1): Select **WebAuth**, and in the WebAuth IP field enter **10.1.1.50**.

Configuration > Auth > Servers > Edit (for Local): Enter **30** in the Timeout field, and then click **Apply**.

2. User Group

Objects > Users > Local Groups > New: Enter **auth_grp3** in the Group Name field, do the following, and then click **OK**:

Select ***user name*** and use the **<<** button to move that user from the Available Members column to the Group Members column.

Repeat the selection process, adding auth users until the group is complete.

3. Policy

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Any

Service: ANY

Action: Permit

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

Authentication: (select)

WebAuth: (select)

User Group: (select), Local Auth Group - auth_grp3

CLI

1. WebAuth

```
set webauth server Local
set interface ethernet1 webauth-ip 10.1.1.50
set interface ethernet1 webauth
set auth-server Local timeout 30
```

2. User Group

```
set user-group auth_grp3 location local
```

> ***Note:*** *The NetScreen device determines a local user group type by the type of members that you add to it. To make auth_grp3 an auth user group, add an auth user to the group.*

Use the following command to add auth users to the user group you have just created:

```
set user-group auth_grp3 user name_str
```

3. Policy

```
set policy top from trust to untrust any any any permit webauth user-group
    auth_grp3
save
```

Example: WebAuth (External User Group)

WebAuth is a method for pre-authenticating users before they initiate traffic across the firewall. In this example, you create a policy requiring authentication via the WebAuth method for all outgoing traffic.

You create an auth user group named “auth_grp4” on both the RADIUS server “radius1” and on the NetScreen device. On the RADIUS server, you create user accounts for everyone in the Trust zone and add them to “auth_grp4”.

Note: *Nearly the same RADIUS server settings are used here as in “Example: Defining an Auth Server Object for RADIUS” on page 388, except that in this example you only specify “auth” as the user account type.*

The Trust zone interface uses ethernet1 and has IP address 10.1.1.1/24. You assign 10.1.1.50 as the WebAuth IP address, and you use the external RADIUS auth-server “radius1” as the default WebAuth server. Consequently, before a user can initiate traffic to the Internet, he or she must first make an HTTP connection to 10.1.1.50 and log on with a user name and password. The NetScreen device then relays all WebAuth user authentication requests and responses between “radius1” and the users attempting to log on.

RADIUS Server

1. Load the NetScreen dictionary file on the RADIUS server.

   Note: *For information on the NetScreen dictionary file, see “NetScreen Dictionary File” on page 381. For instructions on loading the dictionary file onto a RADIUS server, refer to the RADIUS server documentation.*

2. Enter user group “auth_grp4” on the auth-server “radius1“, and then populate it with auth user accounts.

WebUI

1. Auth-Server

   Configuration > Auth > Servers > New: Enter the following, and then click **OK**:

   Name: radius1

   IP/Domain Name: 10.20.1.100

   Backup1: 10.20.1.110

   Backup2: 10.20.1.120

   Timeout: 30

   Account Type: Auth

   RADIUS: (select)

   RADIUS Port: 4500

   Shared Secret: A56htYY97k

2. WebAuth

   Configuration > Auth > WebAuth: Select **radius1** from the WebAuth Server drop-down list, and then click **Apply**.

   Network > Interfaces > Edit (for ethernet1): Select **WebAuth**, in the WebAuth IP field enter **10.10.1.50**, and then click **OK**.

3. User Group

   Objects > Users > External Groups > New: Enter the following, and then click **OK**:

   Group Name: auth_grp4

   Group Type: Auth

4. Policy

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Any

Service: ANY

Action: Permit

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

Authentication: (select)

WebAuth: (select)

User Group: (select), External Auth Group - auth_grp4

CLI

1. Auth-Server

```
set auth-server radius1 type radius
set auth-server radius1 account-type auth
set auth-server radius1 server-name 10.20.1.100
set auth-server radius1 backup1 10.20.1.110
set auth-server radius1 backup2 10.20.1.120
set auth-server radius1 timeout 30
set auth-server radius1 radius port 4500
set auth-server radius1 radius secret A56htYY97kl
```

2. WebAuth

```
set webauth server radius1
set interface ethernet1 webauth-ip 10.1.1.50
set interface ethernet1 webauth
```

3. User Group

```
set user-group auth_grp4 location external
set user-group auth_grp4 type auth
```

4. Policy

```
set policy top from trust to untrust any any any permit webauth user-group
    auth_grp4
save
```

Example: WebAuth + SSL (External User Group)

In this example, you combine WebAuth with Secure Sockets Layer (SSL) technologies to provide security for the user names and passwords that users transmit when logging on. WebAuth makes use of the same certificate that secures administrative traffic to the NetScreen device for management via the WebUI. (For more information about SSL, see “Secure Sockets Layer” on page **3**-7.)

The configuration for WebAuth using an external auth server plus SSL involves the following steps:

- You define an external RADIUS auth-server “radius1” and create an auth user group named “auth_grp5” on both the RADIUS server and on the NetScreen device. On the RADIUS server, you create user accounts for all auth users in the Untrust zone and add them to “auth_grp5”.

  Note: *Nearly identical RADIUS server settings are used here as in “Example: Defining an Auth Server Object for RADIUS” on page 388, except that you only specify “auth” as the user account type here.*

- The Untrust zone interface uses ethernet3 and has IP address 1.1.1.1/24. You assign 1.1.1.50 as the WebAuth IP address, and you use the external RADIUS auth-server “radius1” as the default WebAuth server.
- You specify the following SSL settings:
  - IDX number (1 in this example) of a certificate that you have previously loaded on the NetScreen device[12]
  - DES_SHA-1 ciphers
  - SSL port number 2020
- You enable SSL manageability on ethernet3 so that it does not reject SSL connection attempts to that interface.
- You then configure an incoming policy requiring authentication via the WebAuth + SSL method for all traffic from the Untrust to Trust zones.

Consequently, before a user can initiate traffic to the internal network, he or she must first make an HTTPS connection to https://1.1.1.50:2020 and log on with a user name and password. The NetScreen device then relays all WebAuth user authentication requests and responses between “radius1” and the user attempting to log on.

12. For information on how to obtain and load digital certificates onto a NetScreen device, see “Public Key Cryptography” on page **5**-15.

RADIUS Server

1. Load the NetScreen dictionary file on the RADIUS server.

   Note: *For information on the NetScreen dictionary file, see "NetScreen Dictionary File" on page 381. For instructions on loading the dictionary file onto a RADIUS server, refer to the RADIUS server documentation.*

2. Enter user group "auth_grp5" on the auth-server "radius1", and then populate it with auth user accounts.

WebUI

1. Auth-Server

Configuration > Auth > Servers > New: Enter the following, and then click **OK**:

Name: radius1

IP/Domain Name: 10.20.1.100

Backup1: 10.20.1.110

Backup2: 10.20.1.120

Timeout: 30

Account Type: Auth

RADIUS: (select)

RADIUS Port: 4500

Shared Secret: A56htYY97k

2. WebAuth

Configuration > Auth > WebAuth: Select **radius1** from the WebAuth Server drop-down list, and then click **Apply**.

Network > Interfaces > Edit (for ethernet3): Select **WebAuth**, in the WebAuth IP field enter **1.1.1.50**, and then click **OK**.

3. SSL

Configuration > Admin > Management: Enter the following, and then click **OK**:

HTTPS (SSL) Port: 2020

Certificate: (select the certificate that you previously loaded)

Cipher: DES_SHA-1

Network > Interfaces > Edit (for ethernet3): Select **SSL** in the Management Services section, and then click **OK**.

4. User Group

Objects > Users > External Groups > New: Enter the following, and then click **OK**:

Group Name: auth_grp5

Group Type: Auth

5. Policy

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Any

Service: ANY

Action: Permit

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

Authentication: (select)

WebAuth: (select)

User Group: (select), External Auth Group - auth_grp5

CLI

1. Auth-Server

```
set auth-server radius1 type radius
set auth-server radius1 account-type auth
set auth-server radius1 server-name 10.20.1.100
set auth-server radius1 backup1 10.20.1.110
set auth-server radius1 backup2 10.20.1.120
set auth-server radius1 timeout 30
set auth-server radius1 radius port 4500
set auth-server radius1 radius secret A56htYY97kl
```

Load the NetScreen dictionary file on the RADIUS server.

Note: *For information on the NetScreen dictionary file, see "NetScreen Dictionary File" on page 381. For instructions on loading the dictionary file onto a RADIUS server, refer to the RADIUS server documentation.*

2. WebAuth

```
set webauth server radius1
set interface ethernet3 webauth-ip 1.1.1.50
set interface ethernet3 webauth
```

3. SSL

```
set ssl port 2020
set ssl cert 1
set ssl encrypt des sha-1
set ssl enable
```

4. User Group

```
set user-group auth_grp5 location external
set user-group auth_grp5 type auth
```

5. Policy

```
set policy top from untrust to trust any any any permit webauth user-group
    auth_grp5
save
```

IKE Users and User Groups

An IKE user is a remote VPN user with a dynamically assigned IP address. The user—actually, the user's device—authenticates itself by sending either a certificate or preshared key together with an IKE ID during Phase 1 negotiations with the NetScreen device.

The IKE ID can be an e-mail address, an IP address, a domain name, or ASN1-DN string[13]. A NetScreen device authenticates an IKE user if the user sends either of the following:

- A **certificate** in which one or more of the values that appear in the distinguished name (DN) fields or in the SubAltName field is the same as the user's IKE ID configured on the NetScreen device
- A **preshared key** and an **IKE ID**, and the NetScreen device can successfully generate an identical preshared key from the received IKE ID and a preshared key seed value stored on the NetScreen device

You reference an IKE user or user group in an AutoKey IKE gateway configuration. By gathering IKE users that require similar gateway and tunnel configurations into a group, you only need to define one gateway referencing the group (and one VPN tunnel referencing that gateway), instead of one gateway and tunnel for each IKE user.

It is often impractical to create separate user accounts for every host. In such cases, you can create an IKE user group that has only one member, referred to as a group IKE ID user. The IKE ID of that user contains a set of values that must be present in the dialup IKE users' IKE ID definitions. If the IKE ID of a remote dialup IKE user matches the IKE ID of the group IKE ID user, NetScreen authenticates that remote user. For more information, refer to "Group IKE ID" on page **5**-237.

Note: *You can only store IKE user and IKE user group accounts on the local database.*

13. An example of an IKE ID using the Abstract Syntax Notation, version 1, distinguished name (ASN1-DN) format is CN=joe,OU=it,O=netscreen,L=sunnyvale, ST=ca,C=us,E=joe@ns.com.

Example: Defining IKE Users

In this example, you define four IKE users, Amy, Basil, Clara, and Desmond, each with a different kind of IKE ID.

- Amy – e-mail address (user-fully qualified domain name or U-FQDN): amy@ns.com
- Basil – IP address: 3.3.1.1
- Clara – fully qualified domain name (FQDN): www.netscreen.com
- Desmond – ASN1-DN string: CN=des,OU=art,O=netscreen,L=sunnyvale,ST=ca,C=us,E=des@ns.com

WebUI

Objects > Users > Local > New: Enter the following, and then click **OK**:

User Name: Amy

Status: Enable

IKE User: (select)

Simple Identity: (select)

IKE ID Type: AUTO

IKE Identity : amy@ns.com

Objects > Users > Local > New: Enter the following, and then click **OK**:

User Name: Basil

Status: Enable

IKE User: (select)

Simple Identity: (select)

IKE ID Type: AUTO

IKE Identity : 3.3.1.1

Objects > Users > Local > New: Enter the following, and then click **OK**:

User Name: Clara

Status: Enable

IKE User: (select)

Simple Identity: (select)

IKE ID Type: AUTO

IKE Identity : www.netscreen.com

Objects > Users > Local > New: Enter the following, and then click **OK**:

User Name: Desmond

Status: Enable

IKE User: (select)

Use Distinguished Name for ID: (select)

CN: des

OU: art

Organization: netscreen

Location: sunnyvale

State: ca

Country: us

E-mail: des@ns.com

CLI

```
set user Amy ike-id u-fqdn amy@ns.com
set user Basil ike-id ip 3.3.1.1
set user Clara ike-id fqdn www.netscreen.com
set user Desmond ike-id wildcard
    CN=des,OU=art,O=netscreen,L=sunnyvale,ST=ca,C=us,E=des@ns.com
save
```

Example: Creating an IKE User Group

In this example, you create a user group named ike_grp1. It becomes an IKE user group when you add IKE user Amy to it. You then add the other three IKE uses that you defined in the previous example, "Example: Defining IKE Users" on page 432.

WebUI

Objects > Users > Local Groups > New: Enter **ike_grp1** in the Group Name field, do the following, and then click **OK**:

Select **Amy** and use the **<<** button to move her from the Available Members column to the Group Members column.

Select **Basil** and use the **<<** button to move him from the Available Members column to the Group Members column.

Select **Clara** and use the **<<** button to move her from the Available Members column to the Group Members column.

Select **Desmond** and use the **<<** button to move him from the Available Members column to the Group Members column.

CLI

```
set user-group ike_grp1 location local
set user-group ike_grp1 user amy
set user-group ike_grp1 user basil
set user-group ike_grp1 user clara
set user-group ike_grp1 user desmond
save
```

Referencing IKE Users in Gateways

After you define an IKE user or IKE user group, you can then reference it in an IKE gateway configuration when the remote IKE gateway is a dialup user or dialup user group.

To see examples that reference IKE users in gateway configurations, see the following examples:

- “Example: Policy-Based Dialup VPN, AutoKey IKE” on page **5**-201
- “Example: Group IKE ID (Certificates)” on page **5**-243
- “Example: Group IKE ID (Preshared Keys)” on page **5**-252

XAuth Users and User Groups

The XAuth protocol is composed of two components: remote VPN user authentication (user name plus password) and TCP/IP address assignments (IP address, netmask[14], DNS server, and WINS server assignments). NetScreen supports the application of either component by itself or both components in concert.

An XAuth user or user group is one or more remote users who authenticate themselves when connecting to the NetScreen device via an AutoKey IKE VPN tunnel and optionally receive TCP/IP settings from the NetScreen device. Whereas the authentication of IKE users is actually the authentication of VPN gateways or clients, the authentication of XAuth users is the authentication of the individuals themselves. XAuth users must enter information that only they are supposed to know—their user name and password.

The NetScreen-Remote client can use the TCP/IP settings it receives to create a virtual adapter[15] when sending VPN traffic—while using the TCP/IP network adapter settings provided by the ISP or network admin for non-VPN traffic. By assigning known IP addresses to remote users, you can define routes on the NetScreen device to those addresses via specific tunnel interfaces. Then the NetScreen device can ensure that return routing reaches the remote user's IP address through the VPN tunnel, not via the default gateway. Address assignments also allow a downstream firewall to reference those addresses when creating policies. You can control the length of time that an IP address is associated with an individual XAuth user with the XAuth lifetime setting.

14. The assigned netmask is always 255.255.255.255 and cannot be modified.

15. A virtual adapter is the TCP/IP settings (IP address, DNS server addresses, WINS server addresses) that the NetScreen device assigns to a remote user for the duration of a VPN tunnel connection. Only NetScreen-Remote clients support virtual adapter functionality. NetScreen platforms do not.

ScreenOS supports the following aspects of XAuth:

- Authentication of local XAuth users and external XAuth users
- Authentication of local XAuth user groups and external XAuth user groups if stored on a RADIUS auth server
- IP, DNS server, and WINS server address assignments from an IP address pool for local XAuth users and external XAuth users stored on a RADIUS auth server

  To configure the NetScreen device to use default XAuth settings stored on an external RADIUS server, do either of the following:

  – WebUI: On the VPNs > AutoKey Advanced > XAuth Settings page, select **Query Client Settings on Default Server**.

  – CLI: Enter the **set xauth default auth server** *name_str* **query-config** command.

  The NetScreen device can also use gateway-specific XAuth settings stored on an external RADIUS server. When configuring a specific IKE gateway, do either of the following:

  – WebUI: On the VPNs > AutoKey Advanced > Gateway > New > Advanced page, select the name of the RADIUS server from the External Authentication drop-down list, and select **Query Remote Setting**.

  – CLI: Enter the **set ike gateway** *name_str* **xauth server** *name_str* **query-config** command.

XAuth Users in IKE Negotiations

NetScreen supports XAuth, version 6 (v6). To confirm that both parties in Phase 1 IKE negotiations support XAuth v6, they each send the following vendor ID to each other in the first two Phase 1 messages: 0x09002689DFD6B712. This vendor ID number is specified in the XAuth Internet draft, draft-beaulieu-ike-xauth-02.txt.

After the completion of Phase 1 negotiations, the NetScreen device sends a login prompt to the XAuth user at the remote site. If the XAuth user successfully logs on with the correct user name and password, the NetScreen device assigns an IP address, 32-bit netmask, DNS server addresses, and WINS server addresses to the user, and the two parties continue with Phase 2 negotiations.

The XAuth user has 60 seconds to complete the login process. If the first login attempt fails, the XAuth user can make up to four more attempts, having 60 seconds for each attempt. If the user fails after five consecutive attempts, the NetScreen device stops providing a login prompt and severs the session.

At a minimum, the XAuth-assigned IP address belongs to a user for the duration specified as the XAuth address lifetime. The IP address might belong to the XAuth user longer, depending on when the Phase 1 and Phase 2 security associations (SAs) rekey. The following example illustrates the relationship of Phase 1 and Phase 2 rekey operations and the XAuth IP address lifetime.

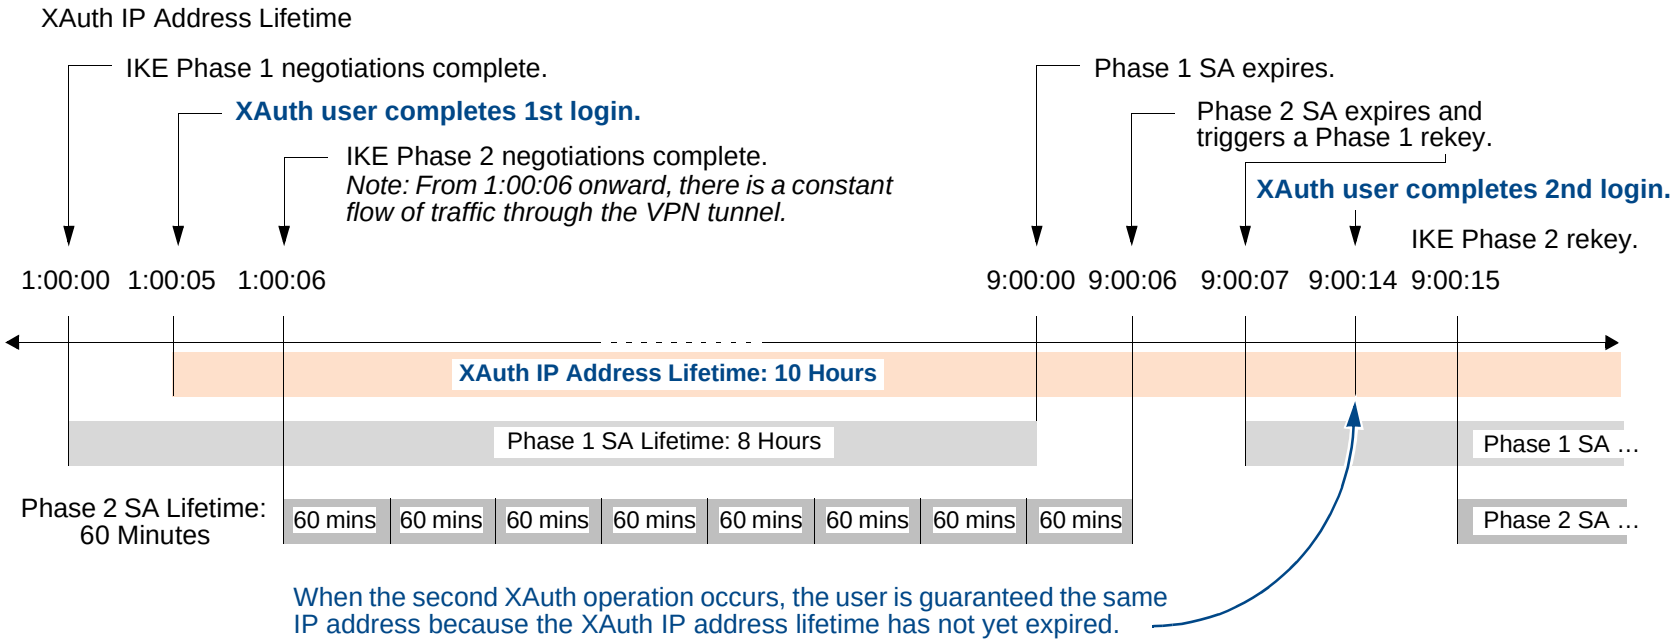

1. The Phase 1 SA is set with an 8-hour lifetime and expires after the first 8 hours.
2. The Phase 2 SA lifetime is set for 60 minutes. Because there is a 5-second delay during the initial IKE negotiations while the XAuth user enters his user name and password, the eighth Phase 2 SA expires 8 hours and 6 seconds (5 seconds for the XAuth login + 1 second for Phase 2 negotiations) after Phase 1 negotiations complete.
3. Because there is active VPN traffic, the expiration of the eighth Phase 2 SA causes the Phase 1 SA, which expired 6 seconds prior, to rekey; that is, Phase 1 IKE negotiations (or “renegotiations”) occur.

4. After Phase 1 IKE renegotiations complete, the NetScreen device prompts the XAuth user to log in again.

   Note: *To avoid repeating further logins after the initial one, configure the VPN tunnel with any idletime other than 0 with the CLI command:* ***set vpn*** *name* ***gateway*** *name* ***idletime*** *number (in minutes). If there is VPN activity at the completion of Phase 1 IKE renegotiations, the NetScreen device does not prompt the XAuth user to log in again. This option enables the user to download large files, transmit or receive streaming media, of participate in Web conferences without interruption.*

5. Because the XAuth address lifetime (10 hours) exceeds the Phase 1 SA lifetime, the user keeps the same IP address—although the user might get a different address after the next Phase 1 rekey occurs.

   If the XAuth address lifetime had been shorter than the Phase 1 SA lifetime, the NetScreen device would have assigned the user another IP address, which might or might not have been the same as the previously assigned address[16].

   Note: *To change the address lifetime, do either of the following:*

   - *(WebUI) VPNs > AutoKey Advanced > XAuth Settings: Enter a number (minutes) in the Reserve Private IP for XAuth User field, and then click* ***Apply****.*
   - *(CLI) set xauth lifetime* number

   To effectively disable the address lifetime feature, enter a value of 1—the minimum value allowed.

16. If it is crucial that a user always be assigned the same IP address, you can specify an address in the user configuration. The NetScreen device then assigns this address instead of assigning one at random from an IP pool. Note that such an address must not be in an IP pool or it might get assigned to another user and be unavailable when needed.

Example: XAuth Authentication (Local User)

In this example, you define an XAuth user named x1 with password aGgb80L0ws on the local database.

You then reference this user in a remote IKE gateway configuration to a peer at IP 2.2.2.2. You name the remote gateway “gw1”, specify Main mode and the proposal pre-g2-3des-sha for Phase 1 negotiations, and use the preshared key “netscreen1”. You name the VPN tunnel “vpn1” and specify the “Compatible” set of proposals for Phase 2 negotiations. You choose the Untrust zone interface ethernet3 as the outgoing interface.

WebUI

1. XAuth User

Objects > Users > Local > New: Enter the following, and then click **OK**:

User Name: x1

Status: Enable

XAuth User: (select)

User Password: iDa84rNk

Confirm Password: iDa84rNk

2. VPN

VPNs > AutoKey Advanced > Gateway > New: Enter the following, and then click **OK**:

Gateway Name: gw1

Security Level: Custom

Remote Gateway Type:

Static IP Address: (select), Address/Hostname: 2.2.2.2

Preshared Key: netscreen1

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click Return to return to the basic Gateway configuration page:

Security Level: Custom: (select)

Phase 1 Proposal: pre-g2-3des-sha

Mode (Initiator): Main (ID Protection)

XAuth Server: (select)

Local Authentication: (select)

User: (select), x1

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: vpn1

Security Level: Compatible

Remote Gateway Tunnel: gw1

CLI

1. XAuth User

```
set user x1 password aGgb80L0ws
set user x1 type xauth
unset user x1 type auth
```
[17]

2. VPN

```
set ike gate gw1 ip 2.2.2.2 main outgoing-interface ethernet3 preshare
    netscreen1 proposal pre-g2-3des-sha
set ike gateway gw1 xauth server Local user x1
set vpn vpn1 gateway gw1 sec-level compatible
save
```

17. The CLI command **set user** *name_str* **password** *pswd_str* creates an auth user. To create an XAuth-only user, you must define the user as an XAuth user (**set user** *name_str* type **xauth**), and then remove the auth user definition (**unset user** *name_str* **type auth**).

Example: XAuth Authentication (Local User Group)

In this example, you create a user group named xa-grp1 on the local database and add the XAuth user “x1” that you created in the previous example, “Example: XAuth Authentication (Local User)” on page 440. When you add that user to the group, it automatically becomes an XAuth user group.

You then reference this group in a remote IKE gateway configuration to a peer at IP 2.2.2.2. You name the remote gateway “gw2”, specify Main mode and the proposal pre-g2-3des-sha for Phase 1 negotiations, and use the preshared key “netscreen2”. You name the VPN tunnel “vpn2” and specify the “Compatible” set of proposals for Phase 2 negotiations. You choose the Untrust zone interface ethernet3 as the outgoing interface.

WebUI

1. XAuth User Group

Objects > Users > Local Groups > New: Enter **xa-grp1** in the Group Name field, do the following, and then click **OK**:

Select **x1** and use the **<<** button to move him from the Available Members column to the Group Members column.

2. VPN

VPNs > AutoKey Advanced > Gateway > New: Enter the following, and then click **OK**:

Gateway Name: gw2

Security Level: Custom

Remote Gateway Type:

Static IP Address: (select), Address/Hostname: 2.2.2.2

Preshared Key: netscreen2

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Gateway configuration page:

Phase 1 Proposal: pre-g2-3des-sha

Mode (Initiator): Main (ID Protection)

XAuth Server: (select)

Local Authentication: (select

User Group: (select), xa-grp1

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: vpn2

Security Level: Compatible

Remote Gateway Tunnel:

Predefined: (select), gw2

CLI

1. XAuth User Group

```
set user-group xa-grp1 location local
set user-group xa-grp1 user x1
```

2. VPN

```
set ike gate gw2 ip 2.2.2.2 main outgoing-interface ethernet3 preshare
    netscreen2 proposal pre-g2-3des-sha
set ike gateway gw2 xauth server Local user-group xa-grp1
set vpn vpn2 gateway gw2 sec-level compatible
save
```

Example: XAuth Authentication (External User)

In this example, you reference an XAuth user named “xa-1” with password iNWw10bd01 that you have previously loaded on an external SecurID auth server. This example uses almost the same configuration of the SecurID auth server as defined in “Example: Defining an Auth Server Object for SecurID” on page 391, except that here you define the account type as XAuth.

You reference XAuth user xa-1 in a remote IKE gateway configuration to a peer at IP 2.2.2.2. You name the remote gateway “gw3”, specify Main mode and the proposal pre-g2-3des-sha for Phase 1 negotiations, and use the preshared key “netscreen3”. You name the VPN tunnel “vpn3” and specify the proposal g2-esp-3des-sha for Phase 2 negotiations. You choose the Untrust zone interface ethernet3 as the outgoing interface.

WebUI

1. External SecurID Auth Server

Configuration > Auth > Servers > New: Enter the following, and then click **OK**:

Name: securid1

IP/Domain Name: 10.20.2.100

Backup1: 10.20.2.110

Timeout: 60

Account Type: XAuth

SecurID: (select)

Client Retries: 3

Client Timeout: 10 seconds

Authentication Port: 15000

Encryption Type: DES

User Duress: No

2. XAuth User

Define the auth user “xa-1” with password iNWw10bd01 on the external SecurID auth server securid1.

3. VPN

VPNs > AutoKey Advanced > Gateway > New: Enter the following, and then click **OK**:

Gateway Name: gw3

Security Level: Custom

Remote Gateway Type:

Static IP Address: (select), Address/Hostname: 2.2.2.2

Preshared Key: netscreen3

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Gateway configuration page:

Phase 1 Proposal: pre-g2-3des-sha

Mode (Initiator): Main (ID Protection)

XAuth Server: (select)

External Authentication: (select), securid1

User: (select)

Name: xa-1

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: vpn3

Security Level: Compatible

Remote Gateway Tunnel:

Predefined: (select), gw3

CLI

1. External SecurID Auth Server

```
set auth-server securid1 type securid
set auth-server securid1 server-name 10.20.2.100
set auth-server securid1 backup1 10.20.2.110
set auth-server securid1 timeout 60
set auth-server securid1 account-type xauth
set auth-server securid1 securid retries 3
set auth-server securid1 securid timeout 10
set auth-server securid1 securid auth-port 15000
set auth-server securid1 securid encr 1
set auth-server securid1 securid duress 0
```

2. XAuth User

Define the auth user “xa-1” with password iNWw10bd01 on the external SecurID auth-server securid1.

3. VPN

```
set ike gate gw3 ip 2.2.2.2 main outgoing-interface ethernet3 preshare
    netscreen3 proposal pre-g2-3des-sha
set ike gateway gw3 xauth server securid1 user xa-1
set vpn vpn3 gateway gw3 sec-level compatible
save
```

Example: XAuth Authentication (External User Group)

In this example, you configure an external RADIUS auth server named “radius1”[18] and define an external auth user group named “xa-grp2”. You define the external XAuth user group xa-grp2 in two places:

1. External RADIUS auth server “radius1”
2. NetScreen device

You populate the XAuth user group “xa-grp2” with XAuth users on the RADIUS server only, leaving the group unpopulated on the NetScreen device. The members in this group are resellers at a remote site who require access to FTP servers in the corporate LAN. You add an entry in the Untrust zone address book for the remote site with IP address 10.2.2.0/24 and name “reseller1”. You also enter an address in the Trust zone address book for the FTP server “rsl-srv1” with IP address 10.1.1.5/32.

You configure a VPN tunnel to 2.2.2.2 to authenticate XAuth users in the user group xa-grp2. You name the remote gateway “gw4”, specify Main mode and the proposal pre-g2-3des-sha for Phase 1 negotiations, and use the preshared key “netscreen4”. You name the VPN tunnel “vpn4” and specify the “Compatible” set of proposals for Phase 2 negotiations. You choose the Untrust zone interface ethernet3 as the outgoing interface.

Finally, you set up create a policy permitting FTP traffic from that reseller1 in the Untrust zone via vpn4 to rsl-svr1 in the Trust zone.

RADIUS Server

1. Load the NetScreen dictionary file on the RADIUS server.

   Note: *For information on the NetScreen dictionary file, see “NetScreen Dictionary File” on page 381. For instructions on loading the dictionary file onto a RADIUS server, refer to the RADIUS server documentation.*

2. Enter auth user group “xa-grp2” on the external auth server “radius1“, and then populate it with XAuth user accounts.

18. The RADIUS auth server configuration is nearly identical to that in “Example: Defining an Auth Server Object for RADIUS” on page 388, except that in this example you only specify “xauth” as the user account type.

WebUI

1. Auth Server

Configuration > Auth > Servers > New: Enter the following, and then click **OK**:

Name: radius1
IP/Domain Name: 10.20.1.100
Backup1: 10.20.1.110
Backup2: 10.20.1.120
Timeout: 30
Account Type: XAuth
RADIUS: (select)
RADIUS Port: 4500
Shared Secret: A56htYY97kl

2. External User Group

Objects > Users > External Groups > New: Enter the following, and then click **OK**:

Group Name: xa-grp2
Group Type: XAuth

3. Address

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: reseller1
IP Address/Domain Name:
IP/Netmask: (select), 10.2.2.0/24
Zone: Untrust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: rsl-svr1

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.5/32

Zone: Trust

4. XAuth User

Define the auth user “xa-1” with password iNWw10bd01 on the external SecurID auth server securid1.

5. VPN

VPNs > AutoKey Advanced > Gateway > New: Enter the following, and then click **OK**:

Gateway Name: gw4

Security Level: Custom

Remote Gateway Type:

Static IP Address: (select), Address/Hostname: 2.2.2.2

Preshared Key: netscreen4

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Gateway configuration page:

Phase 1 Proposal: pre-g2-3des-sha

Mode (Initiator): Main (ID Protection)

XAuth Server: (select)

External Authentication: (select), securid1

User Group: (select)

Name: xa-grp2

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: vpn4

Security Level: Compatible

Remote Gateway:

Predefined: (select), gw4

6. Policy

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), reseller1

Destination Address:

Address Book Entry: (select), rsl-svr1

Service: FTP-Get

Action: Tunnel

Tunnel VPN: vpn4

Modify matching bidirectional VPN policy: (clear)

Position at Top: (select)

CLI

1. Auth Server

```
set auth-server radius1 type radius
set auth-server radius1 account-type xauth
set auth-server radius1 server-name 10.20.1.100
set auth-server radius1 backup1 10.20.1.110
set auth-server radius1 backup2 10.20.1.120
set auth-server radius1 timeout 30
set auth-server radius1 radius port 4500
set auth-server radius1 radius secret A56htYY97kl
```

2. External User Group

```
set user-group xa-grp2 location external
set user-group xa-grp2 type xauth
```

3. Address

```
set address untrust reseller1 10.2.2.0/24
set address trust rsl-svr1 10.1.1.5/32
```

4. VPN

```
set ike gate gw4 ip 2.2.2.2 main outgoing-interface ethernet3 preshare
    netscreen4 proposal pre-g2-3des-sha
set ike gateway gw4 xauth server radius1 user-group xa-grp2
set vpn vpn4 gateway gw4 sec-level compatible
```

5. Policy

```
set policy top from untrust to trust reseller1 rsl-svr1 ftp-get tunnel vpn vpn4
save
```

Example: XAuth Authentication and Address Assignments (Local User Group)

In this example, you set up both authentication and IP, DNS server, and WINS server IP address assignments for an IKE/XAuth user group stored on the local database[19]. When an IKE/XAuth user makes a dialup VPN connection to the NetScreen device, the NetScreen device authenticates the IKE user (that is, the client device) using an IKE ID and an RSA certificate during Phase 1 negotiations. The NetScreen device then authenticates the XAuth user (that is, the individual using the device) using a user name and password and assigns IP, DNS server, and WINS server IP addresses between Phase 1 and Phase 2 negotiations.

You create a local user group ixa-grp1. You then define two IKE/XAuth users named “ixa-u1” (password: ccF1m84s) and “ixa-u2” (password: C113g1tw) and add them to the group, thereby defining the group type as IKE/XAuth. (The addition of other IKE/XAuth users to the group is not included in the example.)

You create a DIP pool named xa-pool1 with an address range from 10.2.2.1 to 10.2.2.100. This is the pool of addresses from which the NetScreen device draws an IP address when assigning one to an XAuth user.

Note: *The DIP pool must be in a different address space than that of the zone to which the XAuth user directs traffic to avoid routing problems and duplicate address assignments.*

19. You can also use an external RADIUS auth server for XAuth user authentication and address assignments.You can use an external SecurID or LDAP auth server for XAuth authentication only (not for address assignments). For IKE user authentication, you can only use the local database.

You configure the following XAuth default settings:

- Set the XAUTH address timeout to 480 minutes.
- Select the local database as the default auth server.
- Enable CHAP, Challenge Handshake Authentication Protocol, in which the NetScreen device sends a challenge (encryption key) to the remote client, who uses the key to encrypt his or her login name and password.
- Select xa-pool1 as the default DIP pool.
- Define the primary and secondary DNS server IP addresses as 10.1.1.150 and 10.1.1.151 respectively.
- Define the primary and secondary WINS server IP addresses as 10.1.1.160 and 10.1.1.161 respectively.

You configure an IKE gateway named “ixa-gw1”, referencing user group ixa-grp1 and using the default XAuth auth server settings. You then configure a VPN tunnel name named “ixa-tun1” and a policy permitting traffic from ixa-grp1 to the Trust zone (IP address 10.1.1.0/24) via VPN tunnel ixa-tun1.

WebUI

1. IKE/XAuth Users and User Group

Objects > Users > Local > New: Enter the following, and then click **OK**:

User Name: ixa-u1
Status: Enable
IKE User: (select)
Simple Identity: (select)
IKE ID Type: AUTO
IKE Identity : u1@ns.com
XAuth User: (select)
User Password: ccF1m84s
Confirm Password: ccF1m84s

Objects > Users > Local > New: Enter the following, and then click **OK**:

User Name: ixa-u2

Status: Enable

IKE User: (select)

Simple Identity: (select)

IKE ID Type: AUTO

IKE Identity : u2@ns.com

XAuth User: (select)

User Password: C113g1tw

Confirm Password: C113g1tw

Objects > Users > Local Groups > New: Enter **ixa-grp1** in the Group Name field, do the following, and then click **OK**:

Select **ixa-u1** and use the **<<** button to move him from the Available Members column to the Group Members column.

Select **ixa-u2** and use the **<<** button to move him from the Available Members column to the Group Members column.

2. IP Pool

Objects > IP Pools > New: Enter the following, and then click **OK**:

IP Pool Name: xa-pool1

Start IP: 10.2.2.1

End IP: 10.2.2.100

3. Default XAuth Auth Server

VPNs > AutoKey Advanced > XAuth Settings: Enter the following, and then click **Apply**:

Reserve Private IP for XAuth User: 480 Minutes
Default Authentication Server: Local
Query Client Settings on Default Server: (clear)
CHAP: (select)
IP Pool Name: xa-pool1
DNS Primary Server IP: 10.1.1.150
DNS Secondary Server IP: 10.1.1.151
WINS Primary Server IP: 10.1.1.160
WINS Secondary Server IP : 10.1.1.161

4. Address

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Trust_zone
IP Address/Domain Name:
  IP/Netmask: (select), 10.1.1.0/24
Zone: Trust

5. VPN

VPNs > AutoKey Advanced > Gateway > New: Enter the following, and then click **OK**:

Gateway Name: ixa-gw1
Security Level: Custom
Remote Gateway Type:
  Dialup User Group: (select)
    Group: ixa-grp1

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Gateway configuration page:

Phase 1 Proposal: rsa-g2-3des-sha

Mode (Initiator): Aggressive

Outgoing Interface: ethernet3

XAuth Server: (select)

Use Default: (select)

User Group: (select), ixa-grp1

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: ixa-vpn1

Security Level: Compatible

Remote Gateway:

Predefined: (select), ixa-gw1

6. Policy

Policies > (From: Untrust; To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Dial-Up VPN

Destination Address:

Address Book Entry: (select), Trust_zone

Service: ANY

Action: Tunnel

Tunnel VPN: ixa-vpn1

Modify matching bidirectional VPN policy: (clear)

Position at Top: (select)

CLI

1. IKE/XAuth Users and User Group

```
set user-group ixa-grp1 location local
set user ixa-u1 type ike xauth
set user ixa-u1 ike-id u-fqdn u1@ns.com
set user ixa-u1 password ccF1m84s
unset user ixa-u1 type auth
set user ixa-u2 type ike xauth
set user ixa-u2 ike-id u-fqdn u2@ns.com
set user ixa-u2 password C113g1tw
unset user ixa-u2 type auth
```

2. IP Pool

```
set ippool xa-pool1 10.2.2.1 10.2.2.100
```

3. Default XAuth Auth Server

```
set xauth lifetime 480
set xauth default auth server Local chap
set xauth default ippool xa-pool1
set xauth default dns1 10.1.1.150
set xauth default dns2 10.1.1.151
set xauth default wins1 10.1.1.160
set xauth default wins210.1.1.161
```

4. Address

```
set address trust Trust_zone 10.1.1.0/24
```

5. VPN

```
set ike gateway ixa-gw1 dialup ixa-grp1 aggressive outgoing-interface ethernet3
    proposal rsa-g2-3des-sha
set ike gateway ixa-gw1 xauth server Local user-group ixa-grp1
set vpn ixa-vpn1 gateway ixa-gw1 sec-level compatible
```

6. Policy

```
set policy top from untrust to trust “Dial-Up VPN” Trust_zone any tunnel vpn
    ixa-vpn1
save
```

XAuth Client

An XAuth client is a remote user or device that connects to an XAuth server via an AutoKey IKE VPN tunnel. A NetScreen device can act as an XAuth client, responding to authentication requests from a remote XAuth server. After the completion of Phase 1 negotiations, the remote XAuth server sends a login prompt to the NetScreen device. If the NetScreen device acting as an XAuth client successfully logs in with the correct user name and password, Phase 2 negotiations commence.

To configure the NetScreen device as an XAuth client, you must specify the following:

- IKE gateway name
- XAuth user name and password

You can configure the following types of XAuth authentication:

- Any — Allows either Challenge Handshake Authentication Protocol (CHAP) or Password Authentication Protocol (PAP)
- CHAP — Allows CHAP only

Example: NetScreen Device as an XAuth Client

In this example, you first configure a remote IKE gateway *gw1* with IP address 2.2.2.2. You specify the standard security level and use the preshared key *netscreen1*. You then configure an XAuth client for the IKE gateway with the username *beluga9* and the password *1234567*. You also require CHAP authentication for the client.

WebUI

VPN > AutoKey Advanced > Gateway > New: Enter the following, and then click **OK**:

Gateway Name: gw1

Security Level: Standard (select)

Remote Gateway Type:

Static IP Address: (select), Address/Hostname: 2.2.2.2

Preshared Key: netscreen1

Outgoing Interface: Untrust

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Gateway configuration page:

XAuth Client: (select)

User Name: beluga9

Password: 1234567

Allowed Authentication Type: (select), CHAP Only

CLI

```
set ike gateway gw1 ip 2.2.2.2 Main outgoing-interface untrust preshare
    netscreen1 sec-level standard
set ike gateway gw1 xauth client chap username beluga1 password 1234567
save
```

L2TP Users and User Groups

Layer 2 Tunneling Protocol (L2TP) provides a means for authenticating remote users and assigning IP, DNS server, and WINS server addresses. You can configure the NetScreen device to use either the local database or an external auth server to authenticate L2TP users. To make IP, DNS server, and WINS server address assignments, you can configure the NetScreen device to use either the local database or a RADIUS server (loaded with the NetScreen dictionary file—see "NetScreen Dictionary File" on page 381).

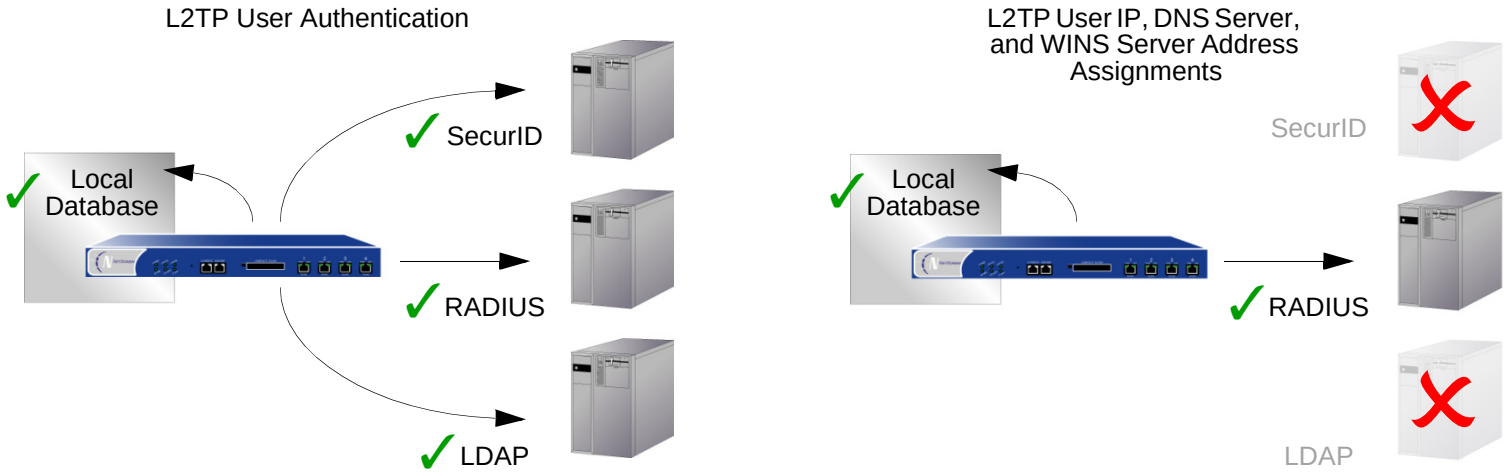

You can even use a combination of auth servers, a different one for each of the two aspects of L2TP. For example, you might use a SecurID server to authenticate an L2TP user but make the address assignments from the local database. The following example illustrates the application of two auth servers to handle both components of L2TP. For other examples, along with a detailed examination of L2TP, see "L2TP" on page **5**-269.

Example: Local and External L2TP Auth Servers

In this example, you set up an external SecurID auth server to authenticate L2TP users, and you use the local database to assign L2TP users with IP, DNS server, and WINS server addresses.

The external SecurID auth server is securid1. It is nearly identical to the auth server configuration in “Example: Defining an Auth Server Object for SecurID” on page 391 except that the account type is L2TP. The SecurID auth server parameters are as follows:

- Name: securid1
- IP Address: 10.20.2.100
- Backup1 IP Address: 10.20.2.110
- Port: 15000
- Encryption: DES
- Client Retries: 3
- Client Timeout: 10 seconds
- Idle Timeout: 60 minutes
- Account Type: L2TP

The L2TP default settings are as follows:

- IP Pool: l2tp1 (172.168.1.1 – 172.168.1.100)
- DNS Primary Server IP: 10.20.2.50
- DNS Secondary Server IP: 10.20.2.51
- PPP Authentication: CHAP
- WINS Primary Server IP: 10.20.2.60
- WINS Primary Server IP: 10.20.2.61

After configuring the NetScreen device with the above settings, you create an L2TP tunnel named “l2tp-tun1” that references securid1 for authentication and uses the default settings for address assignments.

You must also set up the SecurID server as shown above and populate it with L2TP users.

Note: *An L2TP-only configuration is not secure. To add security to an L2TP tunnel, it is recommended that you combine it with an IPSec tunnel, which must be in Transport mode, as done in “Example: Configuring L2TP-over-IPSec” on page* ***5*** *-286.*

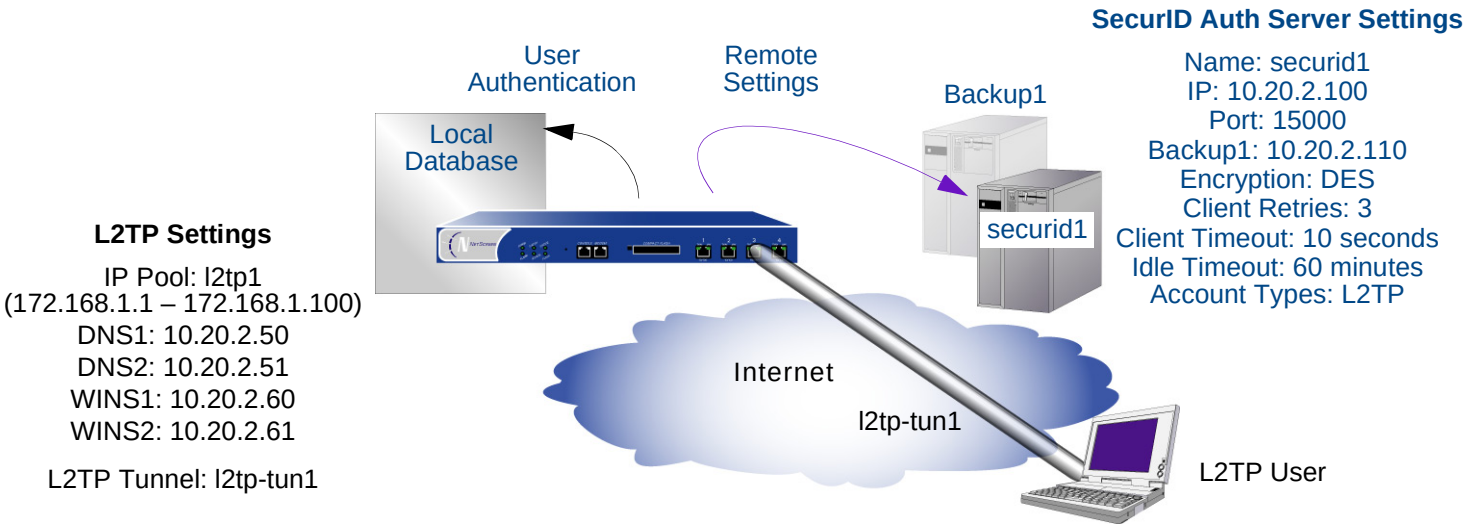

WebUI

1. Auth Server

Configuration > Auth > Servers > New: Enter the following, and then click **OK**:

Name: securid1
IP/Domain Name: 10.20.2.100
Backup1: 10.20.2.110
Timeout: 60
Account Type: L2TP
SecurID: (select)
  Client Retries: 3
  Client Timeout: 10 seconds
  Authentication Port: 15000
  Encryption Type: DES
  Use Duress: No

2. IP Pool

Objects > IP Pools > New: Enter the following, and then click **OK**:

IP Pool Name: l2tp1

Start IP: 172.168.1.1

End IP: 172.168.1.100

3. L2TP Default Settings

VPNs > L2TP > Default Settings: Enter the following, and then click **Apply**:

Default Authentication Server: Local

IP Pool Name: l2tp1

PPP Authentication: CHAP

DNS Primary Server IP: 10.20.2.50

DNS Secondary Server IP: 10.20.2.51

WINS Primary Server IP: 10.20.2.60

WINS Secondary Server IP: 10.20.2.61

4. L2TP Tunnel

VPNs > L2TP > Tunnel > New: Enter the following, and then click **OK**:

Name: l2tp-tun1

Use Custom Settings: (select)

Authentication Server: securid1

Query Remote Settings: (clear)

Dialup User: (select), Allow Any

CLI

1. Auth Server

```
set auth-server securid1 type securid
set auth-server securid1 server-name 10.20.2.100
set auth-server securid1 backup1 10.20.2.110
set auth-server securid1 timeout 60
set auth-server securid1 account-type l2tp
set auth-server securid1 securid retries 3
set auth-server securid1 securid timeout 10
set auth-server securid1 securid auth-port 15000
set auth-server securid1 securid encr 1
set auth-server securid1 securid duress 0
```

2. IP Pool

```
set ippool l2tp1 172.168.1.1 172.168.1.100
```

3. L2TP Default Settings

```
set l2tp default auth server Local
set l2tp default ippool l2tp1
set l2tp default ppp-auth chap
set l2tp dns1 10.20.2.50
set l2tp dns1 10.20.2.51
set l2tp wins1 10.20.2.60
set l2tp wins2 10.20.2.61
```

4. L2TP Tunnel

```
set l2tp l2tp-tun1
set l2tp l2tp-tun1 auth server securid1
save
```

Admin Users

Admin users are the administrators of a NetScreen device. There are five kinds of admin users:

- Root admin
- Root-level read/write admin
- Root-level read-only admin
- Vsys admin
- Vsys read-only admin

> ***Note:*** *For information regarding the privileges of each type of admin user and for examples of the creation, modification and removal of admin users, see "Administration" on page* ***3****-1.*

Although the profile of the root user of a NetScreen device must be stored in the local database, you can store vsys users and root-level admin users with read/write and read-only privileges either in the local database or on an external auth server.

If you store admin user accounts on an external RADIUS auth server and you load the NetScreen dictionary file on the auth server (see "NetScreen Dictionary File" on page 381), you can elect to query admin privileges defined on the server. Optionally, you can specify a privilege level to be applied globally to all admin users stored on that auth server. You can specify either read/write or read-only privileges. If you store admin users on an external SecurID or LDAP auth server, or on a RADIUS server without the NetScreen dictionary file, you cannot define their privilege attributes on the auth server. Therefore, you must assign a privilege level to them on the NetScreen device.

If set on the NetScreen device:	and the RADIUS server is loaded with the NetScreen dictionary file, then:	and a SecurID, LDAP, or RADIUS server without the NetScreen dictionary file, then:
Get privileges from RADIUS server	Assign appropriate privileges	Root- or vsys-level admin login fails
Assign read/write privileges to external admin	Assign root- or vsys-level read/write privileges	Assign root-level read/write privileges Vsys admin login fails
Assign read-only privileges to external admin	Assign root- or vsys-level read-only privileges	Assign root-level read-only privileges Vsys admin login fails

The admin authentication process proceeds as shown in the following illustration:

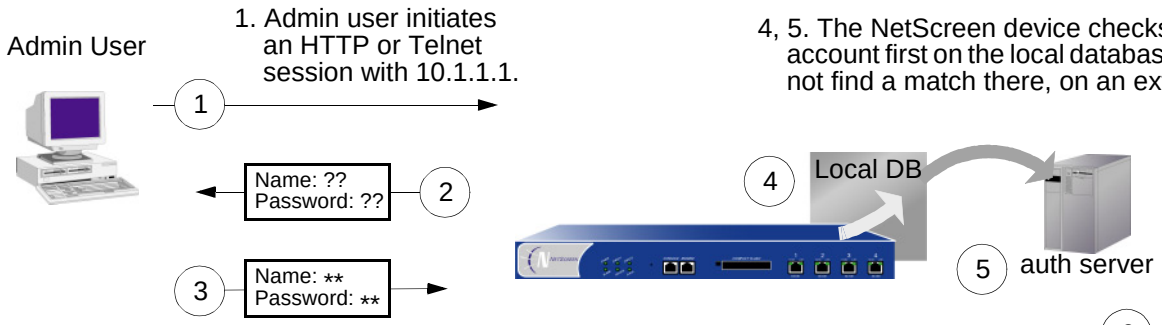

MULTIPLE-TYPE USERS

You can combine auth, IKE, L2TP, XAuth users to create the following combinations to store on the local database:

- Auth/IKE User
- Auth/L2TP User
- Auth/IKE/L2TP User
- IKE/L2TP User
- Auth/XAuth User
- Auth/IKE/XAuth User
- IKE/XAuth User
- L2TP/XAuth User
- IKE/L2TP/XAuth User
- Auth/IKE/L2TP/XAuth User

Although you can make all of the above combinations when defining multiple-type user accounts on the local database, consider the following points before creating them:

- Combining an IKE user type with any other user type limits the potential to scale. You must store an IKE user account on the local database. If you create auth/IKE, IKE/L2TP, and IKE/XAuth user accounts and then the number of users grows beyond the capacity of the local database, you will not be able to relocate these accounts to an external auth server. If you separate IKE user accounts from other types of accounts, you have the flexibility to move the non-IKE user accounts to an external auth server should the need arise to do so.
- L2TP and XAuth provide the same services: remote user authentication and IP, DNS server, and WINS server address assignments. It is not recommended to use L2TP and XAuth together for an L2TP-over-IPSec tunnel. Not only do the two protocols accomplish the same goals, but the L2TP address assignments overwrite the XAuth address assignments after Phase 2 IKE negotiations complete and L2TP negotiations take place.
- If you create a multiple-type user account on the local database combining auth/L2TP or auth/XAuth, the same user name and password must be used for both logins.

  Although it is more convenient to create a single multiple-type user account, separating the user types into two single accounts allows you to increase security. For example, you can store an auth user account on an external auth server and an XAuth user account on the local database. You can then assign different login user names and passwords to each account, and reference the XAuth user in the IKE gateway configuration and the auth user in the policy configuration. The dialup VPN user must authenticate himself twice, potentially with two completely different user names and passwords.

GROUP EXPRESSIONS

A group expression is a statement that you can use in policies to conditionalize the requirements for authentication. Group expressions allow you to combine users, user groups, or other group expressions as alternatives for authentication (“a” OR “b”), or as requirements for authentication (“a” AND “b”). You can also use group expressions to exclude a user, user group, or another group expression (NOT “c”).

Note: *Although you define group expressions on the NetScreen device (and store them on the local database), the users and user groups that you reference in the group expressions must be stored on an external RADIUS server. A RADIUS server allows a user to belong to more than one user group. The local database does not permit this.*

Group expressions make use of the three operators OR, AND, and NOT. The objects in the expression to which OR, AND, and NOT relate can be an auth user, an auth user group, or a previously defined group expression.

Users

OR – If the authentication aspect of a policy specifies that the user be “a” OR “b”, then the NetScreen device authenticates the user if he or she is either one.

AND – The use of AND in a group expression requires that at least one of the two expression objects be either a user group or a group expression. (It is illogical to require a user to be user “a” AND user “b”.) If the authentication aspect of a policy requires that the user be “a” AND a member of group “b”, then the NetScreen device authenticates the user only if those two conditions are met.

NOT – If the authentication aspect of a policy specifies that the user be anyone other than user “c” (NOT “c”), then the NetScreen device authenticates the user as long as he or she is not that user.

User Groups

OR – If the authentication aspect of a policy specifies that the user belong to group “a” OR group “b”, then the NetScreen device authenticates the user if he or she belongs to either group.

AND – If the authentication aspect of a policy requires that the user belong to group “a” AND group “b”, then the NetScreen device authenticates the user only if he or she belongs to both groups.

NOT – If the authentication aspect of a policy specifies that the user belong to any group other than group “c” (NOT “c”), then the NetScreen device authenticates the user if he or she does not belong to that group.

Group Expressions

OR – If the authentication aspect of a policy specifies that the user fit the description of group expression “a” OR group expression “b”, then the NetScreen device authenticates the user if either group expression applies to him or her.

AND – If the authentication aspect of a policy specifies that the user fit the description of group expression “a” AND group expression “b”, then the NetScreen device authenticates the user only if both group expressions apply to him or her.

NOT – If the authentication aspect of a policy specifies that the user not fit the description of group expression “c” (NOT “c”), then the NetScreen device authenticates the user only if he or she does not fit that group expression.

Example: Group Expressions (AND)

In this example, you create a group expression “s+m” that states “sales AND marketing”. You have previously created the auth user groups “sales” and “marketing” on an external RADIUS auth server named “radius1” and populated them with users. (For an example on how to configure an external RADIUS auth server, see “Example: Defining an Auth Server Object for RADIUS” on page 388.) You then use that group expression in an intrazone policy[20] whose authentication component requires a user be a member of both user groups to be able to access the confidential contents on a server named “project1” (10.1.1.70).

WebUI

1. Address

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: project1

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.70/32

Zone: Trust

2. Group Expression

Objects > Group Expressions > New: Enter the following, and then click **OK**:

Group Expression: s+m

AND: (select), sales AND marketing

20. For an intrazone policy to work properly, the source and destination addresses must be in different subnets connected to the NetScreen device through interfaces that are both bound to the same zone. There cannot be any other routing device beside the NetScreen device that can route traffic between the two addresses. For more information about intrazone policies, see “Policies” on page 197.

3. Policy

Policies > (From: Trust, To: Trust) New: Enter the following, and then click **OK**:

Source Address

Address Book Entry: (select), Any

Destination Address

Address Book Entry: (select), project1

Service: ANY

Action: Permit

Position at Top: (select)

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

Authentication: (select)

Auth Server: (select)

Use: radius1

Group Expression: (select), External Group Expression - s+m

CLI

1. Address

```
set address trust project1 10.1.1.70/32
```

2. Group Expression

```
set group-expression s+m sales and marketing
```

3. Policy

```
set policy top from trust to trust any project1 any permit auth server radius1
    group-expression s+m
save
```

Example: Group Expressions (OR)

In this example, you create a group expression “a/b” that states “amy OR basil”. You have previously created auth user accounts “amy” and “basil” on an external RADIUS auth server named “radius1”. (For an example on how to configure an external RADIUS auth server, see “Example: Defining an Auth Server Object for RADIUS” on page 388.) You then use that group expression in a policy from the Trust zone to the DMZ. The authentication component of the policy requires the user to be either amy or basil to be able to access the Web server named “web1” at 210.1.1.70.

WebUI

1. Address

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: web1

IP Address/Domain Name

IP/Netmask: (select), 210.1.1.70/32

Zone: DMZ

2. Group Expression

Objects > Group Expressions > New: Enter the following, and then click **OK**:

Group Expression: a/b

OR: (select), amy OR basil

3. Policy

Policies > (From: Trust, To: DMZ) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), web1

Service: ANY

Action: Permit

Position at Top: (select)

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

Authentication: (select)

Auth Server: (select)

Use: radius1

Group Expression: (select), External Group Expression - a/b

CLI

1. Address

```
set address trust project1 210.1.1.70/32
```

2. Group Expression

```
set group-expression a/b amy or basil
```

3. Policy

```
set policy top from trust to dmz any web1 any permit auth server radius1
    group-expression a/b
save
```

Example: Group Expressions (NOT)

In this example, you create a group expression “-temp” that states “NOT temp”. You have previously created a local auth user group “temp” on an external RADIUS auth server named “radius1”. (For an example on how to configure an external RADIUS auth server, see “Example: Defining an Auth Server Object for RADIUS” on page 388.) You then use that group expression in a policy from the Trust zone to the Untrust zone that allows Internet access to all full-time employees, but not to temporary contractors. The authentication component of the policy requires everyone in the Trust zone to be authenticated except the users in “temp”, who are denied access to the Untrust zone.

WebUI

1. Group Expression

   Objects > Group Expressions > New: Enter the following, and then click **OK**:

   Group Expression: -temp

   OR: (select), NOT temp

2. Policy

   Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

   Source Address:

   Address Book Entry: (select), Any

   Destination Address:

   Address Book Entry: (select), Any

   Service: HTTP

   Action: Permit

   Position at Top: (select)

   > Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

   Authentication: (select)

   Auth Server: (select)

   Use: Local

   Group Expression: (select), External Group Expression - -temp

CLI

1. Group Expression

```
set group-expression -temp not temp
```

2. Policy

```
set policy top from trust to untrust any any any permit auth server radius1
    group-expression -temp
save
```

BANNER CUSTOMIZATION

A banner is the message that appears onscreen in the following places during the following types of logins:

- At the top of a Telnet or console display when an admin user connects to log on to the NetScreen device
- At the top of a Web browser screen after an auth user has successfully logged on to a WebAuth address
- On a Telnet, FTP, or HTTP login prompt, success message, and fail message for auth users

All of the banners, except that for a console login, already have default messages. You can customize the messages that appear on the banners to better suit the network environment in which you use the NetScreen device.

Example: Customizing the WebAuth Success Message

In this example, you change the message that appears in the Web browser to indicate that an auth user has successfully authenticated himself after successfully logging on via WebAuth. The new message is “Authentication approved”.

WebUI

Configuration > Banners > WebAuth: In the Success Banner field, type **Authentication approved**, and then click **Apply**.

CLI

```
set webauth banner success “Authentication approved”
save
```

10

Traffic Shaping

This chapter discusses the various ways you can use your NetScreen device to manage limited bandwidth without compromising quality and availability of the network to all of your users.

The topics discussed include:

- "Applying Traffic Shaping" on page 478
  - "Managing Bandwidth at the Policy Level" on page 478
- "Setting Service Priorities" on page 485

APPLYING TRAFFIC SHAPING

Traffic shaping is the allocation of the appropriate amount of network bandwidth to every user and application on an interface. The appropriate amount of bandwidth is defined as cost-effective carrying capacity at a guaranteed Quality of Service (QoS). You can use a NetScreen device to shape traffic by creating policies and by applying appropriate rate controls to each class of traffic going through the NetScreen device.

Note: *You can only apply traffic shaping to policies whose destination zone has a single physical interface bound to it. NetScreen does not support traffic shaping if the destination zone contains one or more subinterfaces or more than one physical interface.*

Managing Bandwidth at the Policy Level

To classify traffic, you create a policy which specifies the amount of guaranteed bandwidth, the maximum bandwidth, and the priority for each class of traffic. The physical bandwidth of every interface is allocated to the guaranteed bandwidth parameter for all policies. If there is any bandwidth left over, it is sharable by any other traffic. In other words, each policy gets its guaranteed bandwidth and shares whatever is left over on a priority basis (up to the limit of its maximum bandwidth specification).

The traffic shaping function applies to traffic from all policies. If you turn off traffic shaping for a specific policy, while traffic shaping is still turned on for other policies, the system applies a default traffic shaping policy to that particular policy, with the following parameters:

- Guaranteed bandwidth 0
- Unlimited maximum bandwidth
- Priority of 7 (the lowest priority setting)[1]

If you do not want the system to assign this default traffic shaping policy to policies for which you have turned off traffic shaping, then turn off traffic shaping system wide via the CLI command **set traffic-shaping mode off**. You can set traffic shaping to automatic: **set traffic-shaping mode auto**. This allows the system to turn on traffic shaping when a policy requires it, and turn off traffic shaping when policies do not require it.

1. You can enable a mapping of the NetScreen priority levels to the DiffServ Codepoint Marking system. For more information about DS Codepoint Marking, see "Traffic Shaping" on page 215.

Example: Traffic Shaping

In this example, you partition 45Mbps of bandwidth on a T3 interface among three departments on the same subnet. The interface ethernet1 is bound to the Trust zone and ethernet3 is bound to the Untrust zone.

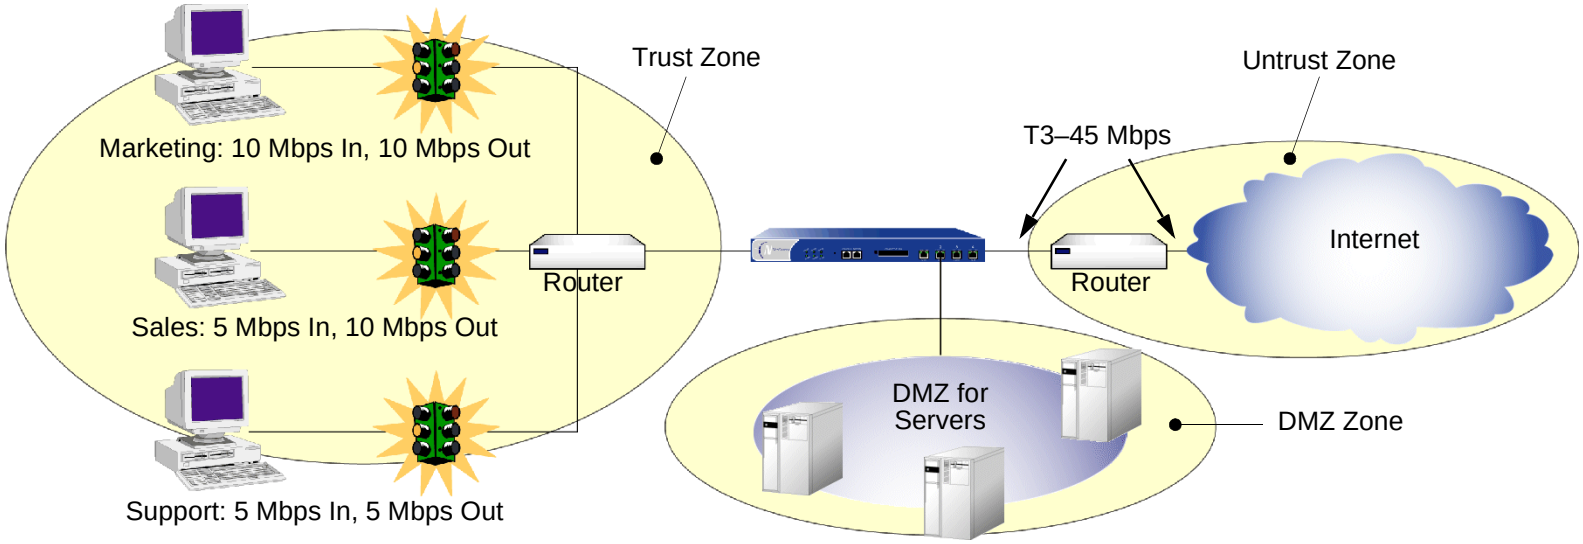

WebUI

1. Bandwidth on Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **OK**:

Traffic Bandwidth: 45000[2]

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Traffic Bandwidth: 45000

2. If you do not specify bandwidth settings on an interface, NetScreen uses whatever the available physical bandwidth is.

2. Bandwidth in Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Name: Marketing Traffic Shaping

Source Address:

Address Book Entry: (select), Marketing

Destination Address:

Address Book Entry: (select), Any

Service: Any

Action: Permit

VPN Tunnel: None[3]

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

Traffic Shaping: (select)

Guaranteed Bandwidth: 10000

Maximum Bandwidth: 15000

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Name: Sales Traffic Shaping Policy

Source Address:

Address Book Entry: (select), Sales

Destination Address:

Address Book Entry: (select), Any

Service: Any

3. You can also enable traffic shaping in policies referencing VPN tunnels.

Action: Permit

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

Traffic Shaping: (select)

Guaranteed Bandwidth: 10000

Maximum Bandwidth: 10000

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Name: Support Traffic Shaping Policy

Source Address:

Address Book Entry: (select), Support

Destination Address:

Address Book Entry: (select), Any

Service: Any

Action: Permit

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

Traffic Shaping: (select)

Guaranteed Bandwidth: 5000

Maximum Bandwidth: 10000

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Name: Allow Incoming Access to Marketing

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Marketing

Service: Any

Action: Permit

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

Traffic Shaping: (select)

Guaranteed Bandwidth: 10000

Maximum Bandwidth: 10000

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Name: Allow Incoming Access to Sales

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Sales

Service: Any

Action: Permit

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

Traffic Shaping: (select)

Guaranteed Bandwidth: 5000

Maximum Bandwidth: 10000

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Name: Allow Incoming Access to Support

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Support

Service: Any

Action: Permit

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

Traffic Shaping: (select)

Guaranteed Bandwidth: 5000

Maximum Bandwidth: 5000

CLI

To enable traffic shaping by policy, do the following:

1. **Bandwidth on Interfaces**

```
set interface ethernet1 bandwidth 45000[4]
set interface ethernet3 bandwidth 45000
```

2. **Bandwidth in Policies**

```
set policy name “Marketing Traffic Shaping” from trust to untrust marketing any
    any permit traffic gbw 10000 priority 0 mbw 15000
set policy name “Sales Traffic Shaping Policy” from trust to untrust sales any
    any permit traffic gbw 10000 priority 0 mbw 10000
set policy name “Support Traffic Shaping Policy” from trust to untrust support
    any any permit traffic gbw 5000 priority 0 mbw 10000
set policy name “Allow Incoming Access to Marketing” from untrust to trust any
    marketing any permit traffic gbw 10000 priority 0 mbw 10000
set policy name “Allow Incoming Access to Sales” from untrust to trust any
    sales any permit traffic gbw 5000 priority 0 mbw 10000
set policy name “Allow Incoming Access to Support” from untrust to trust any
    support any permit traffic gbw 5000 priority 0 mbw 5000
save
```

4. If you do not specify bandwidth settings on an interface, NetScreen uses whatever the available physical bandwidth is.

SETTING SERVICE PRIORITIES

The traffic shaping feature on NetScreen devices allows you to perform priority queuing on the bandwidth that is not allocated to guaranteed bandwidth, or unused guaranteed bandwidth. Priority queuing is a feature that allows all your users and applications to have access to available bandwidth as they need it, while ensuring that important traffic can get through, if necessary at the expense of less important traffic. Queuing allows NetScreen to buffer traffic in up to eight different priority queues. These eight queues are:

- High priority
- 2nd priority
- 3rd priority
- 4th priority
- 5th priority
- 6th priority
- 7th priority
- Low priority (default)

The priority setting for a policy means that the bandwidth not already guaranteed to other policies is queued on the basis of high priority first and low priority last. Policies with the same priority setting compete for bandwidth in a round robin fashion. The NetScreen device processes all of the traffic from all of the policies with high priority before processing any traffic from policies with the next lower priority setting, and so on, until all traffic requests have been processed. If traffic requests exceed available bandwidth, the lowest priority traffic is dropped.

Caution: *Be careful not to allocate more bandwidth than the interface can support. The policy configuration process does not prevent you from creating unsupported policy configurations. You can lose data if the guaranteed bandwidth on contending policies surpasses the traffic bandwidth set on the interface.*

If you do not allocate any guaranteed bandwidth, then you can use priority queuing to manage all of traffic on your network. That is, all high priority traffic is sent before any 2nd priority traffic is sent, and so on. The NetScreen device processes low priority traffic only after all other traffic has been processed.

Example: Priority Queuing

In this example, you configure the guaranteed and maximum bandwidth for three departments—Support, Sales, and Marketing— as follows:

	Outbound Guaranteed	Inbound Guaranteed	Combined Guaranteed	Priority
Support	5*	5	10	High
Sales	2.5	3.5	6	2
Marketing	2.5	1.5	4	3
Total	10	10	20	

* Megabits per second (Mbps)

If all three departments send and receive traffic concurrently through the NetScreen firewall, the NetScreen device must allocate 20 Mbps of bandwidth to fulfill the guaranteed policy requirements. The interface ethernet1 is bound to the Trust zone and ethernet3 is bound to the Untrust zone.

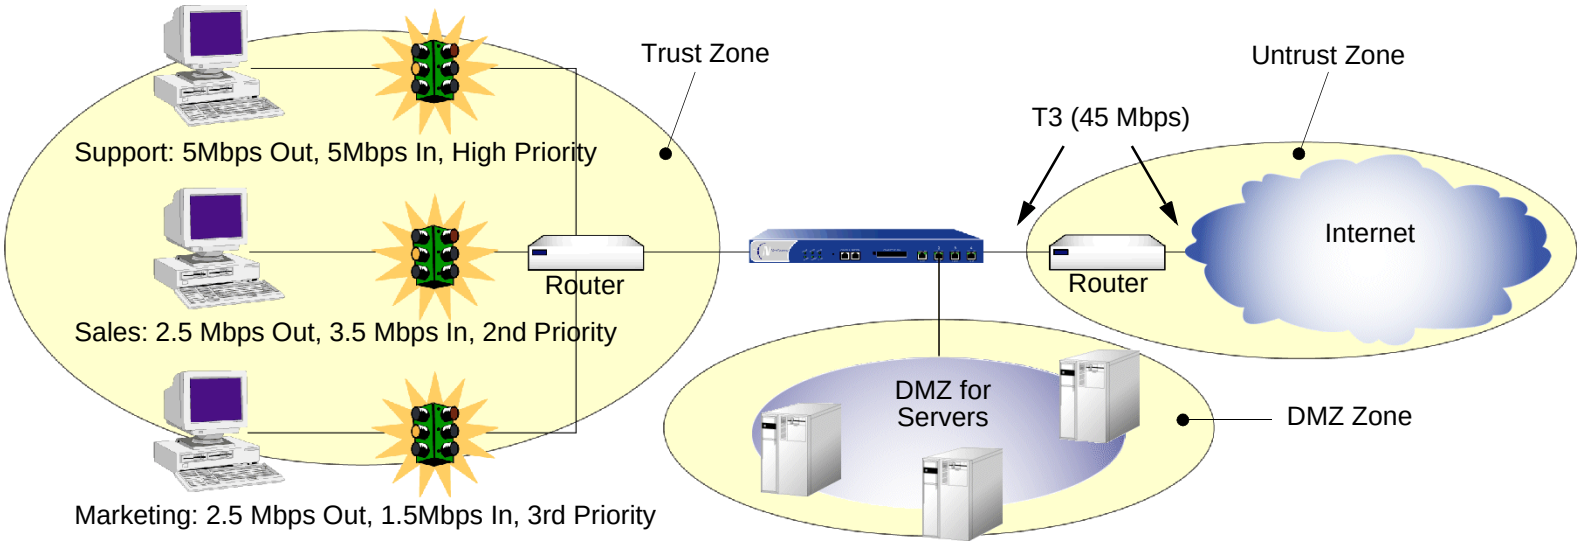

WebUI

1. Bandwidth on Interfaces

Interfaces > Edit (for ethernet1): Enter the following, and then click **OK**:

Traffic Bandwidth: 40000

Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Traffic Bandwidth: 40000

2. Bandwidth in Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Name: Sup-out

Source Address:

Address Book Entry: (select), Support

Destination Address:

Address Book Entry: (select), Any

Service: Any

Action: Permit

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

Traffic Shaping: (select)

Guaranteed Bandwidth: 5000

Maximum Bandwidth: 40000

Traffic Priority: High priority

DiffServ Codepoint Marking[5]: (select)

5. Differentiated Services (DS) is a system for tagging (or “marking”) traffic at a position within a hierarchy of priority. DS Codepoint Marking maps the NetScreen priority level of the policy to the first three bits of codepoint in the DS field in the IP packet header. For more information about DS Codepoint Marking, see “Traffic Shaping” on page 215.

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Name: Sal-out

Source Address:

Address Book Entry: (select), Sales

Destination Address:

Address Book Entry: (select), Any

Service: Any

Action: Permit

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

Traffic Shaping: (select)

Guaranteed Bandwidth: 2500

Maximum Bandwidth: 40000

Traffic Priority: 2nd priority

DiffServ Codepoint Marking: Enable

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Name: Mar-out

Source Address:

Address Book Entry: (select), Marketing

Destination Address:

Address Book Entry: (select), Any

Service: Any

Action: Permit

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

Traffic Shaping: (select)

Guaranteed Bandwidth: 2500

Maximum Bandwidth: 40000

Traffic Priority: 3rd priority

DiffServ Codepoint Marking: (select)

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Name: Sup-in

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Support

Service: Any

Action: Permit

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

Traffic Shaping: (select)

Guaranteed Bandwidth: 5000

Maximum Bandwidth: 40000

Traffic Priority: High priority

DiffServ Codepoint Marking: (select)

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Name: Sal-in

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Sales

Service: Any

Action: Permit

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

Traffic Shaping: (select)

Guaranteed Bandwidth: 3500

Maximum Bandwidth: 40000

Traffic Priority: 2nd priority

DiffServ Codepoint Marking: (select)

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Name: Mar-in

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Marketing

Service: Any

Action: Permit

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

Traffic Shaping: (select)

Guaranteed Bandwidth: 1500

Maximum Bandwidth: 40000

Traffic Priority: 3rd priority

DiffServ Codepoint Marking: (select)

CLI

1. Bandwidth on Interfaces

```
set interface ethernet1 bandwidth 40000
set interface ethernet3 bandwidth 40000
```

2. Bandwidth in Policies

```
set policy name sup-out from trust to untrust support any any permit traffic
    gbw 5000 priority 0 mbw 40000 dscp enable
set policy name sal-out from trust to untrust sales any any permit traffic gbw
    2500 priority 2 mbw 40000 dscp enable
set policy name mar-out from trust to untrust marketing any any permit traffic
    gbw 2500 priority 3 mbw 40000 dscp enable
set policy name sup-in from untrust to trust any support any permit traffic gbw
    5000 priority 0 mbw 40000 dscp enable
set policy name sal-in from untrust to trust any sales any permit traffic gbw
    3500 priority 2 mbw 40000 dscp enable
set policy name mar-in from untrust to trust any marketing any permit traffic
    gbw 1500 priority 3 mbw 40000 dscp enable
save
```

11

System Parameters

This chapter focuses on the concepts involved in establishing system parameters affecting the following areas of a NetScreen security appliance:

- "Domain Name System Support" on page 495
  - "DNS Lookup" on page 496
  - "DNS Status Table" on page 497
- "DHCP" on page 500
  - "DHCP Server" on page 502
  - "DHCP Relay Agent" on page 510
  - "DHCP Client" on page 516
  - "TCP/IP Settings Propagation" on page 518
- "PPPoE" on page 521
- "Downloading/Uploading Settings and Software" on page 528
  - "Saving and Importing Settings" on page 528
  - "Uploading and Downloading Software" on page 530
  - "Configuration Rollback" on page 531
  - "Locking the Configuration File" on page 534
  - "Adding Comments to a Configuration File" on page 535
- "License Keys" on page 536
- "Registration and Activation of Signature Services" on page 538
  - "Temporary Service" on page 538
  - "AV and DI Bundled with a New Device" on page 538

  - "AV Upgrade with DI" on page 539
  - "DI Upgrade Only" on page 540
- "System Clock" on page 541
  - "Date and Time" on page 541
  - "Time Zone" on page 541
  - "NTP" on page 542

DOMAIN NAME SYSTEM SUPPORT

The NetScreen device incorporates Domain Name System (DNS) support allowing you to use domain names as well as IP addresses for identifying locations. A DNS server keeps a table of the IP addresses associated with domain names. Using DNS makes it possible to reference locations by domain name (such as www.netscreen.com) in addition to using the routable IP address, which for www.netscreen.com is 209.125.148.135. DNS translation is supported in all the following programs:

- Address Book
- Syslog
- E-mail
- WebTrends
- Websense
- LDAP
- SecurID
- RADIUS
- NetScreen Security Manager

Before you can use DNS for domain name/address resolution, you must enter the addresses for DNS servers (the primary and secondary DNS servers) in the NetScreen device.

Note: *When enabling the NetScreen device as a Dynamic Host Configuration Protocol (DHCP) server (see* *"DHCP" on page 500**), you must also enter the IP addresses for DNS servers in the DHCP page on the WebUI or through the* ***set interface*** *interface* ***dhcp*** *command in the CLI.*

DNS Lookup

The NetScreen device refreshes all the entries in its DNS table by checking them with a specified DNS server at the following times:

- After an HA failover occurs
- At a regularly scheduled time of day and at regularly scheduled intervals throughout the day
- When you manually command the device to perform a DNS lookup
  - WebUI: Network > DNS: Click Refresh DNS cache.
  - CLI: exec dns refresh

In addition to the existing method of setting a time for a daily automatic refresh of the DNS table, you can also define an interval of time from 4 hours to 24 hours.

Note: *When you add a fully-qualified domain name (FQDN) such as an address or IKE gateway through the WebUI, the NetScreen device resolves it when you click* ***Apply*** *or* ***OK****. When you type a CLI command that references an FQDN, the NetScreen device attempts to resolve it when you enter it.*

When the NetScreen device connects to the DNS server to resolve a domain name/IP address mapping, it stores that entry in its DNS status table. The following list contains some of the details involved in a DNS lookup:

- When a DNS lookup returns multiple entries, the address book accepts all entries. The other programs listed on page 495 accept only the first one.
- The NetScreen device reinstalls all policies if it finds that anything in the domain name table has changed when you refresh a lookup using the **Refresh** button in the WebUI or enter the **exec dns refresh** CLI command.
- If a DNS server fails, the NetScreen device looks up everything again.
- If a lookup fails, the NetScreen device removes it from the cache table.
- If the domain name lookup fails when adding addresses to the address book, the NetScreen device displays an error message stating that you have successfully added the address but the DNS name lookup failed.

The NetScreen device must do a new lookup once a day, which you can schedule the NetScreen device to do at a specified time:

WebUI

Network > DNS: Enter the following, and then click **Apply**:

DNS refresh every day at: Select check box and enter time <hh:mm>

CLI

```
set dns host schedule time_str
save
```

DNS Status Table

The DNS status table reports all the domain names looked up, their corresponding IP addresses, whether the lookup was successful, and when each domain name/IP address was last resolved. The report format looks like the example below:

Name	IP Address	Status	Last Lookup
www.yahoo.com	204.71.200.74 204.71.200.75 204.71.200.67 204.71.200.68	Success	8/13/2000 16:45:33
www.hotbot.com	209.185.151.28 209.185.151.210 216.32.228.18	Success	8/13/2000 16:45:38

To view the DNS status table, do either of the following:

WebUI

Network > DNS > Show DNS Table

CLI

```
get dns host report
```

Example: Defining DNS Server Addresses and Scheduling Lookups

To implement DNS functionality, the IP addresses for the DNS servers at 24.1.64.38 and 24.0.0.3 are entered in the NetScreen device, protecting a single host in a home office. The NetScreen device is scheduled to refresh the DNS settings stored in its DNS status table everyday at 11:00 P.M.

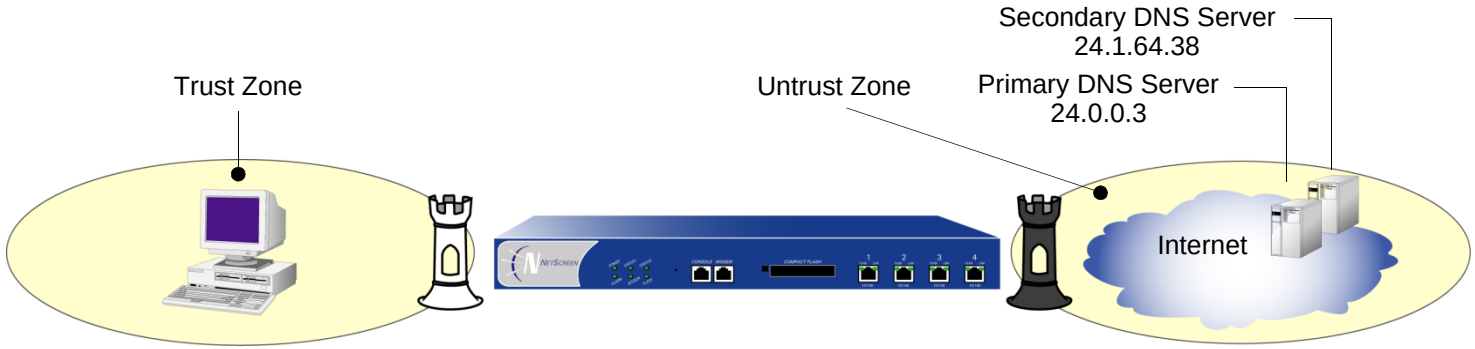

WebUI

Network > DNS: Enter the following, and then click **Apply**:

Primary DNS Server: 24.0.0.3

Secondary DNS Server: 24.1.64.38

DNS Refresh: (select)

Every Day at: 23:00

CLI

```
set dns host dns1 24.0.0.3
set dns host dns2 24.1.64.38
set dns host schedule 23:00
save
```

Example: Setting a DNS Refresh Interval

In this example, you configure the NetScreen device to refresh its DNS table every 4 hours beginning at 12:01 AM every day.

WebUI

Network > DNS: Enter the following, and then click **Apply**:

DNS Refresh: (select)

Every Day at: 12:01

Interval: 4

CLI

```
set dns host schedule 12:01 interval 4
save
```

DHCP

Dynamic Host Configuration Protocol (DHCP) was designed to reduce the demands on network administrators by automatically assigning the TCP/IP settings for the hosts on a network. Instead of requiring administrators to assign, configure, track, and change (when necessary) all the TCP/IP settings for every machine on a network, DHCP does it all automatically. Furthermore, DHCP ensures that duplicate addresses are not used, reassigns unused addresses, and automatically assigns IP addresses appropriate for the subnet on which a host is connected.

Different NetScreen devices support different DHCP roles:

- **DHCP Client:** Some NetScreen devices can act as DHCP clients, receiving a dynamically assigned IP address for any physical interface in any zone.
- **DHCP Server:** Some NetScreen devices can also act as DHCP servers, allocating dynamic IP addresses to hosts (acting as DHCP clients) on any physical or VLAN interface in any zone.

  Note: *While using the DHCP server module to assign addresses to hosts such as workstations in a zone, you can still use fixed IP addresses for other machines such as mail servers and WINS servers.*

- **DHCP Relay Agent:** Some NetScreen devices can also act as DHCP relay agents, receiving DHCP information from a DHCP server and relaying that information to hosts on any physical or VLAN interface in any zone.
- **DHCP Client/Server/Relay Agent:** Some NetScreen devices can simultaneously act as a DHCP client, server, and relay agent. Note that you can only configure one DHCP role on a single interface. For example, you cannot configure the DHCP client and server on the same interface. Optionally, you can configure the DHCP client module to forward TCP/IP settings that it receives to the DHCP server module, for use when providing TCP settings to hosts in the Trust zone acting as DHCP clients.

DHCP consists of two components: a protocol for delivering host-specific TCP/IP configuration settings and a mechanism for allocating IP addresses. When the NetScreen device acts as a DHCP server, it provides the following TCP/IP settings to each host when that host boots up:

- Default gateway IP address and netmask. If you leave these settings as 0.0.0.0/0, the DHCP server module automatically uses the IP address and netmask of the default Trust zone interface[1].
- The IP addresses of the following servers:
  - WINS servers (2):[2] A Windows Internet Naming Service (WINS) server maps a NetBIOS name used in a Windows NT network environment to an IP address used on an IP-based network.
  - NetInfo servers (2): NetInfo is an Apple network service used for the distribution of administrative data within a LAN.
  - NetInfo tag (1): The identifying tag used by the Apple NetInfo database.
  - DNS servers (3): A Domain Name System (DNS) server maps a uniform resource locator (URL) to an IP address.
  - SMTP server (1): A Simple Mail Transfer Protocol (SMTP) server delivers SMTP messages to a mail server, such as a POP3 server, which stores the incoming mail.
  - POP3 server (1): A Post Office Protocol version 3 (POP3) server stores incoming mail. A POP3 server must work conjointly with an SMTP server.
  - News server (1): A news server receives and stores postings for news groups.

Note: *If a DHCP client to which the NetScreen device is passing the above parameters has a specified IP address, that address overrides all the dynamic information received from the DHCP server.*

1. On devices that can have multiple interfaces bound to the Trust zone, the default interface is the first interface bound to that zone and assigned an IP address.
2. The number in parentheses indicates the number of servers supported.

DHCP Server

A NetScreen appliance can support up to eight DHCP servers on any physical or VLAN interface in any zone. When acting as a DHCP server, a NetScreen device allocates IP addresses and subnet masks in two modes:

- In Dynamic mode, the NetScreen device, acting as a DHCP server, assigns (or “leases”) an IP address from an address pool[3] to a host DHCP client. The IP address is leased for a determined period of time or until the client relinquishes the address. (To define an unlimited lease period, enter 0.)
- In Reserved mode, the NetScreen device assigns a designated IP address from an address pool exclusively to a specific client every time that client goes online.

> ***Note:*** *The NetScreen device saves every IP address assigned through DHCP in flash memory. Consequently, rebooting the NetScreen device does not affect address assignments.*

Example: NetScreen Device as DHCP Server

Using DHCP, the 172.16.10.0/24 network in the Trust zone is sectioned into three IP address pools.

- 172.16.10.10 through 172.16.10.19
- 172.16.10.120 through 172.16.10.129
- 172.16.10.210 through 172.16.10.219

The DHCP server assigns all IP addresses dynamically, except for two workstations with reserved IP addresses, and four servers that have static IP addresses. The interface ethernet1 is bound to the Trust zone, has IP address 172.16.10.1/24, and is in NAT mode. The domain name is dynamic.com.

3. An address pool is a defined range of IP addresses within the same subnet from which the NetScreen device can draw DHCP address assignments. You can group up to 255 IP addresses.

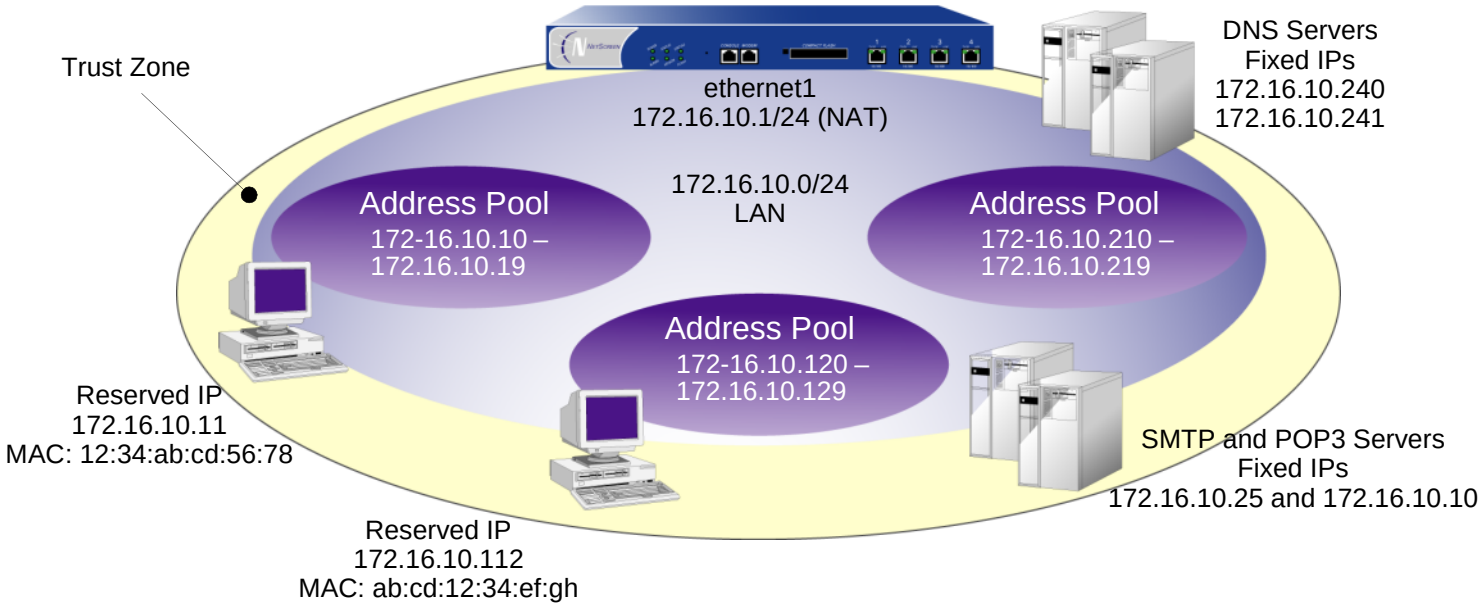

WebUI

1. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: DNS#1

Comment: Primary DNS Server

IP Address/Domain Name:

IP/Netmask: (select), 172.16.10.240/32

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: DNS#2

Comment: Secondary DNS Server

IP Address/Domain Name:

IP/Netmask: (select), 172.16.10.241/32

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: SMTP

Comment: SMTP Server

IP Address/Domain Name:

IP/Netmask: (select), 172.16.10.25/32

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: POP3

Comment: POP3 Server

IP Address/Domain Name:

IP/Netmask: (select), 172.16.10.110/32

Zone: Trust

2. DHCP Server

Network > DHCP > Edit (for ethernet1) > DHCP Server: Enter the following, and then click **Apply**:[4]

Lease: Unlimited (select)

WINS#1: 0.0.0.0

DNS#1: 172.16.10.240

> Advanced Options: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

WINS#2: 0.0.0.0

DNS#2: 172.16.10.241

DNS#3: 0.0.0.0

SMTP: 172.16.10.25

POP3: 172.16.10.110

NEWS: 0.0.0.0

NetInfo Server #1: 0.0.0.0

NetInfo Server #2: 0.0.0.0

NetInfo Tag: (leave field empty)

Domain Name: dynamic.com

> Addresses > New: Enter the following, and then click **OK**:

Dynamic: (select)

IP Address Start: 172.16.10.10

IP Address End: 172.16.10.19

4. If you leave the Gateway and Netmask fields as 0.0.0.0, the DHCP server module sends the IP address and netmask set for ethernet1 to its clients (172.16.10.1 and 255.255.255.0 in this example). However, if you enable the DHCP client module to forward TCP/IP settings to the DHCP server module (see "TCP/IP Settings Propagation" on page 518), then you must manually enter 172.16.10.1 and 255.255.255.0 in the Gateway and Netmask fields.

> Addresses > New: Enter the following, and then click **OK**:

Dynamic: (select)

IP Address Start: 172.16.10.120

IP Address End: 172.16.10.129

> Addresses > New: Enter the following, and then click **OK**:

Dynamic: (select)

IP Address Start: 172.16.10.210

IP Address End: 172.16.10.219

> Addresses > New: Enter the following, and then click **OK**:

Reserved: (select)

IP Address: 172.16.10.11

Ethernet Address: 1234 abcd 5678

> Addresses > New: Enter the following, and then click **OK**:

Reserved: (select)

IP Address: 172.16.10.112

Ethernet Address: abcd 1234 efgh

CLI

1. Addresses

```
set address trust dns1 172.16.10.240/32 “primary dns server”
set address trust dns2 172.16.10.241/32 “secondary dns server”
set address trust snmp 172.16.10.25/32 “snmp server”
set address trust pop3 172.16.10.110/32 “pop3 server”
```

2. DHCP Server

```
set interface ethernet1 dhcp server option domainname dynamic.com[5]
set interface ethernet1 dhcp server option lease 0
set interface ethernet1 dhcp server option dns1 172.16.10.240
set interface ethernet1 dhcp server option dns2 172.16.10.241
set interface ethernet1 dhcp server option smtp 172.16.10.25
set interface ethernet1 dhcp server option pop3 172.16.10.110
set interface ethernet1 dhcp server ip 172.16.10.10 to 172.16.10.19
set interface ethernet1 dhcp server ip 172.16.10.120 to 172.16.10.129
set interface ethernet1 dhcp server ip 172.16.10.210 to 172.16.10.219
set interface ethernet1 dhcp server ip 172.16.10.11 mac 1234abcd5678
set interface ethernet1 dhcp server ip 172.16.10.112 mac abcd1234efgh
set interface ethernet1 dhcp server service
save
```

5. If you do not set an IP address for the gateway or a netmask, the DHCP server module sends its clients the IP address and netmask for ethernet1 (172.16.10.1 and 255.255.255.0 in this example). However, if you enable the DHCP client module to forward TCP/IP settings to the DHCP server module (see “TCP/IP Settings Propagation” on page 518), then you must manually set these options: **set interface ethernet1 dhcp server option gateway 172.16.10.1** and **set interface ethernet1 dhcp server option netmask 255.255.255.0**.

DHCP Server in an NSRP Cluster

When the master unit in a redundant NSRP cluster functions as a DHCP server, all members in the cluster maintain all DHCP configurations and IP address assignments. In the event of a failover, the new master unit maintains all the DHCP assignments. However, termination of HA communication disrupts synchronization of existing DHCP assignments among the cluster members. After restoring HA communication, you can resynchronize the DHCP assignments by using the following CLI command on both units in the cluster: **set nsrp rto-mirror sync**.

DHCP Server Detection

When a DHCP server on a NetScreen device starts up, the system can first check to see if there is already a DHCP server on the interface. ScreenOS automatically stops the local DHCP server process from starting if another DHCP server is detected on the network. To detect another DHCP server, the device sends out DHCP boot requests at two-second intervals. If the device does not receive any response to its boot requests, it then starts its local DHCP server process.

If the NetScreen device receives a response from another DHCP server, the system generates a message indicating that the DHCP service is enabled on the NetScreen device but not started because another DHCP server is present on the network. The log message includes the IP address of the existing DHCP server.

You can set one of three operational modes for DHCP server detection on an interface: Auto, Enable, or Disable[6]. Auto mode causes the Netscreen device to always check for an existing DHCP server at bootup. You can configure the device to not attempt to detect another DHCP server on an interface by setting the NetScreen DHCP server to Enable or Disable mode. In Enable mode, the DHCP server is always on and the device does not check if there is an existing DHCP server on the network. In Disable mode, the DHCP server is always off.

6. Auto mode is the default DHCP server detection mode for NetScreen-5XP and NetScreen-5XT devices. For other NetScreen devices that support the DHCP server, Enable mode is the default DHCP server detection mode.

Example: Turning on DHCP Server Detection

In this example, you set the DHCP server on the ethernet1 interface to check for an existing DHCP server on the interface first before starting up.

WebUI

Network > DHCP > Edit (for ethernet1) > DHCP Server: Enter the following, and then click **OK**:

Server Mode: Auto (select)

CLI

```
set interface ethernet1 dhcp server auto
save
```

Example: Turning off DHCP Server Detection

In this example, you set the DHCP server on the ethernet1 interface to start up without checking to see if there is an existing DHCP server on the network.

WebUI

Network > DHCP > Edit (for ethernet1) > DHCP Server: Enter the following, and then click **OK**:

Server Mode: Enable (select)

CLI

```
set interface ethernet1 dhcp server enable
save
```

Note: *Issuing the CLI command* ***set interface*** interface ***dhcp server service*** *command activates the DHCP server. If the DHCP server detection mode for the interface is set to Auto, the DHCP server on the NetScreen device starts only if it does not find an existing server on the network. Issuing the* ***unset interface*** interface ***dhcp server service*** *command disables the DHCP server on the NetScreen device and also deletes any existing DHCP configuration.*

DHCP Relay Agent

When acting as a DHCP relay agent, the NetScreen device forwards DHCP requests and assignments between hosts in one zone and a DHCP server in another zone. The DHCP messages between the NetScreen device and the DHCP server can be transmitted in the open or through a VPN tunnel.

You can configure up to three DHCP servers for the DHCP relay agent. The relay agent unicasts an address request from a DHCP client to all configured DHCP servers. The relay agent forwards to the client the first response received from a server.

Note: *When a NetScreen device acts as a DHCP relay agent, the NetScreen device does not generate DHCP allocation status reports because the remote DHCP server controls all the IP address allocations.*

The following simplified illustration presents the process involved in using a NetScreen device as a DHCP relay agent. Note that to ensure security, the DHCP messages pass through a VPN tunnel when traveling over the untrusted network.

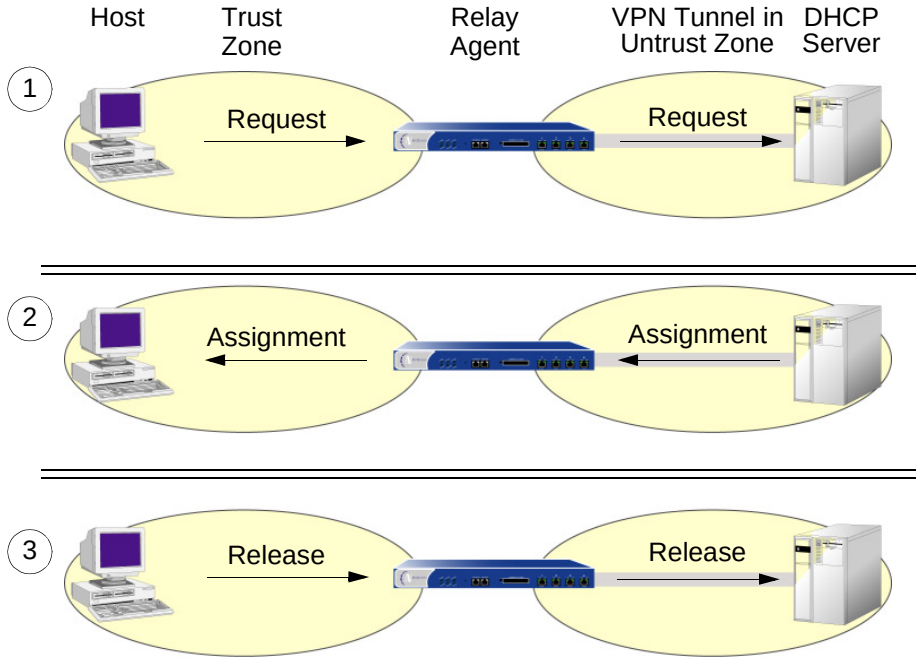

When the NetScreen device functions as a DHCP relay agent, its interfaces must be in either Route mode or Transparent mode. For interfaces in Route mode, you must configure a policy from one zone to another zone for the predefined service DHCP-Relay. For interfaces in Transparent mode, the DHCP client must reside in the V1-Trust zone, while the DHCP server can reside in either the V1-Untrust or V1-DMZ zone. No policy is needed for interfaces in Transparent mode.

You can configure the DHCP relay agent on any physical or VLAN interface. You cannot configure DHCP relay agent and DHCP server or client functions on the same interface.

Example: NetScreen Device as DHCP Relay Agent

In this example, a NetScreen device receives its DHCP information from a DHCP server at 194.2.9.10 and relays it to hosts in the Trust zone. The hosts receive IP addresses from an IP pool defined on the DHCP server. The address range is 180.10.10.2—180.10.10.254. The DHCP messages pass through a VPN tunnel between the local NetScreen device and the DHCP server, located behind a remote NetScreen device whose Untrust zone interface IP address is 2.2.2.2/24. The interface ethernet1 is bound to the Trust zone, has the IP address 180.10.10.1/24, and is in Route mode. The interface ethernet3 is bound to the Untrust zone and has the IP address 1.1.1.1/24. All security zones are in the trust-vr routing domain.

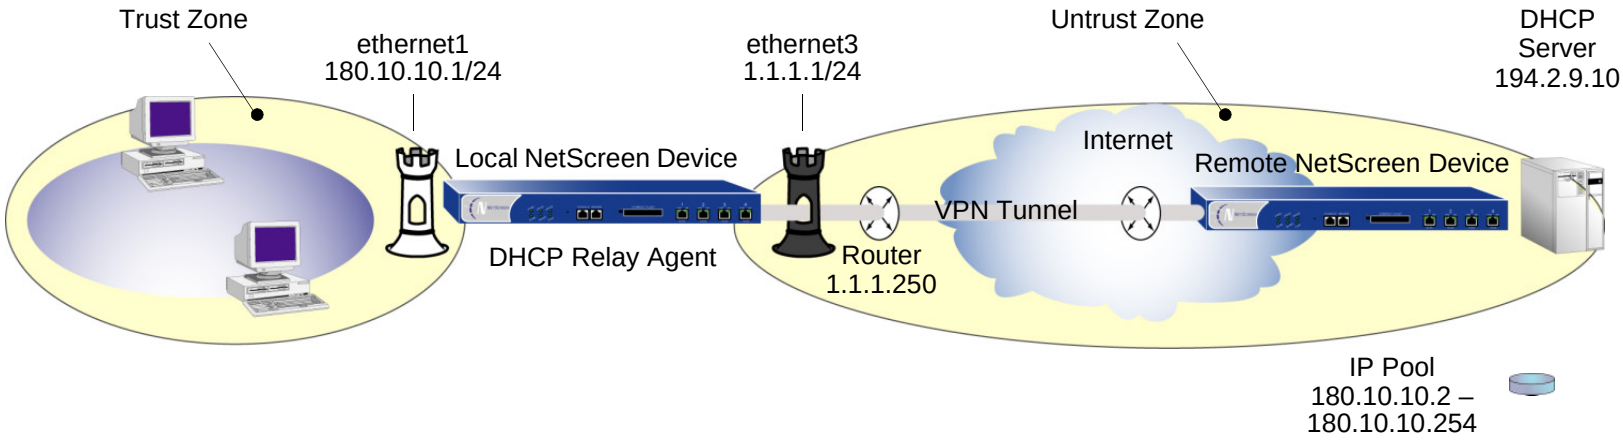

WebUI

1. Interfaces

Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone: Trust

Static IP: (select this option when present)

IP Address/Netmask: 180.10.10.1/24

Enter the following, and then click **OK**:

Interface Mode: Route

Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

2. Address

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: DHCP Server

IP Address/Domain Name:

IP/Netmask: (select), 194.2.9.10/32

Zone: Untrust

3. VPN

VPNs > AutoKey Advanced > Gateway > New: Enter the following, and then click **OK**:

Gateway Name: dhcp server

Security Level: Custom

Remote Gateway Type:

Static IP: (select), Address/Hostname: 2.2.2.2

Outgoing Interface: ethernet3

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

Security Level:

User Defined: Custom (select)

Phase1 Proposal: rsa-g2-3des-sha

Mode (Initiator): Main (ID Protection)

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: to_dhcp

Security Level: Compatible

Remote Gateway:

Predefined: (select), to_dhcp

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

Bind to: None

4. DHCP Relay Agent

Network > DHCP > Edit (for ethernet1) > DHCP Relay Agent: Enter the following, and then click **Apply**:

Relay Agent Server IP or Domain Name: 194.2.9.10

Use Trust Zone Interface as Source IP for VPN: (select)

5. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 1.1.1.250[7]

6. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), DHCP Server

Service: DHCP-Relay

Action: Tunnel

Tunnel VPN: to_dhcp

Modify matching outgoing VPN policy: (select)

7. Setting a route to the external router designated as the default gateway is essential for both outbound VPN and network traffic. In this example, the NetScreen device sends encapsulated VPN traffic to this router as the first hop along its route to the remote NetScreen device. In the illustration for this example, the concept is presented by depicting the tunnel passing through the router.

CLI

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 180.10.10.1/24
set interface ethernet1 route
set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
```

2. Address

```
set address untrust dhcp_server 194.2.9.10/32
```

3. VPN

```
set ike gateway “dhcp server” ip 2.2.2.2 main outgoing-interface ethernet3
    proposal rsa-g2-3des-sha
set vpn to_dhcp gateway “dhcp server” proposal g2-esp-3des-sha
```

4. DHCP Relay Agent

```
set interface ethernet1 dhcp relay server-name 194.2.9.10
set interface ethernet1 dhcp relay vpn
```

5. Route

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
```

6. Policies

```
set policy from trust to untrust any dhcp_server dhcp-relay tunnel vpn to_dhcp
set policy from untrust to trust dhcp_server any dhcp-relay tunnel vpn to_dhcp
save
```

DHCP Client

When acting as a DHCP client, the NetScreen device receives an IP address dynamically from a DHCP server for any physical interface in any security zone. If there are multiple interfaces bound to a single security zone, you can configure a DHCP client for each interface as long as each interface is not connected to the same network segment. If you configure a DHCP client for two interfaces that are connected to the same network segment, the first address assigned by a DHCP server is used. (If the DHCP client receives an address update to the same IP address, IKE is not rekeyed.)

Note: *While some NetScreen devices can act as a DHCP server, DHCP relay agent, or a DHCP client at the same time, you cannot configure more than one DHCP role on a single interface.*

Example: NetScreen Device as DHCP Client

In this example, the interface bound to the Untrust zone has a dynamically assigned IP address. When the NetScreen device requests its IP address from its ISP, it receives its IP address, subnet mask, gateway IP address, and the length of its lease for the address. The IP address of the DHCP server is 2.2.2.5.

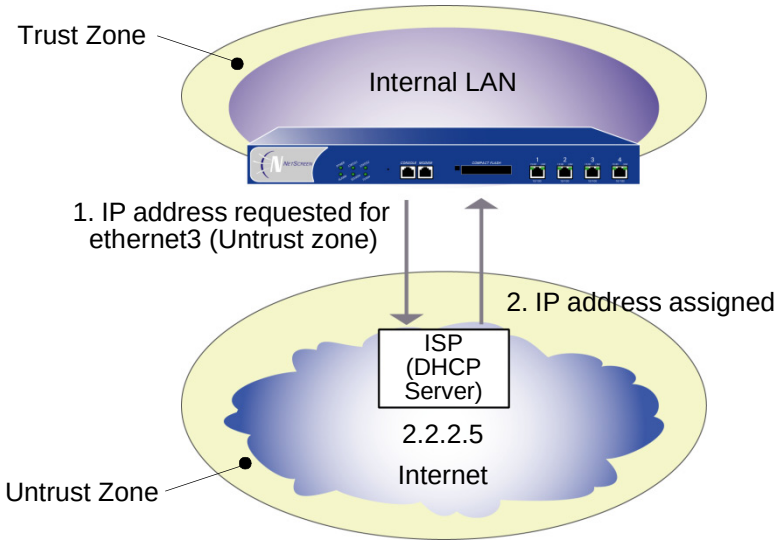

Note: *Before setting up a site for DHCP service, you must have the following:*

- *Digital subscriber line (DSL) modem and line*
- *Account with ISP*

WebUI

Network > Interfaces > Edit (for ethernet3): Select **Obtain IP using DHCP**[8], and then click **OK**.

CLI

```
set interface ethernet3 dhcp client
set interface ethernet3 dhcp settings server 2.2.2.5
save
```

8. You cannot specify the IP address of the DHCP server through the WebUI; however, you can do so through the CLI.

TCP/IP Settings Propagation

Some NetScreen devices can act as a Dynamic Host Control Protocol (DHCP) client, receiving its TCP/IP settings and the IP address for any physical interface in any security zone from an external DHCP server. Some NetScreen devices can act as a DHCP server, providing TCP/IP settings and IP addresses to clients in any zone. When a NetScreen device acts both as a DHCP client and a DHCP server simultaneously, it can transfer the TCP/IP settings learned through its DHCP client module to its default DHCP server module[9]. TCP/IP settings include the IP address of the default gateway and a subnet mask, and IP addresses for any or all of the following servers:

- DNS (3)
- WINS (2)
- NetInfo (2)
- SMTP (1)
- POP3 (1)
- News (1)

The NetScreen device is both a client of the DHCP server in the Untrust zone and a DHCP server to the clients in the Trust zone.

It takes the TCP/IP settings that it receives as a DHCP client and forwards them as a DHCP server to the clients in the Trust zone.

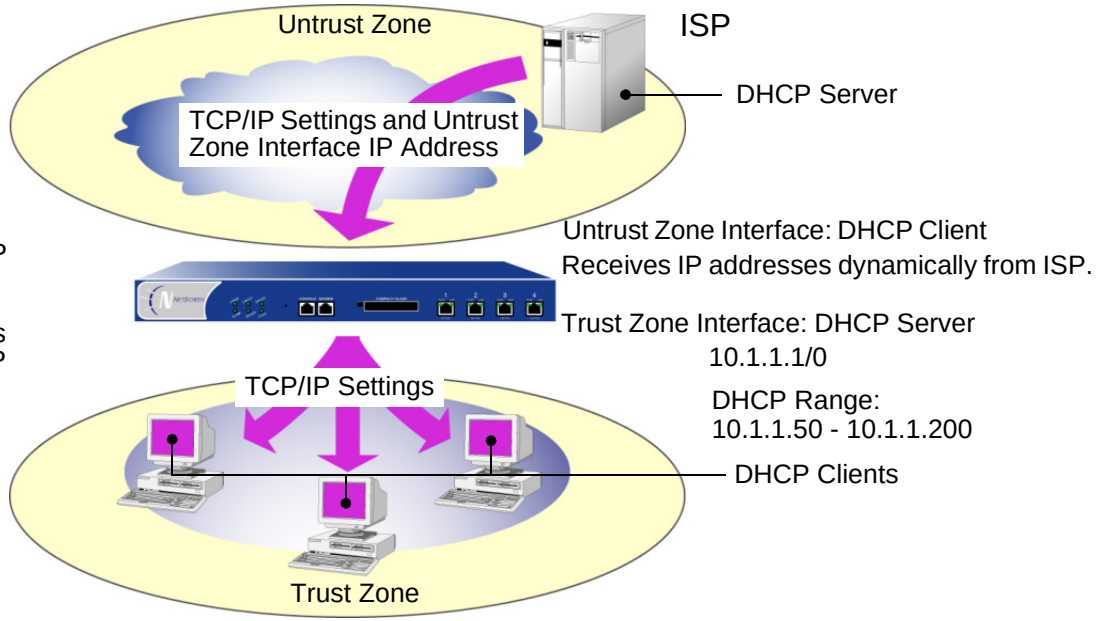

9. While you can configure up to eight DHCP servers on any physical or VLAN interface, the default DHCP server on the device resides on a specific interface on each platform. On the NetScreen-5XP, the default DHCP server resides on the Trust interface. On the NetScreen-5XT, the default DHCP server resides on the Trust interface for Trust-Untrust port mode, the ethernet1 interface for Dual-Untrust port mode, and the ethernet2 interface for Home-Work and Combined port modes. For other devices, the default DHCP server resides on the ethernet1 interface.

You can configure the DHCP server module to propagate all TCP/IP settings that it receives from the DHCP client module using the **set interface** *interface* **dhcp-client settings update-dhcpserver** command. You can also override any setting with a different one.

Example: Forwarding TCP/IP Settings

In this example, you configure the NetScreen device to act both as a DHCP client on the ethernet3 interface and as a DHCP server on the ethernet1 interface. (The default DHCP server is on the ethernet1 interface.)

As a DHCP client, the NetScreen device receives an IP address for the ethernet3 interface and its TCP/IP settings from an external DHCP server at 211.3.1.6. You enable the DHCP client module in the NetScreen device to transfer the TCP/IP settings it receives to the DHCP server module.

You configure the NetScreen DHCP server module to do the following with the TCP/IP settings that it receives from the DHCP client module:

- Forward the DNS IP addresses to its DHCP clients in the Trust zone.
- Override the default gateway[10], netmask, and SMTP server and POP3 server IP addresses with the following:
  - 10.1.1.1 (this is the IP address of the ethernet1 interface)
  - 255.255.255.0 (this is the netmask for the ethernet1 interface)
  - SMTP: 211.1.8.150
  - POP3: 211.1.8.172

You also configure the DHCP server module to deliver the following TCP/IP settings that it does not receive from the DHCP client module:

- Primary WINS server: 10.1.2.42
- Secondary WINS server: 10.1.5.90

Finally, you configure the DHCP server module to assign IP addresses from the following IP Pool to the hosts acting as DHCP clients in the Trust zone: 10.1.1.50 – 10.1.1.200.

10. If the DHCP server is already enabled on the Trust interface and has a defined pool of IP addresses (which is default behavior on some NetScreen devices), you must first delete the IP address pool before you can change the default gateway and netmask.

WebUI

> **Note:** *You can only set this feature through the CLI.*

CLI

1. DHCP Client

```
set interface ethernet3 dhcp-client settings server 211.3.1.6
set interface ethernet3 dhcp-client settings update-dhcpserver
set interface ethernet3 dhcp-client settings autoconfig
set interface ethernet3 dhcp-client enable
```

2. DHCP Server

```
set interface ethernet1 dhcp server option gateway 10.1.1.1
set interface ethernet1 dhcp server option netmask 255.255.255.0
set interface ethernet1 dhcp server option wins1 10.1.2.42
set interface ethernet1 dhcp server option wins2 10.1.5.90
set interface ethernet1 dhcp server option pop3 211.1.8.172
set interface ethernet1 dhcp server option smtp 211.1.8.150
set interface ethernet1 dhcp server ip 10.1.1.50 to 10.1.1.200
set interface ethernet1 dhcp server service
save
```

PPPoE

PPP-over-Ethernet (PPPoE) merges the Point-to-Point Protocol (PPP), which is usually used for dialup connections, with the Ethernet protocol, which can connect multiple users at a site to the same customer premises equipment. While many users can share the same physical connection, access control, billing, and type of service are handled on a per-user basis. Some NetScreen devices support a PPPoE client, allowing them to operate compatibly on DSL, Ethernet Direct, and cable networks run by ISPs using PPPoE for their clients' Internet access.

On devices that support PPPoE, you can configure a PPPoE client instance on any or all interfaces. You configure a specific instance of PPPoE with a user name and password and other parameters, and bind the instance to an interface. When there are two Ethernet interfaces (a primary and a backup) bound to the Untrust zone, you can configure one or both interfaces for PPPoE. For example, in Dual Untrust port mode[11], you can configure the primary interface (ethernet3) for DHCP and the backup interface (ethernet2) for PPPoE. Or, you can configure PPPoE for both the primary and backup interfaces.

Example: Setting Up PPPoE

The following example illustrates how to define the untrusted interface of a NetScreen device for PPPoE connections, and how to initiate PPPoE service.

In this example, the NetScreen device receives a dynamically assigned IP address for its Untrust zone interface (ethernet3) from the ISP, and the NetScreen device also dynamically assigns IP addresses for the three hosts in its Trust zone. In this case, the NetScreen device acts both as a PPPoE client and a DHCP server. The Trust zone interface must be in either NAT mode or Route mode. In this example, it is in NAT mode.

11. Port modes are supported on some NetScreen appliances, such as the NetScreen-5XT.

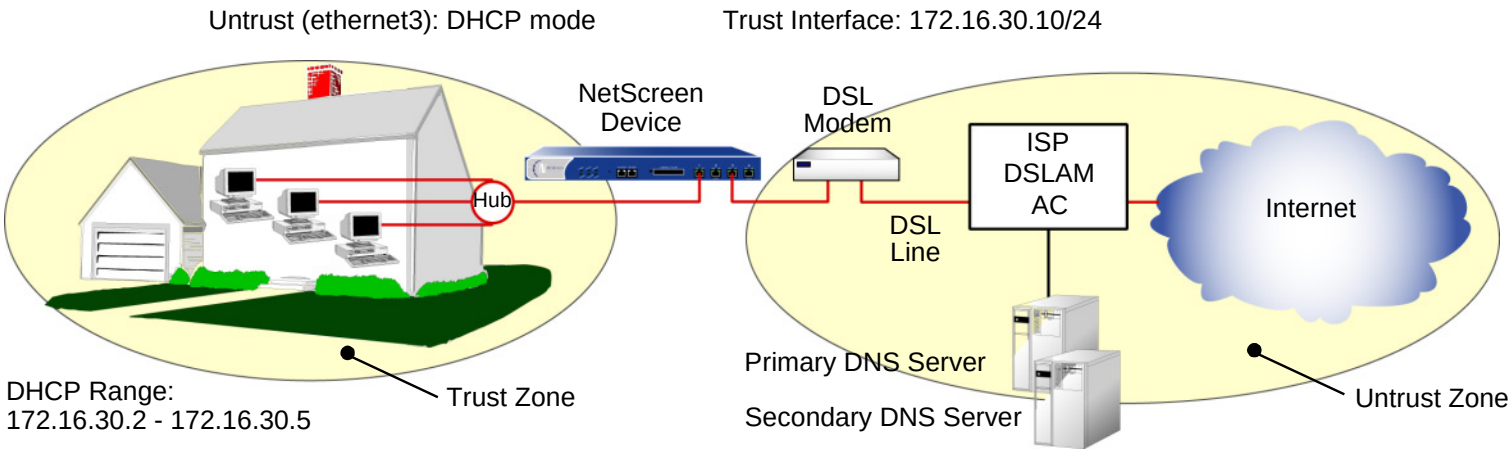

Before setting up the site in this example for PPPoE service, you must have the following:

- Digital subscriber line (DSL) modem and line
- Account with ISP
- User name and password (obtained from the ISP)

WebUI

1. Interfaces and PPPoE

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **OK**:

Zone: Trust

Static IP: (select this option when present)

IP Address/Netmask: 172.16.30.10/24

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone: Untrust

Obtain IP using PPPoE: (select)

User Name/Password: <name>/<password>

Network > Interfaces > Edit (for ethernet3): To test your PPPoE connection, click **Connect**.

Note: *When you initiate a PPPoE connection, your ISP automatically provides the IP addresses for the Untrust zone interface and the IP addresses for the Domain Name Service (DNS) servers. When the NetScreen device receives DNS addresses via PPPoE, the new DNS settings overwrite the local settings by default. If you do not want the new DNS settings to replace the local settings, you can use the CLI command* ***unset pppoe dhcp-updateserver*** *to disable this behavior.*

If you use a static IP address for the Untrust zone interface, you must obtain the IP addresses of the DNS servers and manually enter them on the NetScreen device and on the hosts in the Trust zone.

2. DHCP Server

Network > Interfaces > Edit (for ethernet1) > DHCP: Select **DHCP Server**, and then click **Apply**.

Network > Interfaces > Edit (for ethernet1) > DHCP: Enter the following, and then click **Apply**:

Lease: 1 hour

Gateway: 0.0.0.0

Netmask: 0.0.0.0

DNS#1: 0.0.0.0

> Advanced: Enter the following, and then click **Return**:

DNS#2: 0.0.0.0

Domain Name: (leave blank)

Network > Interfaces > DHCP (for ethernet1) > New Address: Enter the following, and then click **OK**:

Dynamic: (select)

IP Address Start: 172.16.30.2

IP Address End: 172.16.30.5

3. Activating PPPoE on the NetScreen Device

Turn off the power to the DSL modem, the NetScreen device, and the three workstations.

Turn on the DSL modem.

Turn on the NetScreen device.

The NetScreen device makes a PPPoE connection to the ISP and, through the ISP, gets the IP addresses for the DNS servers.

4. Activating DHCP on the Internal Network

Turn on the workstations.

The workstations automatically receive the IP addresses for the DNS servers. They get an IP address for themselves when they attempt a TCP/IP connection.

> ***Note:*** *When you use DHCP to assign IP addresses to hosts in the Trust zone, the NetScreen device automatically forwards the IP addresses of the DNS servers that it receives from the ISP to the hosts.*
>
> *If the IP addresses for the hosts are not dynamically assigned through DHCP, you must manually enter the IP addresses for the DNS servers on each host.*

Every TCP/IP connection that a host in the Trust zone makes to the Untrust zone automatically goes through the PPPoE encapsulation process.

CLI

1. Interfaces and PPPoE

```
set interface ethernet1 zone trust
set interface ethernet1 ip 172.16.30.10/24
set interface ethernet3 zone untrust
set pppoe interface ethernet3
set pppoe username name_str password pswd_str
```

To test your PPPoE connection:

```
exec pppoe connect
get pppoe
```

2. DHCP Server

```
set interface ethernet1 dhcp server service
set interface ethernet1 dhcp server ip 172.16.30.2 to 172.16.30.5
set interface ethernet1 dhcp server option lease 60
save
```

3. Activating PPPoE on the NetScreen Device

Turn off the power to the DSL modem, the NetScreen device, and the three workstations.

Turn on the DSL modem.

Turn on the NetScreen device.

4. Activating DHCP on the Internal Network

Turn on the workstations.

The workstations automatically receive the IP addresses for the DNS servers. They get an IP address for themselves when they attempt a TCP/IP connection.

Every TCP/IP connection that a host in the Trust zone makes to the Untrust zone automatically goes through the PPPoE encapsulation process.

Example: Configuring PPPoE on Primary and Backup Untrust Interfaces

For this example, the NetScreen-5XT is in Dual Untrust mode. In the following example, you configure PPPoE for both the primary (ethernet3) and backup (ethernet2) interfaces to the Untrust zone.

WebUI

PPPoE Configuration for ethernet3 Interface

Network > PPPoE > New: Enter the following, and then click **OK**:

PPPoE instance: eth3-pppoe

Bound to interface: ethernet3 (select)

Username: user1

Password: 123456

Authentication: Any (select)

Access Concentrator: ac-11

PPPoE Configuration for ethernet2 Interface

Network > PPPoE > New: Enter the following, and then click **OK**:

PPPoE instance: eth2-pppoe

Bound to interface: ethernet2 (select)

Username: user2

Password: 654321

Authentication: Any (select)

Access Concentrator: ac-22

CLI

1. PPPoE Configuration for ethernet3 Interface

```
set pppoe name eth3-pppoe username user1 password 123456
set pppoe name eth3-pppoe ac ac-11
set pppoe name eth3-pppoe authentication any
set pppoe name eth3-pppoe interface ethernet3
```

2. PPPoE Configuration for ethernet2 Interface

```
set pppoe name eth2-pppoe username user2 password 654321
set pppoe name eth2-pppoe ac ac-22
set pppoe name eth2-pppoe authentication any
set pppoe name eth2-pppoe interface ethernet2
save
```

DOWNLOADING/UPLOADING SETTINGS AND SOFTWARE

You can upload and download configuration settings and software to and from a NetScreen device. The kinds of location that you upload from and download to depend on whether you use the WebUI or the CLI to perform the operation. Using the WebUI and Web browser support, you can upload and download configuration settings and upload ScreenOS software from any local directory. Through the CLI, you can upload and download settings and software from and to a TFTP server or PC card.

Saving and Importing Settings

It is good practice to backup your settings after every significant change you make. Through the WebUI, you can download the configuration to any local directory as a backup precaution. With some NetScreen devices, you can use the CLI to download the configuration to a TFTP server or flash card. Should you need the saved backup configuration, you can then simply upload it to the NetScreen device.

The ability to download and upload a configuration also provides the means for mass distribution of configuration templates.

To download a configuration:

WebUI

1. Configuration > Update > Config File: Click **Save to File**.

   A system message prompts you to open the file or save it to your computer.
2. Click **Save**.
3. Browse to the location where you want to save the configuration file, and then click **Save**.

CLI

```
save config from flash to { tftp ip_addr | slot } filename [ from interface ]
```

Note: *On some NetScreen devices, you must specify slot 1 or slot 2.*

To upload a configuration:

WebUI

Configuration > Update > Config File: Enter the following, and then click **Apply**:

Select **Merge to Current Configuration** if you want to combine both the new and the current configurations, or **Replace Current Configuration** if you want the new configuration to overwrite the current configuration.

> New Configuration File: Enter the configuration file location or click **Browse** and navigate to the file location, select the file, and then click **Open**.

CLI

```
save config from { tftp ip_addr | slot } filename to flash [ merge [ from
    interface ] ]
```

Note: *On some NetScreen devices, you must specify slot 1 or slot 2.*

Uploading and Downloading Software

When a new NetScreen ScreenOS version becomes available, you can purchase it and download it from the NetScreen download site. Then you can upload the new software using the **save** command, or use the WebUI to upload software from a local directory. Through the CLI, you can upload the software from a TFTP server or PC card, and you can download software to a TFTP server.

Note: *After the software is upgraded, the NetScreen device reboots. This process takes a few minutes.*

WebUI

Configuration > Update > ScreenOS/Keys: Enter the following, and then click **Apply**:

Select what you want to update: Firmware, Image Key, or License Key.

> Load File: Enter the location of the file you want to update or click **Browse** and navigate to the file location, select the file, and then click **Open**.

CLI

```
save software from { flash | slot1 filename | tftp ip_addr filename ) to flash
```

Note: *On some NetScreen devices, you must specify slot 1 or slot 2.*

Through the CLI, you can also download software to a TFTP server, using the **save** command:

```
save software from flash to tftp ip_addr filename [ from interface ]
```

Configuration Rollback

In the event that you load a configuration file that causes problems, such as the failure of the NetScreen device or remote users losing the ability to manage the device, you can perform a configuration rollback to revert to a last-known-good configuration file that you previously saved in flash memory. We refer to the reverted configuration as the LKG (Last Known Good) configuration.

Note: *Not all NetScreen devices support configuration rollback. To see if your NetScreen device supports this feature, please refer to the relevant data sheet for your platform.*

Last-Known-Good Configuration

Before performing a configuration rollback, make sure you have a LKG configuration file saved in flash memory to which the NetScreen device can revert. To do this, use the **get config rollback** CLI command. The filename for a LKG configuration is *lkg.cfg*. If you do not see this file, it means that it does not exist and you must create it.

To save a configuration file to flash as the last-known-good:

1. Ensure that the current configuration on the NetScreen device is good.
2. Save the current configuration to flash memory using the **save config to last-known-good** CLI command. This command overwrites the existing LKG configuration in flash memory with the current configuration file.

Regularly saving the configuration on the NetScreen device as the LKG configuration file is a good way to backup your latest changes to the configuration and maintain an up-to-date copy of the configuration.

Automatic and Manual Configuration Rollback

You can enable the NetScreen device to revert automatically to the last-known-good (LKG) configuration, or you can perform the rollback manually. The automatic configuration rollback feature enables the NetScreen device to rollback to the LKG configuration in case of a problem with a newly loaded configuration.

The automatic configuration rollback feature is disabled by default. Furthermore, it is disabled after every startup, regardless of whether it was enabled or disabled before starting up the device. To enable automatic configuration rollback, use the **exec config rollback enable** command. To disable the feature, use the **exec config rollback disable** command.

To perform a manual configuration rollback, use the **exec config rollback** command.

Note: *The WebUI does not support the configuration rollback feature.*

After you enable the configuration rollback feature, the command prompt changes to indicate this state:

```
ns-> exec config rollback enable
ns(rollback enabled)->
```

When you disable the configuration rollback feature, the command prompt returns to just the device host name:

```
ns(rollback enabled)-> exec config rollback disable
ns->
```

To verify that the automatic configuration rollback feature is enabled, use the **get config rollback** command. If it is enabled, the first line of the **get config rollback** output is:

```
config rollback is enabled
```

Otherwise, the first line of the output is:

```
config rollback is disabled
```

If an LKG configuration file exists, the second line of the **get config rollback** output is:

```
Last-known-good config file flash:/$lkg$.cfg exists in the flash.
```

Following this line, the size and contents of the file are displayed.

If an LKG configuration file does not exist, the second—and final—line of the output is:

```
Last-known-good config file flash:/$lkg$.cfg does not exist.
```

When the configuration rollback feature is enabled, you can trigger the rollback operation by any of the following actions:

- Restarting the NetScreen device (by turning the power off and then on again)
- Resetting the NetScreen device (by entering the **reset** command)
- Entering the **exec config rollback** command

Loading a New Configuration File

The following describes how to load a new configuration file, enable the configuration rollback feature, and what to do in case the new configuration file causes problems.

1. Save the current configuration as the LKG using the **Save config to last-known-good** CLI command.
2. Enable automatic configuration rollback on the NetScreen device using the **exec config rollback enable** CLI command. Enabling this feature simultaneously locks the LKG file to prevent other users from overwriting it, and consequently disrupting an ongoing configuration rollback.
3. Load the new configuration file using the WebUI or CLI. For more information, see "Uploading and Downloading Software" on page 530.
4. Test the new configuration file by issuing commands. A few scenarios can occur:
   - The new configuration is running correctly.
   - The new configuration is defective and as a result, you can no longer reach and manage the NetScreen device. In this case, you have to power off the device. When you power it on, the NetScreen device reads the flash memory file, which indicates that the configuration rollback feature is enabled. That information prompts the NetScreen device to load the LKG file automatically.
   - You notice problems with or errors in the new configuration file. In this case, you need to reset the NetScreen device using the **reset** CLI command. When the device restarts, it reads the flash memory file, which indicates that the configuration rollback feature is enabled. That information prompts the NetScreen device to load the LKG file automatically.
   - The new configuration is defective and causes the NetScreen device to become inoperable. In this case, the NetScreen device restarts automatically. Upon restart, it reads the flash memory file, which indicates that the configuration rollback feature is enabled. That information prompts the NetScreen device to load the LKG file automatically.

Note: *NSRP—In an active/active setup, if loading a new configuration file fails, both NetScreen devices revert to the LKG. In an active/passive setup, if loading a new configuration file fails, only the master unit reverts to the LKG. Only after you save the configuration to file does the master unit synchronize the backup unit.*

Locking the Configuration File

You can lock a configuration file in flash memory to prevent it from being overwritten by other admins or before importing a new configuration file. When you lock the configuration file, the device starts a lock timer. If the device does not receive a CLI command within a previously specified lockout period, it automatically restarts, using the configuration that was locked in flash memory. It is good practice to lock the current configuration of the device before you start importing a configuration file. This prevents the device from freezing for an indefinite period of time due to a failure in the import process.

When you lock the configuration file, you and any other admin connected to the device (for example, through telnet or the WebUI) cannot save to the configuration file. You must first unlock the configuration, then save the new configuration commands with the **save** command.

Note: *You can lock/unlock a configuration file through the CLI only. This feature is not available on the WebUI.*

CLI

To lock the configuration file:

exec config lock start

To unlock the file:

exec config lock end

To abort the lockout and immediately reboot the device with the configuration that was previously locked in flash:

exec config lock abort

To change the default lockout period (5 minutes):

set config lock timeout <*number*>

Adding Comments to a Configuration File

You can add comments to an external configuration file. The comments can be in a separate line of text or at the tail end of one line. The comment must begin with # (the number sign) and be followed by a space. When the comment is at the tail end of a line, a space must also precede the number sign. When you save the file onto a NetScreen device—either by merging the new configuration with the existing one or by completely replacing the existing configuration with the new one—the device parses the configuration for lines beginning with the number sign and removes any comments that it finds.

Note: *If the number sign appears within quotation marks, the NetScreen device does not treat it as a special marker but as part of an object name and does not remove it. For example, the NetScreen device does not delete* ***"#5 server"*** *in the command* ***set address trust "#5 server" 10.1.1.5/32*** *because it appears within quotation marks.*

The NetScreen device does not save any comments introduced with the number sign in either RAM of flash memory. For example, if an external configuration file contains the following lines:

```
set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24 # change IP address
# add new MIP addresses
set interface ethernet3 mip 1.1.1.10 host 10.1.1.10 netmask 255.255.255.255
set interface ethernet3 mip 1.1.1.11 host 10.1.1.11 netmask 255.255.255.255
set interface ethernet3 mip 1.1.1.12 host 10.1.1.12 netmask 255.255.255.255
# all MIPs use the trust-vr routing domain by default
```

When you view the configuration after you load the file, you see the following (the comments are gone):

```
set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
set interface ethernet3 mip 1.1.1.10 host 10.1.1.10 netmask 255.255.255.255
set interface ethernet3 mip 1.1.1.11 host 10.1.1.11 netmask 255.255.255.255
set interface ethernet3 mip 1.1.1.12 host 10.1.1.12 netmask 255.255.255.255
```

Also, if you paste a block of commands that includes comments into a console or Telnet session, the NetScreen device discards all comments immediately upon running the commands.

LICENSE KEYS

The license key feature allows you to expand the capabilities of your NetScreen device without having to upgrade to a different device or system image. You can purchase a key that unlocks specified features already loaded in the software, such as the following:

- User capacity
- Virtual systems, zones, and virtual routers
- HA

Each NetScreen device ships with a standard set of features enabled and might support the activation of optional features or the increased capacity of existing features. For information regarding which features are currently available for upgrading, refer to the latest marketing literature from NetScreen.

The procedure for obtaining and applying a license key is as follows:

1. Contact the value-added reseller (VAR) who sold you the NetScreen device or contact NetScreen Technologies directly.
2. Provide the serial number of your device and state the feature option you want.

   The license key is generated and then sent to you via e-mail.
3. Enter the key through either the WebUI or CLI. (See the following example.)

Example: Expanding User Capacity

A small company using a single NetScreen device with a license for 10 users has grown to the point where it now needs an unrestricted user license. The NetScreen administrator expands the capabilities of the device by obtaining a software key for an unrestricted number of users. The license key number is 6a48e726ca050192 and is in a text file named “A2010002.txt” located at C:\netscreen\keys.

WebUI

Configuration > Update > ScreenOS/Keys: Do the following, and then click **Apply**:

License Key Update: (select)

Load File: C:\netscreen\keys\A2010002.txt

Or

Click **Browse** and navigate to C:\netscreen\keys, select A2010002.txt, and then click **Open**.

CLI

```
exec license-key capacity 6a48e726ca050192
reset
```

REGISTRATION AND ACTIVATION OF SIGNATURE SERVICES

Before your NetScreen device can receive regular signature service for AV (antivirus) patterns or DI (Deep Inspection) signatures, you must purchase a subscription to the service, register for the service, and then retrieve the subscription. Retrieving the subscription activates your services on the device. How the service activation process works depends upon the way you purchased the services, and what the services are.

Temporary Service

To allow you time to for AV or DI services, the NetScreen device provides a temporary grace period. During this period, the device can obtain services on a temporary basis.

- No NetScreen device comes with DI already enabled. To obtain temporary DI service, you must start a WebUI session and click the Retrieve Subscriptions Now in the Configuration > Update > ScreenOS/Keys page. This provides a one-time, one-day DI key.
- If your device has AV service bundled at time of purchase, the device has pre-installed temporary service.

Warning! *To avoid service interruption, you must perform registration as soon as possible after purchasing your subscription. Registration ensures continuation of the subscription.*

AV and DI Bundled with a New Device

If you purchased a new NetScreen device that already has the AV and DI services, perform the following steps to activate the services.

1. Configure the device for internet connectivity. (For instructions, refer to the *Getting Started* sheet and the user's guide for your NetScreen device.)
2. Register the device at the following site:

   www.netscreen.com/cso

   Devices with bundled AV services come with a temporary, pre-installed subscription, so you can go ahead and use the service immediately. However, you *must* register the device to receive your full paid subscription.

3. Retrieve the subscription on the device. You can do this either of two ways:
   - In WebUI, click Retrieve Subscriptions Now from the Configuration > Update > ScreenOS/Keys page.
   - Using the CLI, run the following command:

     ```
     exec license-key update
     ```

You can now configure the device to automatically retrieve or manually retrieve the signature services. For instructions on configuring your NetScreen device for these services, refer to "Deep Inspection Overview" on page 124, and "Antivirus Scanning" on page 76.

AV Upgrade with DI

If you purchase AV and DI services separately from the NetScreen device, perform the following steps to activate the services.

1. After ordering the services, you should receive a support certificate, via e-mail, from NetScreen or your authorized NetScreen reseller. This certificate is a readable document that contains information you need to register your device.
2. Make sure the device is registered. If it is not currently registered, go to the following site:

   www.netscreen.com/cso
3. Register the support certificate to the device.
4. If you are subscribing and registering for the DI service only, go on to Step 5 immediately.

   If you are subscribing and registering for the AV service, you must wait up to four hours for the system to process the registration before proceeding with Step 5.
5. Confirm that your device has internet connectivity.
6. Retrieve the subscription on the device. You can do this either of two ways:
   - In WebUI, click Retrieve Subscriptions Now from the Configuration > Update > ScreenOS/Keys page.
   - Using the CLI, run the following command:

     ```
     exec license-key update
     ```

You can now configure the device to automatically retrieve or manually retrieve the signature services. For instructions on configuring your NetScreen device for these services, refer to "Deep Inspection Overview" on page 124, and "Antivirus Scanning" on page 76.

DI Upgrade Only

If you purchased DI services only, and you purchased your NetScreen device separately from the DI service, perform the following steps to activate the service.

1. After ordering the service, you should receive a support certificate, via e-mail, from NetScreen or your authorized NetScreen reseller. This certificate is a readable document that contains information you need to register your device.
2. Make sure the device is registered. If it is not currently registered, go to the following site:

   www.netscreen.com/cso
3. Register the support certificate to the device. It may be necessary to wait up to four hours for the system to process the registration before proceeding.
4. Confirm that your device has internet connectivity.
5. Retrieve the subscription on the device. You can do this either of two ways:
   - In WebUI, click Retrieve Subscriptions Now from the Configuration > Update > ScreenOS/Keys page.
   - Using the CLI, run the following command:

     ```
     exec license-key update
     ```

You can now configure the device to automatically retrieve or manually retrieve the DI signature service. For instructions on configuring your NetScreen device for this service, refer to "Deep Inspection Overview" on page 124, and "Antivirus Scanning" on page 76.

SYSTEM CLOCK

It is important that your NetScreen device always be set to the right time. Among other things, the time on your NetScreen device affects the set up of VPN tunnels and the timing of schedules. There are many ways that you can ensure that the NetScreen device always maintains the accurate time. First, you must set the system clock to the current time. Next, you can enable the daylight saving time option and you can configure up to three NTP servers (one primary and two backups) from which the NetScreen device can regularly update its system clock.

Date and Time

To set the clock to the current time and date, you can use the WebUI or the CLI. Through the WebUI, you do this by synchronizing the system clock with the clock on your computer:

1. Configuration > Date/Time: Click the **Sync Clock with Client** button.

   A pop-up message prompts you to specify if you have enabled the daylight saving time option on your computer clock.

2. Click **Yes** to synchronize the system clock and adjust it according to daylight saving time or **No** to synchronize the system clock without adjusting it for daylight saving time.

Through the CLI, you set the clock by manually entering the date and time using this command "**set clock mm/dd/yyyy hh:mm:ss**".

Time Zone

You set the time zone by specifying the number of hours by which the local time for the NetScreen device is behind or ahead of GMT (Greenwich Mean Time). For example, if the local time zone for the NetScreen device is Pacific Standard Time, it is 8 hours behind GMT. Therefore, you have to set the clock to **-8**.

If you set the time zone using the WebUI:

Configuration > Date/Time > Set Time Zone_hours_minutes from GMT

If you set the time zone using the CLI:

ns -> set clock timezone *number* (a number from -12 to 12)

or

ns-> set ntp timezone *number* (a number from -12 to 12)

NTP

To ensure that the NetScreen device always maintains the right time, it can use NTP (Network Time Protocol) to synchronize its system clock with that of an NTP server over the Internet. You can do this manually or configure the NetScreen device to perform this synchronization automatically at time intervals that you specify.

Multiple NTP Servers

You can configure up to three NTP servers on a NetScreen device: one primary server and two backup servers. When you configure the NetScreen device to synchronize its system clock automatically, it queries each configured NTP server sequentially. The device always queries the primary NTP server first. If the query is not successful, the device then queries the first backup NTP server and so on until it gets a valid reply from one of the NTP servers configured on the NetScreen device. The device makes four attempts on each NTP server before it terminates the update and logs the failure.

When you manually synchronize the system clock, and you can only do this using the CLI, you can specify a particular NTP server or none at all. If you specify a NTP server, the NetScreen device queries that server only. If you do not specify a NTP server, the NetScreen device queries each NTP server configured on the NetScreen device sequentially. You can specify a NTP server using its IP address or its domain name.

Maximum Time Adjustment

For automatic synchronization, you can specify a maximum time adjustment value (in seconds). The maximum time adjustment value represents the acceptable time difference between the NetScreen device system clock and the time received from an NTP server. The NetScreen device only adjusts its clock with the NTP server time if the time difference between its clock and the NTP server time is within the maximum time adjustment value that you set. For example, if the maximum time adjustment value is 3 seconds, and the time on the device system clock is 4:00:00

and the NTP server sends 4:00:02 as the time, the difference in time between the two is acceptable and the NetScreen device can update its clock. If the time adjustment is greater than the value you set, the NetScreen device does not synchronize its clock and proceeds to try the first backup NTP server configured on the NetScreen device. If the NetScreen device does not receive a valid reply after trying all the configured NTP servers, it generates an error message in the event log.

The default value for this feature is 3 seconds and the range is 0 (no limit) to 3600 (one hour).

When you manually synchronize the system clock, and you can only do this using the CLI, the NetScreen device does not verify the maximum time adjustment value. Instead, if it receives a valid reply, the NetScreen device displays a message informing you of which NTP server it reached, what the time adjustment is, and the type of authentication method used. The message also asks you to confirm or cancel the system clock update.

If the NetScreen device does not receive a reply, it displays a timeout message. This message appears only after the device unsuccessfully attempted to reach all NTP servers configured on the NetScreen device.

Note: *When issuing requests using the CLI, you can cancel the current request by pressing Ctrl-C on the keyboard.*

NTP and NSRP

The NetScreen Redundancy Protocol (NSRP) contains a mechanism for synchronizing the system clock of NSRP cluster members. Although the resolution for synchronization is in seconds, NTP has sub-second resolution. Because the time on each cluster member might differ by a few seconds due to processing delays, NetScreen recommends that you disable NSRP time synchronization when NTP is enabled on both cluster members and they can each update their system clock from an NTP server. To disable the NSRP time synchronization function, enter the following command:

```
set ntp no-ha-sync
```

Example: Configuring NTP Servers and a Maximum Time Adjustment Value

In the following example you configure the NetScreen device to update its clock at five-minute intervals from NTP servers at IP addresses 1.1.1.1, 1.1.1.2, and 1.1.1.3. You also set a maximum time adjustment value of 2 seconds.

WebUI

Configuration > Date/Time: Enter the following, and then click **Apply**:

Automatically synchronize with an Internet Time Server (NTP): (select)

Update system clock every minutes: 5

Maximum time adjustment seconds: 2

Primary Server IP/Name: 1.1.1.1

Backup Server1 IP/Name: 1.1.1.2

Backup Server2 IP/Name: 1.1.1.3

CLI

1. set clock ntp
2. set ntp server 1.1.1.1
3. set ntp server backup1 1.1.1.2
4. set ntp server backup2 1.1.1.3
5. set ntp interval 5
6. set ntp max-adjustment 2
7. save

Secure NTP Servers

You can secure NTP traffic by using MD5-based checksum to provide authentication of NTP packets. You do not need to create an IPSec tunnel. This type of authentication ensures the integrity of NTP traffic. It does not prevent outside parties from viewing the data, but prevents anyone from tampering with it.

To enable the authentication of NTP traffic, you must assign a unique key id and preshared key to each NTP server you configure on a NetScreen device. The key id and preshared key serve to create a checksum, with which the NetScreen device and the NTP server can authenticate the data.

Authentication Types

There are two types of authentication for NTP traffic: required and preferred.

When you select **Required** authentication, the NetScreen device must include the authentication information—key id and checksum—in every packet it sends to a NTP server and must authenticate all NTP packets it receives from a NTP server. Before authentication can occur between a NetScreen device and a NTP server, the administrators of the NetScreen device and the NTP server must first exchange a key id and a preshared key. They have to exchange these manually and can do so in different ways such as via e-mail or telephone.

When you select **Preferred** authentication, the NetScreen device must first operate as in Required mode by trying to authenticate all NTP traffic. If all attempts to authenticate fail, the NetScreen device then operates as if you selected no authentication. It sends out packets to a NTP server without including a key id and checksum. Essentially, although there is a preference for authentication, if authentication fails, the NetScreen device still permits NTP traffic.

Index

A

access policies
    See policies
address book
    adding addresses 127
    editing group entries 132
    entries 127
    groups 129
    modifying addresses 128
    removing addresses 133
    See also addresses
address groups 129, 206
    creating 131
    editing 132
    options 130
    removing entries 133
address negation 237
address translation
    See NAT, NAT-dst, and NAT-src
addresses
    address book entries 127
    defined 206
    in policies 206
    private 78
    public 77
admin users 465–466
    auth process 466
    privileges from RADIUS 465
    server support 372
    timeout 378
aggregate interfaces 67
alarms
    thresholds 213
ALG 159
ALGs
    for custom services 207
antivirus scanning
    policies 214
application, in policies 207
Application-Layer Gateway
    See ALG
ARP 95
    ingress IP address 98
auth servers 372
    address 377
    authentication process 376
    backup servers 377
    default 395
    defining 388–396
    external 376
    feature support 372
    ID number 377
    in IKE gateways 396
    in policies 396
    LDAP 386–387
    LDAP, defining 393
    maximum number 373
    multiple user types 373
    object name 377
    object properties 377
    RADIUS 379–381
    RADIUS, defining 388
    RADIUS, user type support 380
    SecurID 384–385
    SecurID, defining 391
    timeout 377
    types 377
    user type support 372
    XAuth queries 437
auth users 398–430
    groups 398, 402
    in policies 398
    point of authentication 397
    pre-policy auth 212, 400
    run-time (external user group) 412
    run-time (external user) 409
    run-time (local user group) 406
    run-time (local user) 403
    run-time auth process 211, 399
    run-time authentication 211, 399
    server support 372
    timeout 377
    WebAuth 212, 400
    WebAuth (external user group) 423
    WebAuth (local user group) 420
    WebAuth + SSL (external user group) 427
authentication
    Allow Any 212
    policies 210
    users 210, 371–476
    WebAuth 400
authentication, users 371–476
    accounts 371
    admin 465
    auth servers 372
    auth users 398
    IKE users 372, 431
    L2TP users 460
    local database 374–375
    Manual Key users 372
    multiple-type 467
    point of authentication 397
    profiles 371
    types and applications 397–467
    user types 372
    WebAuth 372
    with different logins 467
    XAuth users 436

B

bandwidth 215
    default priority 485
    guaranteed 215, 478, 486
    managing 478
    maximum 215, 486
    maximum specification 478
    priority levels 485
    priority queues 485
    unlimited maximum 478
banners, customizing 476

C

CHAP 453
character types, ScreenOS supported xiv
CLI
  conventions x
  set vip multi -port 358
clock, system 541–545
  See also system clock
common name 387
configuration
  adding comments 535
  LKG 531
  loading 533
  locking 534
  rollback 531–532, 533
configuration settings
  downloading 528
  uploading 528
conventions
  CLI x
  illustration xiii
  names xiv
  WebUI xi
counting 213
creating
  address groups 131
  MIP addresses 333
  service groups 168
  zones 51

D

defining
  zones 51
DHCP 115, 121, 521
  client 500
  HA 508
  relay agent 500
  server 500
dictionary file 465
DiffServ 215
  See DS Codepoint Marking
DIP 119, 171–174
  fix-port 173
  groups 189–192
  modifying a DIP pool 174
  PAT 172
  pools 210
DIP pools
  address considerations 259
  NAT-src 246
  size 259
distinguished name 387
DNS 495
  lookup 496
  server 523
  status table 497
Domain name system
  See DNS
DS Codepoint Marking 478, 487, 488
DSL 517, 522
Dynamic IP pools
  See DIP pools
dynamic routing 30

E

editing
  address groups 132
  policies 241
  zones 52

F

Function Zone Interfaces 68
  HA Interface 69
  Management Interface 68

G

gatekeeper devices 141
global zone 359
graphs, historical 213
group
  addresses 129
  services 167
group expressions 468–475
  operators 468
  other group expressions 469
  server support 372
  user groups 468
  users 468

H

H.323 protocol 141
HA
  DHCP 508
  Virtual HA Interface 69
  See also NSRP
High Availability
  See HA
historical graphs 213
Home zone 61

I

ICMP services 139
  message code 139
  message type 139
icons
  defined 216
  policy 216
idle session timeout 377
IKE
  IKE ID 431, 452
  user groups, defining 434
  users 431–435
  users, defining 432
  users, groups 431
IKE users
  IKE ID 397, 431
  server support 372
  with other use types 467
illustration
  conventions xiii
interfaces
  addressing 77
  aggregate 67
  binding to zone 76
  default 79
  DIP 171
  HA 69
  L3 security zones 77
  loopback 86
  MGT 68
  MIP 331
  modifying 81
  physical 3
  redundant 67

secondary IP address 84
tunnel 49, 69, 70–73
unbinding from zone 80
viewing interface table 74
VIP 356
Virtual HA 69
VSI 68
IP addresses
defining for each port 127
host ID 78
L3 security zones 77–78
network ID 78
private 77
private address ranges 78
public 77
secondary 84
virtual 356
IP pools
See DIP pools

L

L2TP
address assignment 460
external auth server 461
local database 461
policies 209
user authentication 460
L2TP users 460–464
point of authentication 397
server support 372
with XAuth 467
Last-Known-Good configuration
See LKG configuration
LDAP 386–387
auth server object 393
common name identifier 387
distinguished name 387
server port 387
structure 386
user types supported 387
license keys 536–537
Lightweight Directory Access Protocol
See LDAP
LKG (last-known-good) 531
LKG configuration 531
local database 374–375
IKE users 431
timeout 375
user types supported 374
logging 213
loopback interfaces 86

M

Management interface
See MGT interface
mapped IP
See MIP
MGT interface 68
MIP 12, 331
address range 335
bidirectional translation 252
creating addresses 333
creating on tunnel interface 341
creating on zone interface 333
default netmask 335
default virtual router 335
definition 252
global zone 332
reachable from other zones 336
same-as-untrust interface 342–345
to zone with interface-based NAT 112
multimedia sessions, SIP 156
multiple-type users 467

N

names
conventions xiv
NAT
definition 246
NAT-src with NAT-dst 310–330
NAT mode 110–117
interface settings 113
traffic to Untrust zone 91, 112
NAT-dst 276–330
address range 250
address range to address range 256, 300
address range to single IP 256, 295
address shifting 251, 277, 300
one-to-many translation 291
one-to-one translation 286
packet flow 278–281
port mapping 249, 276, 305
route considerations 277, 282–285
single IP with port mapping 255
single IP, no port mapping 255
unidirectional translation 252, 257
with MIPs or VIPs 248
NAT-src 246, 259–275
address shifting 267–272
address shifting, range considerations 267
DIP pool with address shifting 254
DIP pool with PAT 253, 260–263
DIP pool, fixed port 253
DIP pools 246
egress interface 254, 273–275
fixed port 259, 264–266
interface based 247
port address translation 247
Route mode Route mode
NAT-src 118
unidirectional translation 252, 257
negation, address 237
NetInfo 501
netmasks 206
uses of 78
NetScreen dictionary file 381
network, bandwidth 478
NSRP
configuration rollback 533
DHCP 508
DIP groups 189–192
HA session backup 212
NTP synchronization 543
redundant interfaces 67
VSIs 68
NTP 542–545
authentication types 545
max time adjustment 542
maximum time adjustment 542
multiple servers 542
NSRP synchronization 543
secure servers 545
server configuration 544
servers 542

P

packet flow 11–13
NAT-dst 278–281
PAT 172, 259
PC card 528, 530
pinholes 161
policies 3
actions 207
address groups 206
address negation 237
addresses 206
addresses in 206
alarms 213
antivirus scanning 214
application 207
authentication 210
bidirectional VPNs 208, 216
changing 241
counting 213
Deep Inspection 209
deny 207
DIP groups 190
disabling 241
enabling 241
functions of 197
global 201, 217, 234
HA session backup 212
icons 216
ID 206
internal rules 203
interzone 200, 217, 218, 223
intrazone 201, 217, 231
L2TP 209
L2TP tunnels 209
location 218
lookup sequence 202
management 216
managing bandwidth 478
maximum limit 130
multiple items per component 236
name 208
NAT-dst 210
NAT-src 210
order 243
permit 207
policy context 235
policy set lists 202
policy verification 242
position at top 209, 243
removing 244
reordering 243
required elements 199
root system 203
schedules 214
security zones 206
service book 134
service groups 167
services 206
services in 134, 206
shadowing 242
traffic logging 213
traffic shaping 215
tunnel 207
types 200–201
URL filtering 213
virtual systems 203
VPN dialup user groups 206
VPNs 208
policy-based NAT
See NAT-dst and NAT-src
tunnel interfaces 69
Port Address Translation
See PAT
port mapping 249, 276
port modes 55
ports
port numbers 366
priority queuing 485
private addresses 78
public addresses 77

Q

QoS 478

R

RADIUS 379–381
auth server object 388
NetScreen dictionary file 465
object properties 380
port 380
retry timeout 380
shared secret 380
Remote authentication dial in user service
See RADIUS
RFCs
1349, “Type of Service in the Internet Protocol Suite” 215
1777, “Lightweight Directory Access Protocol” 386
1918, “Address Allocation for Private Internets” 78
2474, “Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers” 215
rollback, configuration 531–532
Route mode 118–123
interface settings 119
routing 30
between secondary IP addresses 84
routing table 31
RSH ALG 140
rules, derived from polices 203
run-time authentication 211, 399

S

schedules 193, 214
SCREEN
MGT zone 48
ScreenOS
function zones 54
global zone 48
Home-Work zone 61
interfaces physical 3
overview 1–27
packet flow 11–13
policies 3
port modes 55
security zone interfaces 3
security zones 2, 48
security zones, global 2
security zones, predefined 2
subinterfaces 4
tunnel zones 49
updating 530

virtual systems 10
zones 45–54
SDP 159–160
secondary IP addresses 84
SecurID 384–385
ACE server 384
auth server object 391
authentication port 385
authenticator 384
client retries 385
client timeout 385
duress 385
encryption type 385
token code 384
user type support 385
security zones 2
destination zone determination 13
global 2
interfaces 3, 66
physical interfaces 66
predefined 2
source zone determination 12
subinterfaces 66
service book
adding service 136
custom service 134
custom service (CLI) 136
modifying entries (CLI) 138
modifying entries (Web UI) 169
pre-configured services 134
removing entries (CLI) 138
service groups (Web UI) 167
service groups 167–170
creating 168
deleting 170
modifying 169
services 134
custom ALGs 207
defined 206
drop-down list 134
ICMP 139
in policies 206
modifying timeout 136
timeout threshold 135
Session Initiation Protocol
See SIP
session timeout
idle timeout 377
settings
downloading 528
importing 528
saving 528
uploading 528
shadowed policies 242
SIP 156–166
ALG 159, 163
connection information 160
defined 156
inactivity timeouts 163
media announcements 160
media inactivity timeout 163, 166
messages 156
multimedia sessions 156
pinholes 159
request method types 157
Request Methods 157
response codes 158
response types 157
responses 157
RTCP 160
RTP 160
SDP 159–160
session inactivity timeout 163
signaling 159
signaling inactivity timeout 163, 166
software
updating 530
uploading and downloading 530
SSL
with WebAuth 427
static routing 30, 33–44
configuring 38
using 36
subinterfaces 4
creating (root system) 82
deleting 83
subscriptions
registration and activation 538–540
temporary service 538
support certificate 539, 540
system clock 541–545
date & time 541
sync with client 541
time zone 541
system, parameters 493–544

T

TFTP server 528, 530
time zone 541
timeout
admin user 378
auth user 377
token code 384
trace-route 98, 101
traffic
counting 213
logging 213
priority 215
shaping 478
traffic shaping 477–491
automatic 478
interface requirement 478
service priorities 485
Transparent mode 92–109
ARP/trace-route 96
blocking non-ARP traffic 94
blocking non-IP traffic 94
broadcast traffic 94
flood 96
routes 94
unicast options 96
tunnel interfaces 69
definition 69
policy-based NAT 69

U

unknown unicast options 95–101
ARP 98–101
flood 96–97
trace-route 98, 101
URL filtering 213
user authentication
See authentication, users

users
groups, server support 372
IKE 431–435
IKE, groups 434
users, admin 465–466
auth process 466
timeout 378
users, IKE
defining 432
groups 431
IKE ID 431
users, L2TP 460–464
users, XAuth 436–458

V

vendor-specific attributes
See VSAs
VIP 12
bidirectional translation 252
configuring 359
custom and multi-port services 363–369
custom services, low port numbers 357
definition 252
editing 362
global zone 359
reachable from other zones 359
removing 362
required information 357
to zone with interface-based NAT 112
virtual adapter 436
Virtual HA interface 69
Virtual IP
See VIP
virtual routers
See VRs
virtual system 10
VLAN zone 93
VLAN1
Interface 93, 102
Zones 93
VLANs
tags 4
voice-over IP communication 141
VPNs
idletime 439
policies 208
to zone with interface-based NAT 112
tunnel zones 49
VRs 35
forwarding traffic between 5
introduction 5
VSAs 381
attribute name 381
attribute number 381
attribute type 381
vendor ID 381

W

WebAuth 372
external user group 423
local user group 420
pre-policy auth process 212, 400
with SSL (external user group) 427
WebUI
conventions xi
Work zone 61

X

XAuth
address assignments 436, 438
address timeout 438
auth and address 452
client authentication 458
defined 436
external auth server queries 437
external user auth 444
external user group auth 447
IP address lifetime 438–439
lifetime 439
local user auth 440
local user group auth 442
query remote settings 437
ScreenOS as client 458
TCP/IP assignments 437
user authentication 436
virtual adapter 436
VPN idletime 439
XAuth users 436–458
point of authentication 397
server support 372
with L2TP 467

Z

zones 45–54
function 54
global 48, 359
Layer 2 93
security 48
tunnel 49
VLAN 54, 93

NetScreen Concepts & Examples

ScreenOS Reference Guide

Volume 3: Administration

ScreenOS 5.0.0

P/N 093-0926-000

Rev. B

Copyright Notice

NetScreen Technologies, Inc.
Building #3
805 11th Avenue
Sunnyvale, CA 94089
www.netscreen.com

FCC Statement

The following information is for FCC compliance of Class A devices: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. The equipment generates, uses, and can radiate radio-frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case users will be required to correct the interference at their own expense.

The following information is for FCC compliance of Class B devices: The equipment described in this manual generates and may radiate radio-frequency energy. If it is not installed in accordance with NetScreen's installation instructions, it may cause interference with radio and television reception. This equipment has been tested and found to comply with the limits for a Class B digital device in accordance with the specifications in part 15 of the FCC rules. These specifications are designed to provide reasonable protection against such interference in a residential installation. However, there is no guarantee that interference will not occur in a particular installation.

If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Consult the dealer or an experienced radio/TV technician for help.
- Connect the equipment to an outlet on a circuit different from that to which the receiver is connected.

Caution: Changes or modifications to this product could void the user's warranty and authority to operate this device.

Disclaimer

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR NETSCREEN REPRESENTATIVE FOR A COPY.

Contents

Preface v
Conventions vi
CLI Conventions vi
WebUI Conventions vii
Illustration Conventions ix
Naming Conventions and Character Types x
NetScreen Documentation xi
Chapter 1 Administration 1
Management via the Web User Interface 3
WebUI Help 4
Copying the Help Files to a Local Drive 4
Pointing the WebUI to the New Help Location 4
HTTP 5
Session ID 5
Secure Sockets Layer 7
Management via the Command Line Interface 9
Telnet 9
Securing Telnet Connections 10
Secure Shell 11
Client Requirements 13
Basic SSH Configuration on the NetScreen Device 13
Authentication 15
SSH and Vsys 17
Host Key 18
Example: SSHv1 with PKA for Automated Logins 19
Secure Copy (SCP) 20
Serial Console 21
Modem Port 22
Management via NetScreen-Security Manager 23
Initiating Connectivity Between Agent and Management System 24
Enabling and Disabling the Agent 25
Example: Enabling the Security Manager Agent 25
Changing Management System Server Address 26
Example: Setting the Primary Server IP Address 26
Setting Report Parameters 26
Example: Enabling Alarm and Statistics Reporting 27
Controlling Administrative Traffic 29
MGT and VLAN1 Interfaces 30
Example: Administration through the MGT Interface 30
Example: Administration through the VLAN1 Interface 31
Administrative Interface 32
Example: Setting Administrative Interface Options 32
Manage IP 34
Example: Setting Manage IPs for Multiple Interfaces 34

Levels of Administration 37
Root Administrator 37
Read/Write Administrator 38
Read-Only Administrator 38
Virtual System Administrator 38
Virtual System Read-Only Administrator 39
Defining Admin Users 39
Example: Adding a Read-Only Admin 39
Example: Modifying an Admin 40
Example: Deleting an Admin 40
Example: Clearing an Admin's Sessions 41
Securing Administrative Traffic 42
Changing the Port Number 43
Example: Changing the Port Number 43
Changing the Admin Login Name and Password 44
Example: Changing an Admin User's Login Name and Password 45
Example: Changing One's Own Password 46
Setting the Minimum Length of the Root Admin Password 47
Resetting the Device to the Factory Default Settings 48
Restricting Administrative Access 49
Example: Restricting Administration to a Single Workstation 49
Example: Restricting Administration to a Subnet 50
Restricting the Root Admin to Console Access 50
VPN Tunnels for Administrative Traffic 51
Example: Administration through a Route-Based Manual Key VPN Tunnel 52
Example: Administration through a Policy-Based Manual Key VPN Tunnel 58
Chapter 2 Monitoring NetScreen Devices 65
Storing Log Information 66
Event Log 67
Viewing the Event Log 68
Example: Viewing the Event Log by Severity Level and Keyword 69
Sorting and Filtering the Event Log 70
Example: Sorting Event Log Entries by IP Address 70
Downloading the Event Log 71
Example: Downloading the Event Log 71
Example: Downloading the Event Log for Critical Events 71
Traffic Log 72
Viewing the Traffic Log 74
Example: Viewing Traffic Log Entries 74
Sorting and Filtering the Traffic Log 75
Example: Sorting the Traffic Log by Time 75
Downloading the Traffic Log 76
Example: Downloading a Traffic Log 76
Self Log 77
Viewing the Self Log 77
Sorting and Filtering the Self Log 78
Example: Filtering the Self Log by Time 79
Downloading the Self Log 80
Example: Downloading the Self Log 80
Asset Recovery Log 81
Example: Downloading the Asset Recovery Log 81

Traffic Alarms 82
Example: Policy-Based Intrusion Detection 83
Example: Compromised System Notification 84
Example: Sending E-mail Alerts 86
Syslog 87
Example: Enabling Multiple Syslog Servers 88
WebTrends 89
Example: Enabling Syslog and WebTrends for Notification Events 89
SNMP 91
Implementation Overview 94
Example: Defining a Read/Write SNMP Community 95
VPN Tunnels for Self-Initiated Traffic 97
Example: Self-Generated Traffic through a Route-Based Tunnel 99
Example: Self-Generated Traffic through a Policy-Based Tunnel 109
Counters 120
Example: Viewing Screen Counters 126
Appendix A SNMP MIB Files ..... A-I
Index ..... IX-I

Preface

NetScreen devices provide different ways for you to manage the devices, either locally or remotely. Volume 3, “Administration” describes the various methods for managing NetScreen devices and explains ScreenOS administrative levels. This volume also describes how to secure local and remote administration of NetScreen devices, and how to monitor device activity. An appendix contains brief descriptions of the NetScreen Management Information Base (MIB) files that support communications between NetScreen devices and SNMP management applications.

CONVENTIONS

This document contains several types of conventions, which are introduced in the following sections:

- "CLI Conventions"
- "WebUI Conventions" on page vii
- "Illustration Conventions" on page ix
- "Naming Conventions and Character Types" on page x

CLI Conventions

The following conventions are used when presenting the syntax of a command line interface (CLI) command:

- Anything inside square brackets [ ] is optional.
- Anything inside braces { } is required.
- If there is more than one choice, each choice is separated by a pipe (|). For example,

  ```
  set interface { ethernet1 | ethernet2 | ethernet3 } manage
  ```

  means "set the management options for the ethernet1, ethernet2, or ethernet3 interface".
- Variables appear in *italic.* For example:

  ```
  set admin user name password
  ```

When a CLI command appears within the context of a sentence, it is in **bold** (except for variables, which are always in *italic*). For example: "Use the **get system** command to display the serial number of a NetScreen device."

Note: *When typing a keyword, you only have to type enough letters to identify the word uniquely. For example, typing* ***set adm u joe j12fmt54*** *is enough to enter the command* ***set admin user joe j12fmt54****. Although you can use this shortcut when entering commands, all the commands documented here are presented in their entirety.*

WebUI Conventions

Throughout this book, a chevron (>) is used to indicate navigation through the WebUI by clicking menu options and links. For example, the path to the address configuration dialog box is presented as **Objects > Addresses > List > New**. This navigational sequence is shown below.

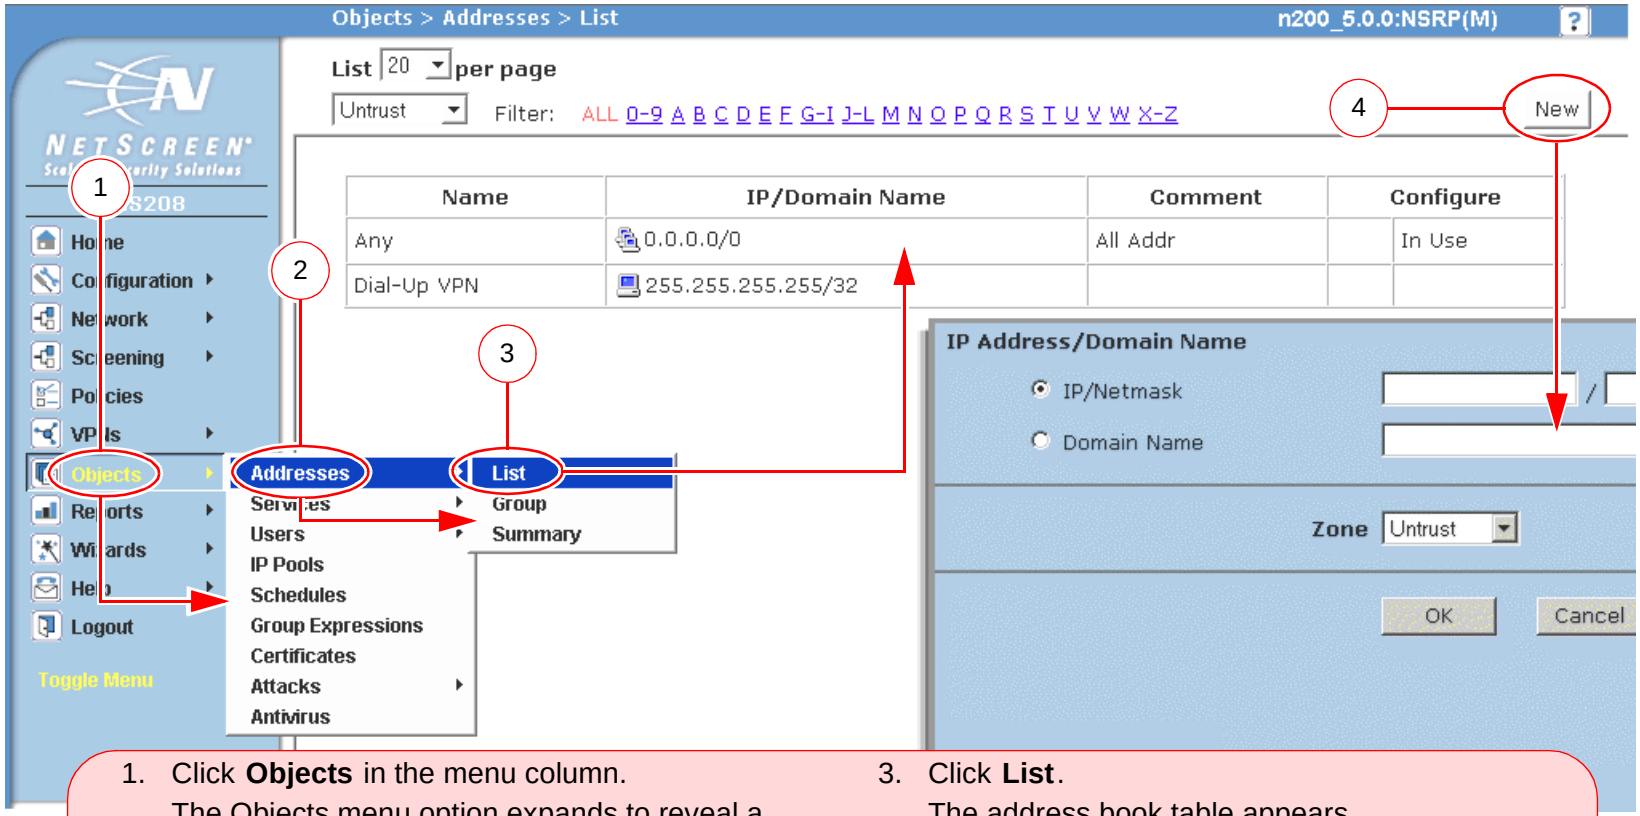

1. Click **Objects** in the menu column.
   The Objects menu option expands to reveal a subset of options for Objects.
2. (Applet menu) Hover the mouse over **Addresses**.
   (DHTML menu) Click **Addresses**.
   The Addresses option expands to reveal a subset of options for Addresses.
3. Click **List**.
   The address book table appears.
4. Click the **New** link.
   The new address configuration dialog box appears.

To perform a task with the WebUI, you must first navigate to the appropriate dialog box where you can then define objects and set parameters. The set of instructions for each task is divided into two parts: a navigational path and configuration details. For example, the following set of instructions includes the path to the address configuration dialog box and the settings for you to configure:

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: addr_1

IP Address/Domain Name:

IP/Netmask: (select), 10.2.2.5/32

Zone: Untrust

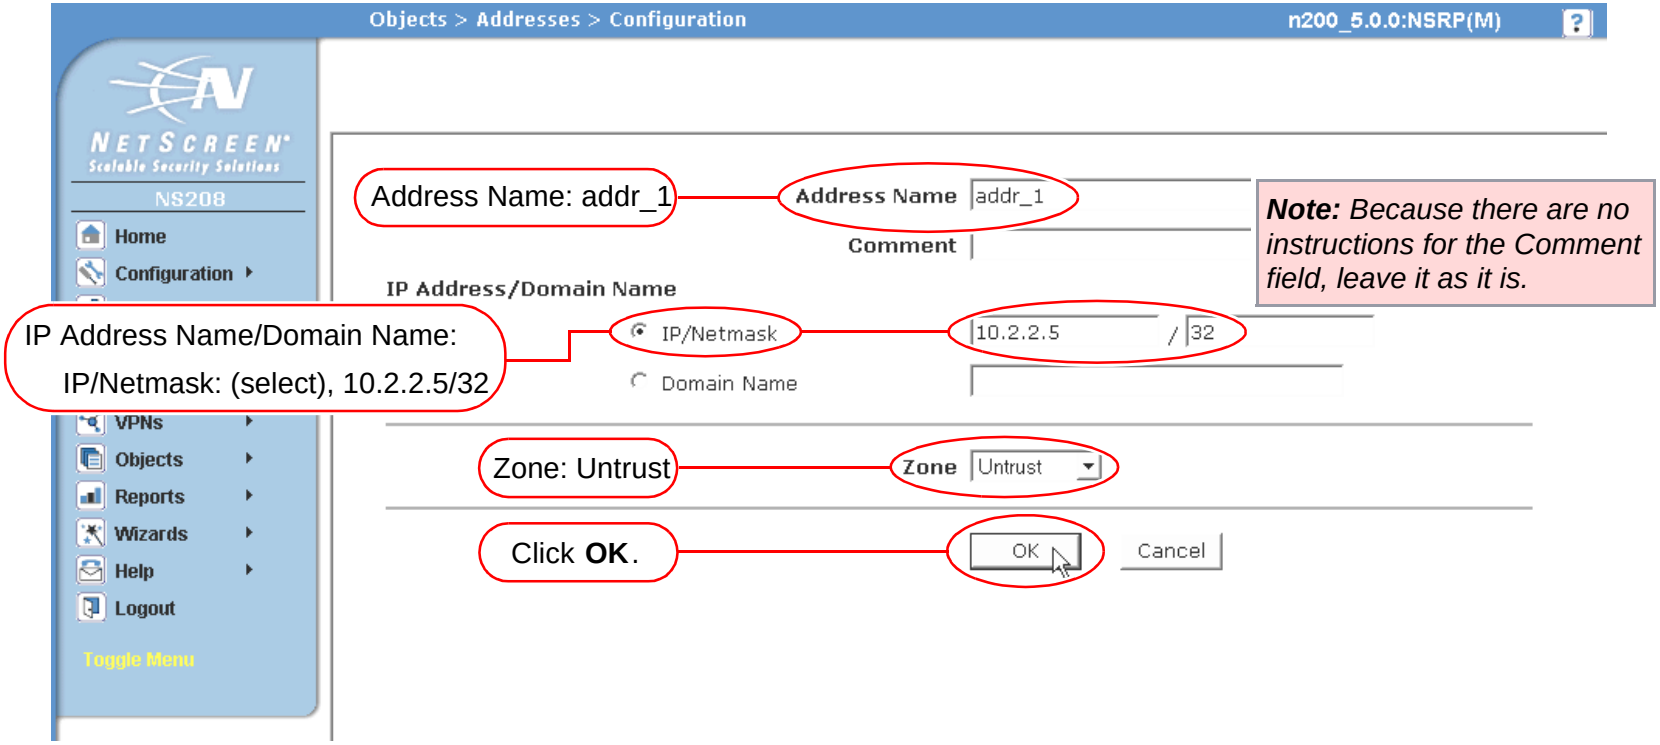

Illustration Conventions

The following graphics make up the basic set of images used in illustrations throughout this book:

Generic NetScreen Device

Virtual Routing Domain

Security Zone

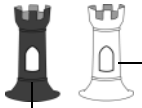

Security Zone Interfaces

White = Protected Zone Interface (example: Trust Zone)

Black = Outside Zone Interface (example: Untrust Zone)

Tunnel Interface

VPN Tunnel

Router Icon

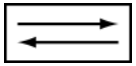

Switch Icon

Local Area Network (LAN) with a Single Subnet

(example: 10.1.1.0/24)

Internet

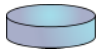

Dynamic IP (DIP) Pool

Desktop Computer

Laptop Computer

Generic Network Device

(examples: NAT server, Access Concentrator)

Server

Naming Conventions and Character Types

ScreenOS employs the following conventions regarding the names of objects—such as addresses, admin users, auth servers, IKE gateways, virtual systems, VPN tunnels, and zones—defined in ScreenOS configurations.

- If a name string includes one or more spaces, the entire string must be enclosed within double quotes ("); for example, **set address trust "local LAN" 10.1.1.0/24**.
- NetScreen trims any spaces leading or trailing text within a set of double quotes; for example, **" local LAN "** becomes **"local LAN"**.
- NetScreen treats multiple consecutive spaces as a single space.
- Name strings are case sensitive, although many CLI key words are case insensitive. For example, **"local LAN"** is different from **"local lan"**.

ScreenOS supports the following character types:

- Single-byte character sets (SBCS) and multiple-byte character sets (MBCS). Examples of SBCS are ASCII, European, and Hebrew. Examples of MBCS—also referred to as double-byte character sets (DBCS)—are Chinese, Korean, and Japanese.

  Note: *A console connection only supports SBCS. The WebUI supports both SBCS and MBCS, depending on the character sets that your Web browser supports.*

- ASCII characters from 32 (0x20 in hexidecimals) to 255 (0xff), except double quotes ("), which have special significance as an indicator of the beginning or end of a name string that includes spaces.

NETSCREEN DOCUMENTATION

To obtain technical documentation for any NetScreen product, visit www.netscreen.com/resources/manuals/.

To obtain the latest software version, visit www.netscreen.com/services/download_soft. Select a category of software product from the dropdown list, then follow the displayed instructions. (You must be a registered user to download NetScreen software.)

If you find any errors or omissions in the following content, please contact us at the e-mail address below:

techpubs@netscreen.com

1 Administration

This chapter describes various management methods and tools, ways to secure administrative traffic, and the administrative privilege levels that you can assign to admin users. This chapter contains the following sections:

- "Management via the Web User Interface" on page 3
  - "WebUI Help" on page 4
  - "HTTP" on page 5
  - "Secure Sockets Layer" on page 7
- "Management via the Command Line Interface" on page 9
  - "Telnet" on page 9
  - "Secure Shell" on page 11
  - "Secure Copy (SCP)" on page 20
  - "Serial Console" on page 21
- "Management via NetScreen-Security Manager" on page 23
  - "Initiating Connectivity Between Agent and Management System" on page 24
  - "Enabling and Disabling the Agent" on page 25
  - "Changing Management System Server Address" on page 26
  - "Setting Report Parameters" on page 26
- "Controlling Administrative Traffic" on page 29
  - "MGT and VLAN1 Interfaces" on page 30
  - "Administrative Interface" on page 32
  - "Manage IP" on page 34

- "Levels of Administration" on page 37
  - "Defining Admin Users" on page 39
- "Securing Administrative Traffic" on page 42
  - "Changing the Port Number" on page 43
  - "Changing the Admin Login Name and Password" on page 44
  - "Resetting the Device to the Factory Default Settings" on page 48
  - "Restricting Administrative Access" on page 49
  - "VPN Tunnels for Administrative Traffic" on page 51

MANAGEMENT VIA THE WEB USER INTERFACE

For administrative ease and convenience, you can use the Web user interface (WebUI). NetScreen devices use Web technology that provides a Web-server interface to configure and manage the software.

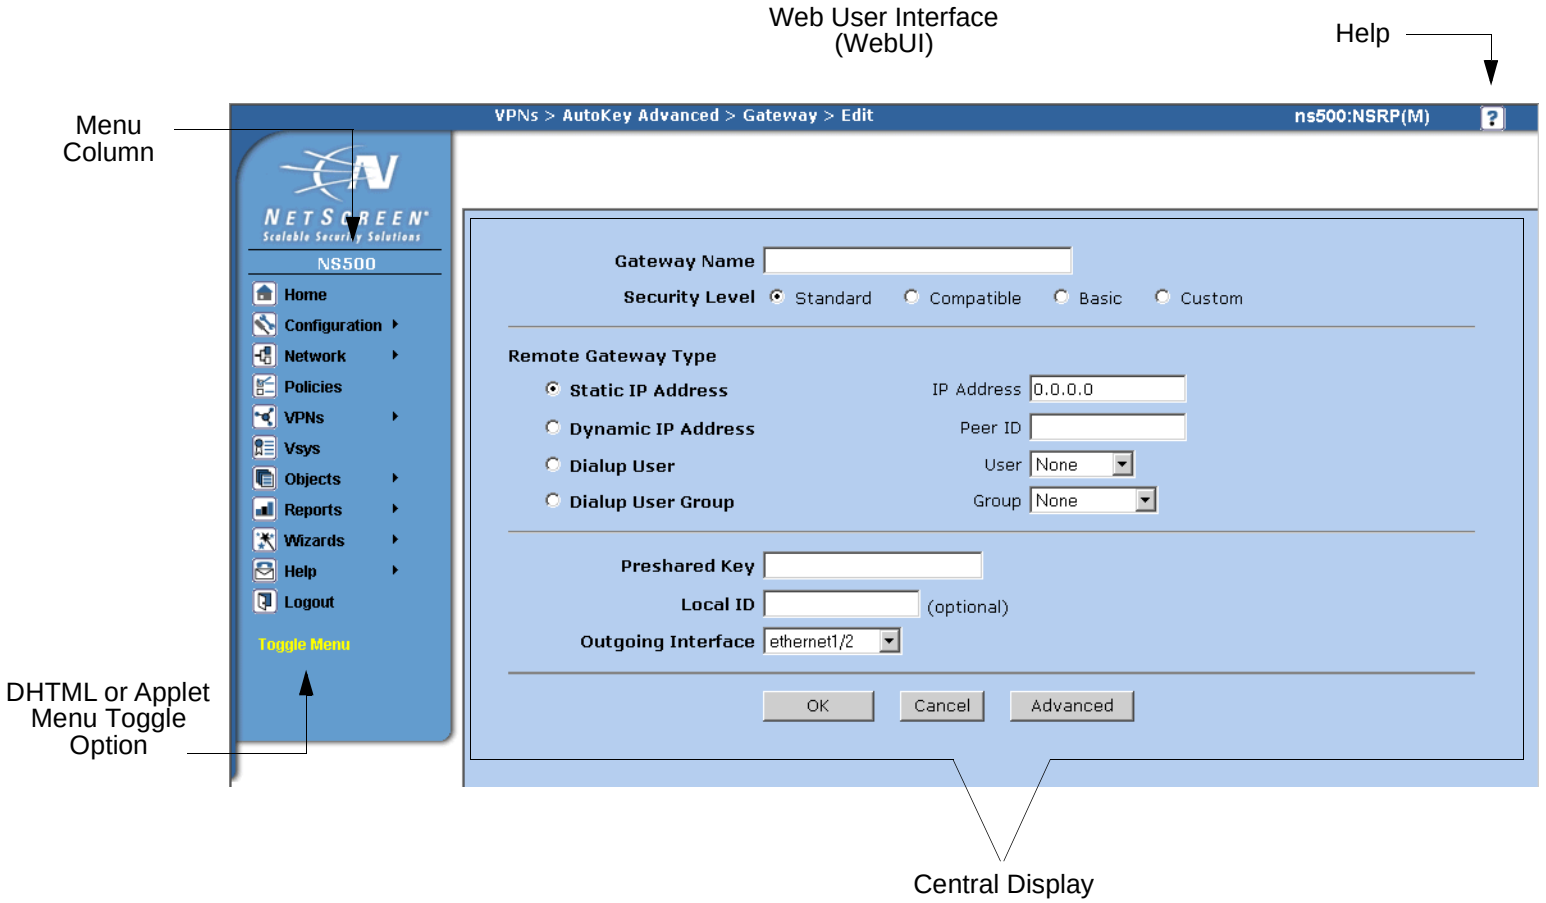

To use the WebUI, you must have the following:

- Netscape Communicator (version 4.7 or later) or Microsoft Internet Explorer (version 5.5 or later)
- TCP/IP network connection to the NetScreen device

WebUI Help

You can view Help files for the WebUI at http://help.netscreen.com/help/english/<screenos_version> /ns<platform_number> (for example, http://help.netscreen.com/help/english/5.0.0/ns500).

You also have the option of relocating the Help files. You might want to store them locally and point the WebUI to either the administrator's workstation or to a secured server on the local network. In case you do not have Internet access, storing the Help files locally provides accessibility to them you otherwise would not have.

Copying the Help Files to a Local Drive

The Help files are available on the documentation CD. You can modify the WebUI to point to the Help files on the CD in your local CD drive. You can also copy the files from the CD to a server on your local network or to another drive on your workstation and configure the WebUI to invoke the Help files from there.

Note: *If you want to run the Help files directly from the documentation CD, you can skip this procedure. Proceed to* *"Pointing the WebUI to the New Help Location" on page 4**.*

1. Load the documentation CD in the CD drive of your workstation.
2. Navigate to that drive and copy the directory named help.

   The Help directory contains the following subdirectories: english/<ScreenOS_number>/ns<platform_number>.
3. Navigate to the location you want to store the Help directory and paste it there.

Pointing the WebUI to the New Help Location

You must now redirect the WebUI to point to the new location of the Help directory. Change the default URL to the new file path, where

- **<path>** is the specific path to the Help directory from the administrator's workstation
- **<screenos_version>** is the version of the ScreenOS loaded on the NetScreen device that you are managing
- **<platform_number>** is the platform number of the NetScreen device

1. Configuration > Admin > Management: In the Help Link Path field, replace the underlined section of the default URL **http://help.netscreen.com/help/english/<screenos_version>/ns<platform_number>**

   with

   (for local drive) **file://<path>/ …**

   or

   (for local server) **http://<server_name>/<path>/ …**

2. Click **Apply**.

   When you click the **help** link in the upper right corner of the WebUI, the device now uses the new path that you specified in the Help Link Path field to locate the appropriate Help file.

HTTP

With a standard Web browser you can access, monitor, and control your network security configurations remotely using the Hypertext Transfer Protocol (HTTP).

You can secure HTTP administrative traffic by encapsulating it in a virtual private network (VPN) tunnel or by using the Secure Sockets Layer (SSL) protocol. You can further secure administrative traffic by completely separating it from network user traffic. To do this, you can run all administrative traffic through the MGT interface—available on some NetScreen devices—or bind an interface to the MGT zone and devote it exclusively to administrative traffic.

Note: *For more information, see "Secure Sockets Layer" on page 7, "VPN Tunnels for Administrative Traffic" on page 51, and "MGT and VLAN1 Interfaces" on page 30.*

Session ID

The NetScreen device assigns each HTTP administrative session a unique session ID. For NetScreen devices that support virtual systems (vsys), the ID is globally unique across all systems—root and vsys.

Each session ID is a 39-byte number resulting from the combination of five pseudo-randomly generated numbers. The randomness of the ID generation—versus a simple numerical incrementation scheme—makes the ID nearly impossible to predict. Furthermore, the randomness combined with the length of the ID makes accidental duplication of the same ID for two concurrent administrative sessions extremely unlikely.

The following are two benefits that a session ID provides to NetScreen administrators:

- The NetScreen device can distinguish concurrent sessions from multiple admins behind a NAT device that assigns the same source IP address to all outbound packets.

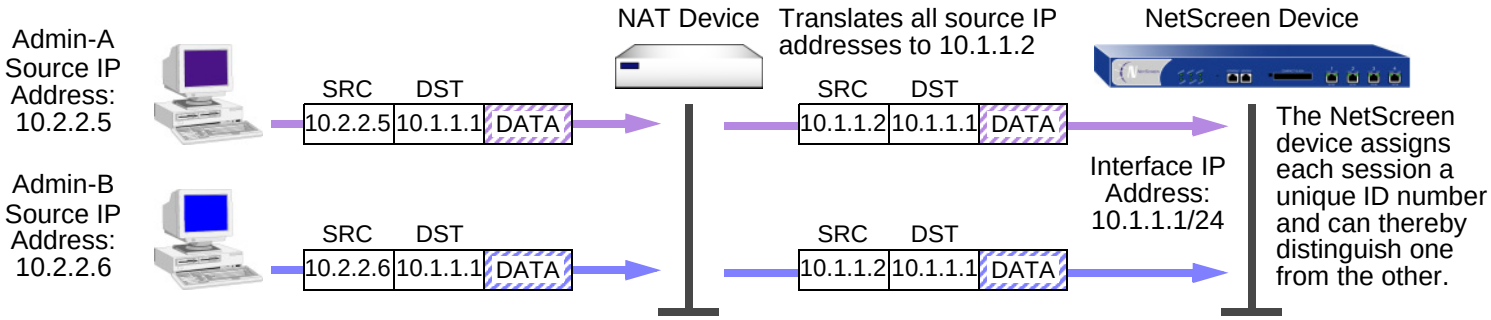

- The NetScreen device can distinguish concurrent root-level admin sessions from the same source IP address to the root system, and from there to different virtual systems.

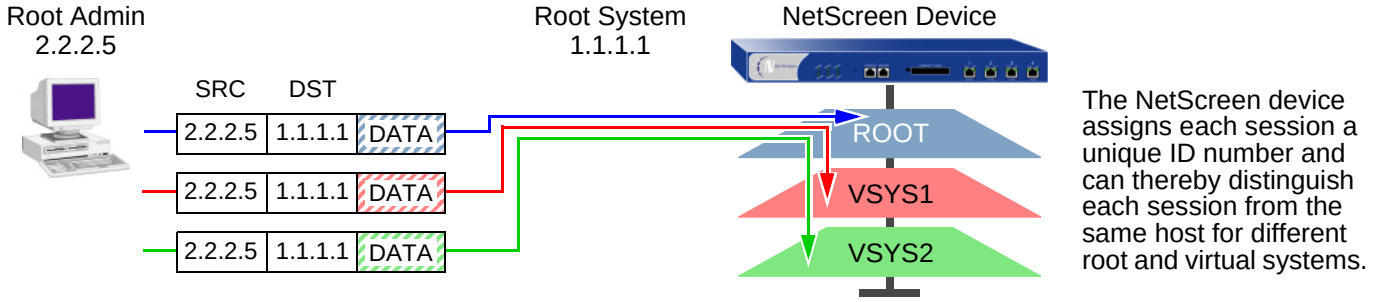

Secure Sockets Layer

Secure Sockets Layer (SSL) is a set of protocols that can provide a secure connection between a Web client and a Web server communicating over a TCP/IP network. NetScreen ScreenOS provides:

- Web SSL support
- SSL version 3 compatibility (not version 2)
- Netscape Communicator 4.7x and Internet Explorer 5.x compatibility[1]
- Public Key Infrastructure (PKI) key management integration (see "Public Key Cryptography" on page **5**-15.)

SSL is not a single protocol, but consists of the SSL Handshake Protocol (SSLHP), which allows the server and client to authenticate each other and negotiate an encryption method, and the SSL Record Protocol (SSLRP), which provides basic security services to higher-level protocols such as HTTP. These two protocols operate at the following two layers in the Open Systems Interconnection (OSI) model:

- SSLHP at the application layer (layer 7)
- SSLRP at the presentation layer (layer 6)

Independent of application protocol, SSL uses TCP to provide secure service. SSL uses certificates to authenticate first the server or both the client and the server, and then encrypt the traffic sent during the session. Before using SSL, you must first create a public/private key pair and then load a certificate. Because SSL is integrated with PKI key/certificate management, you can select the SSL certificate from one of the certificates in the certificate list. You can also use the same certificate for an IPSec VPN.

Note: *For information on obtaining certificates, see "Certificates and CRLs" on page **5**-21.*

1. Check your Web browser to see how strong the ciphers can be and which ones your browser supports. (Both the NetScreen device and your Web browser must support the same kind and size of ciphers you use for SSL.) In Internet Explorer 5x, click **Help**, **About Internet Explorer**, and read "Cipher Strength." To obtain the advanced security package, click the **Update Information** link. In Netscape Communicator, click **Help**, **About Communicator**, and read the section about RSA®. To change the SSL configuration settings, click **Security Info**, **Navigator**, **Configure SSL v3**.

NetScreen supports the following encryption algorithms for SSL:

- RC4 with 40-bit and 128-bit keys
- DES: Data Encryption Standard
- 3DES: Triple DES

NetScreen supports the same authentication algorithms for SSL as for VPNs—Message Digest version 5 (MD5) and Secure Hash Algorithm version 1 (SHA-1). The RC4 algorithms are always paired with MD5; DES and 3DES with SHA-1.

The basic steps for setting up SSL are as follows:

1. Obtain a certificate and load it on the NetScreen device[2].

   For details on requesting and loading a certificate, see *"Certificates and CRLs" on page **5**-21*.

2. Enable SSL management:

   Configuration > Admin > Management: Enter the following, and then click **Apply**:

   Certificate: Select the certificate you intend to use from the drop-down list.

   Cipher: Select the cipher you intend to use from the drop-down list.

3. Configure the interface through which you manage the NetScreen device to permit SSL management:

   Network > Interfaces > Edit (for the interface you want to manage): Enable the **SSL** management service check box, and then click **OK**.

4. Connect to the NetScreen device via the SSL port. That is, when you type the IP address for managing the NetScreen device in your browser's URL field, change "http" to "https", and follow the IP address with a colon and the HTTPS (SSL) port number (for example, https://123.45.67.89:1443).

2. Be sure to specify a bit length that your Web browser also supports.

MANAGEMENT VIA THE COMMAND LINE INTERFACE

Advanced administrators can attain finer control by using the command line interface (CLI). To configure a NetScreen device with the CLI, you can use any software that emulates a VT100 terminal. With a terminal emulator, you can configure the NetScreen device using a console from any Windows, UNIX™, or Macintosh® operating system. For remote administration through the CLI, you can use Telnet or Secure Shell (SSH). With a direct connection through the console port, you can use Hyperterminal®.

Note: *For a complete listing of the ScreenOS CLI commands, refer to the* NetScreen CLI Reference Guide.

Telnet

Telnet is a login and terminal emulation protocol that uses a client/server relationship to connect to and remotely configure network devices over a TCP/IP network. The administrator launches a Telnet client program on the administration workstation and creates a connection with the Telnet server program on the NetScreen device. After logging on, the administrator can issue CLI commands, which are sent to the Telnet program on the NetScreen device, effectively configuring the device as if operating through a direct connection. Using Telnet to manage NetScreen devices requires the following:

- Telnet software on the administrative workstation
- An Ethernet connection to the NetScreen device

The setup procedure to establish a Telnet connection is as follows:

Establishing a Telnet connection

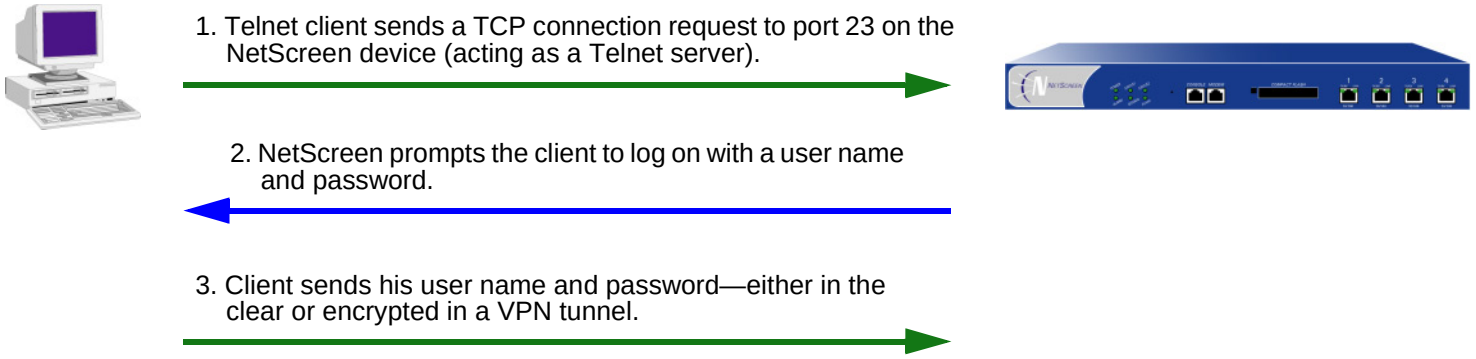

To minimize an unauthorized user's chances of logging in to a device, you can limit the number of unsuccessful login attempts allowed before the NetScreen device terminates a Telnet session. This restriction also protects against certain types of attacks, such as automated dictionary attacks.

Be default, the NetScreen device allows up to three unsuccessful login attempts before it closes the Telnet session. To change this number, enter the following command:

```
set admin access attempts number
```

Note: *You must use the CLI to set this restriction.*

Securing Telnet Connections

You can secure Telnet traffic by completely separating it from network user traffic. Depending upon your NetScreen device model, you can run all administrative traffic through the MGT interface or devote an interface such as the DMZ entirely to administrative traffic.

In addition, to ensure that admin users use a secure connection when they manage a NetScreen device through Telnet, you can require such users to telnet only through a virtual private network (VPN) tunnel[3]. After you have set this restriction, the device denies access if anyone tries to telnet without going through a VPN tunnel.

To restrict Telnet access through a VPN, enter the following command:

```
set admin telnet access tunnel
```

Note: *You must use the CLI to set this restriction .*

3. For information on VPN tunnels, see Volume 5, "VPNs".

Secure Shell

The built-in Secure Shell (SSH) server on a NetScreen device provides a means by which administrators can remotely manage the device in a secure manner using applications that are Secure Shell (SSH) aware. SSH allows you to open a remote command shell securely and execute commands. SSH provides protection from IP or DNS spoofing attacks and password or data interception.

You can choose to run either an SSH version 1 (SSHv1) or SSH version 2 (SSHv2) server on the NetScreen device. SSHv2 is considered more secure than SSHv1 and is currently being developed as the IETF standard. However, SSHv1 has been widely deployed and is commonly used. Note that SSHv1 and SSHv2 are not compatible with each other. That is, you cannot use an SSHv1 client to connect to an SSHv2 server on the NetScreen device, or vice versa. The client console or terminal application must run the same SSH version as the server.

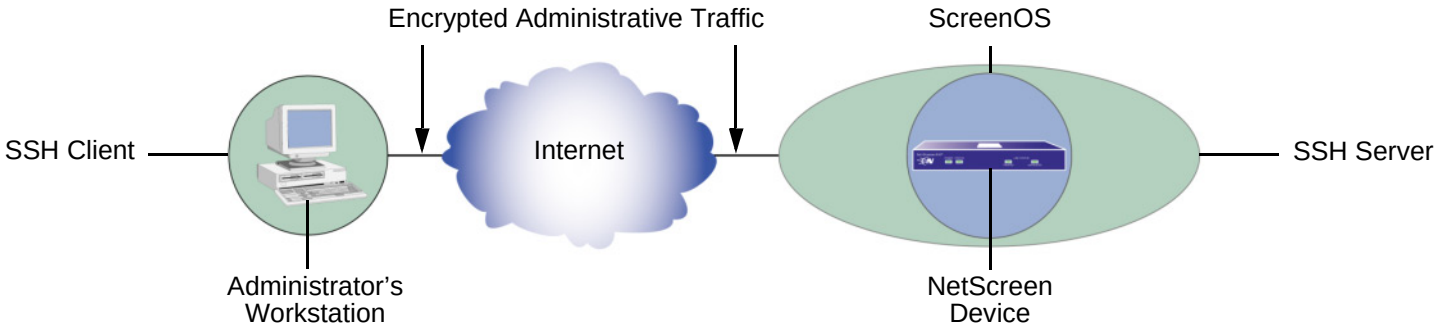

The basic SSH connection procedure is shown below:

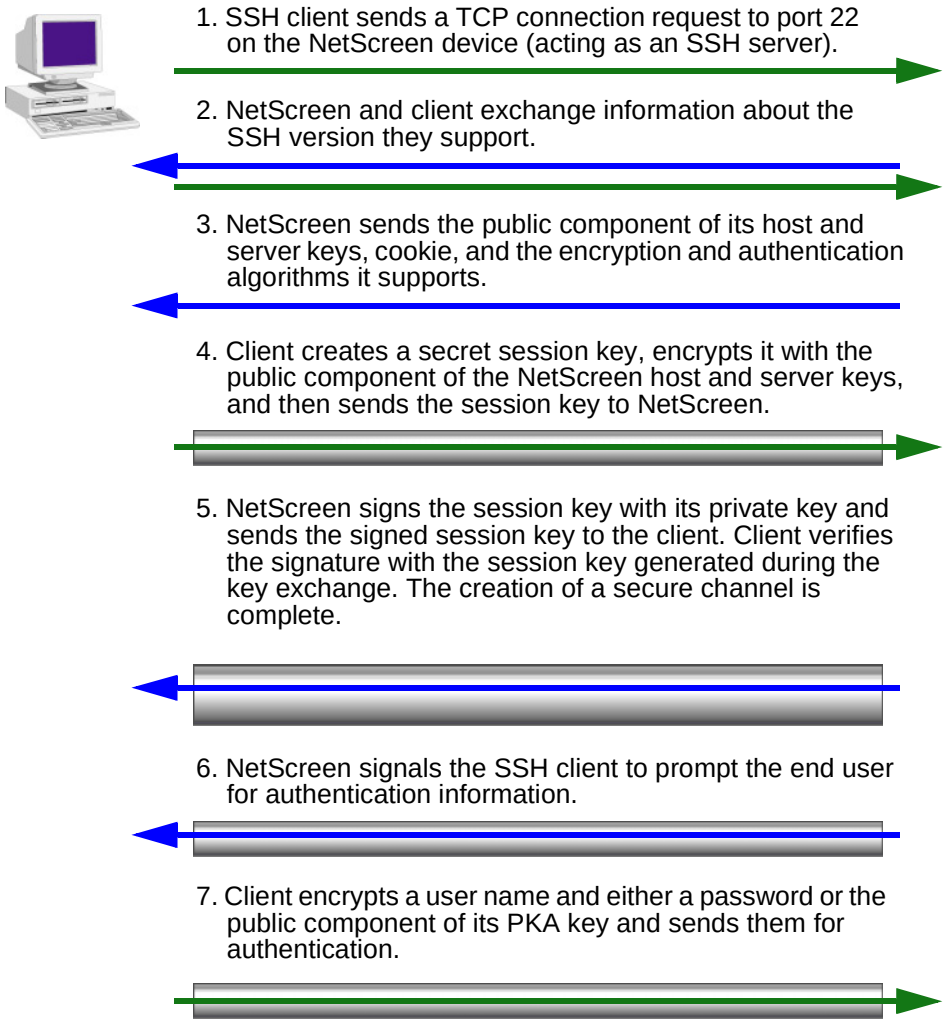

Keys

Host Key: Public key component of a public/private key pair used to authenticate the NetScreen device/vsys to the client and encrypt the session key. (Each vsys has its own host key.) The host key is permanently bound to the device/vsys.

Server Key: Temporary RSA public/private key pair used to encrypt the session key. (NetScreen generates a new one every hour by default for each vsys.)

Session Key: Temporary secret key (DES or 3DES) that the client and NetScreen create together during the connection setup to encrypt communication (when the session ends, it is discarded).

PKA Key: Persistent RSA public/private key pair that resides on the SSH client. The client's public key must also be loaded on the NetScreen device before initiating an SSH connection and the PKA key must be bound to the admin user.

Note: *Public/Private Key Pair = A set of cryptographic keys such that what one encrypts the other (and only the other) can decrypt.*

A maximum of five SSH sessions are allowed on a NetScreen device at any one time.

Client Requirements

As mentioned previously, the client application must run the same SSH version as the server on the NetScreen device. SSHv2 clients must be configured to request the Diffie-Hellman key exchange algorithm and the Digital Signature Algorithm (DSA) for public key device authentication. SSHv1 clients must be configured to request the RSA for public key device authentication.

Basic SSH Configuration on the NetScreen Device

The following are the basic steps for configuring SSH on a NetScreen device:

1. Determine whether you will use password or Public Key Authentication (PKA) for SSH. If PKA will be used, the PKA keys must be bound to an admin before SSH connections can be made. See "Authentication" on page 15 for more information about using passwords or PKA.
2. Determine which version of SSH you need to enable on the NetScreen device. (Remember that the client application and the SSH server on the NetScreen device must run the same SSH version.) If you enabled SSH on the NetScreen device in a previous ScreenOS version, SSHv1 runs when you enable SSH now. To see which version of SSH is active but not enabled on the NetScreen device, enter the CLI **get ssh** command:

```
ns-> get ssh
SSH V1 is active
SSH is not enabled
SSH is not ready for connections
Maximum sessions: 8
Active sessions: 0
```

In the output shown above, SSHv1 is active and runs when you enable SSH. If you want to use a different SSH version, make sure that all keys created with the previous version are removed. For example, to clear SSHv1 keys and to use SSHv2, enter the following CLI commands:

```
ns-> delete ssh device all
```

The following messages appear:

```
SSH disabled for vsys: 1
PKA key deleted from device: 0
Host keys deleted from device: 1
Execute the ‘set ssh version v2’ command to activate SSH v2 for the device
```

To use SSHv2, enter the following CLI command:

```
ns-> set ssh version v2
```

Note: *Setting the SSH version does not enable SSH on the NetScreen device.*

3. If you do not want to use port 22 (the default port) for SSH client connections, you can specify a port number between 1024 and 32767[4].

   ```
   ns-> set admin ssh port 1024
   ```

4. Enable SSH for the root system or for the virtual system. See “SSH and Vsys” on page 17 for additional information about enabling and using SSH on a per-vsys basis.

   To enable SSH for the root system:

   ```
   ns-> set ssh enable
   ```

   To enable SSH for a vsys, you need to first enter the vsys and then enable SSH:

   ```
   ns-> set vsys v1
   ns(v1)-> set ssh enable
   ```

5. Enable SSH on the interface on which the SSH client will connect.

   ```
   ns-> set interface manage ssh
   ```

6. Distribute the host key that is generated on the NetScreen device to the SSH client. See “Host Key” on page 18 for more information.

4. **You can also use the WebUI to change the port number and enable SSHv2 and SCP on the Configuration > Admin > Management page**.

Authentication

An administrator can connect to a NetScreen device with SSH using one of two authentication methods:

- **Password Authentication:** This method is used by administrators who need to configure or monitor a NetScreen device. The SSH client initiates an SSH connection to the NetScreen device. If SSH manageability is enabled on the interface receiving the connection request, the NetScreen device signals the SSH client to prompt the user for a user name and password. When the SSH client has this information, it sends it to the NetScreen device, which compares it with the user name and password in the admin user's account. If they match, the NetScreen device authenticates the user. If they do not match, the NetScreen device rejects the connection request.
- **Public Key Authentication (PKA):** This method provides increased security over the password authentication method and allows you to run automated scripts. Basically, instead of a user name and password, the SSH client sends a user name and the public key component of a public/private key pair[5]. The NetScreen device compares it with up to four public keys that can be bound to an admin. If one of the keys matches, the NetScreen device authenticates the user. If none of them match, the NetScreen device rejects the connection request.

Both authentication methods require the establishment of a secure connection before the SSH client logs on. After an SSH client has established an SSH connection with the NetScreen device, he must authenticate himself either with a user name and password or with a user name and public key.

Both password authentication and PKA require that you create an account for the admin user on the NetScreen device and enable SSH manageability on the interface through which you intend to manage the NetScreen device via an SSH connection. (For information about creating an admin user account, see "Defining Admin Users" on page 39.) The password authentication method does not require any further set up on the SSH client.

On the other hand, to prepare for PKA, you must first perform the following tasks:

5. The supported authentication algorithms are RSA for SSHv1 and DSA for SSHv2.

1. On the SSH client, generate a public and private key pair using a key generation program. (The key pair is either RSA for SSHv1 or DSA for SSHv2. See the SSH client application documentation for more information.)

   Note: *If you want to use PKA for automated logins, you must also load an agent on the SSH client to decrypt the private key component of the PKA public/private key pair and hold the decrypted version of the private key in memory.*

2. Move the public key from the local SSH directory to a directory on your TFTP server[6], and launch the TFTP program.
3. Log on to the NetScreen device so that you can configure it through the CLI.
4. To load the public key from the TFTP server to the NetScreen device, enter one of the following CLI commands:

   For SSHv1:

   ```
   exec ssh tftp pka-rsa [ username name ] file-name name_str ip-addr
        tftp_ip_addr
   ```

   For SSHv2:

   ```
   exec ssh tftp pka-dsa [ user-name name ] file-name name_str ip-addr
        tftp_ip_addr
   ```

   The **username** or **user-name** options are only available to the root admin, so that only the root admin can bind an RSA key to another admin. When you—as the root admin or as a read/write admin—enter the command without a user name, the NetScreen device binds the key to your own admin account; that is, it binds the key to the admin that enters the command.

   Note: *The NetScreen device supports up to four PKA public keys per admin user.*

6. You can also paste the content of the public key file directly into the CLI command **set ssh pka-rsa [username** *name_str* **] key** *key_str* (for SSHv1) or **set ssh pka-dsa [user-name** *name_str* **] key** *key_str* (for SSHv2), pasting it where indicated by the variable *key_str*, or into the Key field in the WebUI (Configuration > Admin > Administrators > SSH PKA). However, the CLI and WebUI have a size restriction: the public key size cannot exceed 512 bits. This restriction is not present when loading the key via TFTP.

When an administrator attempts to log on via SSH on an interface that has SSH manageability enabled, the NetScreen device first checks if a public key is bound to that administrator. If so, the NetScreen device authenticates the administrator using PKA. If a public key is not bound to the administrator, the NetScreen device prompts for a user name and password. (You can use the following command to force an admin to use only the PKA method: **set admin ssh password disable username** *name_str* .) Regardless of the authentication method you intend the administrator to use, when you initially define his or her account, you still must include a password, even though when you later bind a public key to this user, the password becomes irrelevant.

SSH and Vsys

For NetScreen devices that support vsys, you can enable and configure SSH on a per-vsys basis. Each vsys has its own host key (see "Host Key" on page 18) and maintains and manages a PKA key for the admin of the system.

Note that the maximum number of SSH sessions is a device-wide limit and is between 2 and 24, depending upon the device. If the maximum number of SSH clients are already logged into the device, no other SSH client can log in to the SSH server. The root system and the vsys share the same SSH port number. This means that if you change the SSH port from the default port 22, the port is changed for all vsys as well.

Host Key

The host key allows the NetScreen device to identify itself to an SSH client. On NetScreen devices that support virtual systems (vsys), each vsys has its own host key. When SSH is first enabled on a vsys (for devices that support vsys) or on a NetScreen device, a host key is generated that is unique to the vsys or device. The host key is permanently bound to the vsys or device and the same host key is used if SSH is disabled and then enabled again.

The host key on the NetScreen device must be distributed to the SSH client in one of two ways:

- Manually—the root or vsys admin sends the host key to the client admin user via e-mail, phone, etc. The receiving admin stores the host key in the appropriate SSH file on the SSH client system. (The SSH client application determines the file location and format.)
- Automatically—When the SSH client connects to the NetScreen device, the SSH server sends the unencrypted public component of the host key to the client. The SSH client searches its local host key database to see if the received host key is mapped to the address of the NetScreen device. If the host key is unknown (there is no mapping to the NetScreen device address in the client's host key database), the Admin user may be able to decide whether to accept the host key. Otherwise, the connection is terminated. (See the appropriate SSH client documentation for information on accepting unknown host keys.)

To verify that the SSH client has received the correct host key, the Admin user on the client system can generate the SHA hash of the received host key. The client Admin user can then compare this SHA hash with the SHA hash of the host key on the NetScreen device. On the NetScreen device, you can display the SHA hash of the host key by executing the CLI command **get ssh host-key**.

Example: SSHv1 with PKA for Automated Logins

In this example, you (as the root admin) set up SSHv1 public key authentication (PKA) for a remote host that runs an automated script. The sole purpose for this remote host to access the NetScreen device is to download the configuration file every night. Because authentication is automated, no human intervention is necessary when the SSH client logs on to the NetScreen device.

You define an admin user account named cfg, with password cfg and read-write privileges. You enable SSH manageability on interface ethernet1, which is bound to the Untrust zone.

You have previously used a key generation program on your SSH client to generate an RSA public/private key pair, moved the public key file, which has the file name “idnt_cfg.pub”, to a directory on your TFTP server, and launched the TFTP program. The IP address of the TFTP server is 10.1.1.5.

WebUI

Configuration > Admin > Administrators > New: Enter the following, and then click **OK**:

Name: cfg

New Password: cfg

Confirm Password: cfg

Privileges: Read-Write (select)

SSH Password Authentication: (select)

Network > Interfaces > Edit (for ethernet1): Select **SSH** in Service Options, and then click **OK**.

Note: *You can only load a public key file for SSH from a TFTP server via the* ***exec ssh*** *command.*

CLI

```
set admin user cfg password cfg privilege all
set interface ethernet1 manage ssh
exec ssh tftp pka-rsa username cfg file-name idnt_cfg.pub ip-addr 10.1.1.5
save
```

Secure Copy (SCP)

Secure Copy (SCP) provides a way for a remote client to transfer files to or from the NetScreen device using the SSH protocol. (The SSH protocol provides authentication, encryption, and data integrity to the SCP connection.) The NetScreen device acts as an SCP server to accept connections from SCP clients on remote hosts.

SCP requires that the remote client be authenticated before file transfer commences. SCP authentication is exactly the same process used to authenticate SSH clients. The SCP client can be authenticated with either a password or a PKA key. Once the SCP client is authenticated, one or more files can be transferred to or from the NetScreen device. The SCP client application determines the exact method for specifying the source and destination file names; refer to the SCP client application documentation.

SCP is disabled by default on the NetScreen device. To enable SCP, you must also enable SSH.

WebUI

Configuration > Admin > Management > Select the following, and then click **Apply**:

Enable SSH: (select)

Enable SCP: (select)

CLI

```
set ssh enable
set scp enable
save
```

The following is an example of an SCP client command to copy the configuration file from flash memory on a NetScreen device (administrator name is netscreen and IP address is 10.1.1.1) to the file ns_sys_config_backup on the client system:

```
scp netscreen@10.1.1.1:ns_sys_config ns_sys_config_backup
```

You need to consult your SCP client application documentation for information on how to specify the administrator name, device IP address, source file, and destination file.

Serial Console

You can manage a NetScreen device through a direct serial connection from the administrator's workstation to the NetScreen device via the console port. Although a direct connection is not always possible, this is the most secure method for managing the device provided that the location around the NetScreen device is secure.

Note: *To prevent unauthorized users from logging in remotely as the root admin, you can require the root admin to log in to the NetScreen device through the console only. For additional information on this restriction, see "Restricting the Root Admin to Console Access" on page 50.*

Depending on your NetScreen device model, creating a serial connection requires one of the following cables:

- A female DB-9 to male DB-25 straight-through serial cable
- A female DB-9 to male DB-9 straight-through serial cable
- A female DB-9 to male MiniDIN-8 serial cable
- A female DB-9 to RJ-45 adapter with an RJ-45 to RJ-45 straight-through ethernet cable

You will also need Hyperterminal software (or another kind of VT100 terminal emulator) on the management workstation, with the Hyperterminal port settings configured as follows:

- Serial communications 9600 bps
- 8 bit
- No parity
- 1 stop bit
- No flow control

Note: *For more details on using Hyperterminal, see the "Getting Started" chapter in the* NetScreen CLI Reference Guide *or the "Initial Configuration" chapter in one of the installer's guides.*

Modem Port

You can also manage the NetScreen device by connecting the administrator's workstation to the modem port on the device. The modem port functions similarly to the console port, except that you cannot define parameters for the modem port or use this connection to upload an image.

To prevent unauthorized users from managing the device through a direct connection to the console or modem port, you can disable both ports by entering the following commands:

```
set console disable
set console aux disable
```

Note: *On the NS-5XT, you can use the modem port to connect to an external modem only.*

MANAGEMENT VIA NETSCREEN-SECURITY MANAGER

NetScreen-Security Manager (NSM) is an enterprise-level management application that configures multiple devices over a LAN or WAN environment. The Security Manager UI enables administrators to configure many devices from central locations.

Security Manager uses two components to allow remote communication with NetScreen devices.

- The *Management System*, a set of services that reside on an external host. These services process, track, and store device management information exchanged between the device and the Security Manager UI.
- The *Agent*, a service that resides on each managed NetScreen device. The Agent receives configuration parameters from the external Management System and pushes it to ScreenOS. The Agent also monitors the device and transmits reports back to the Management System.

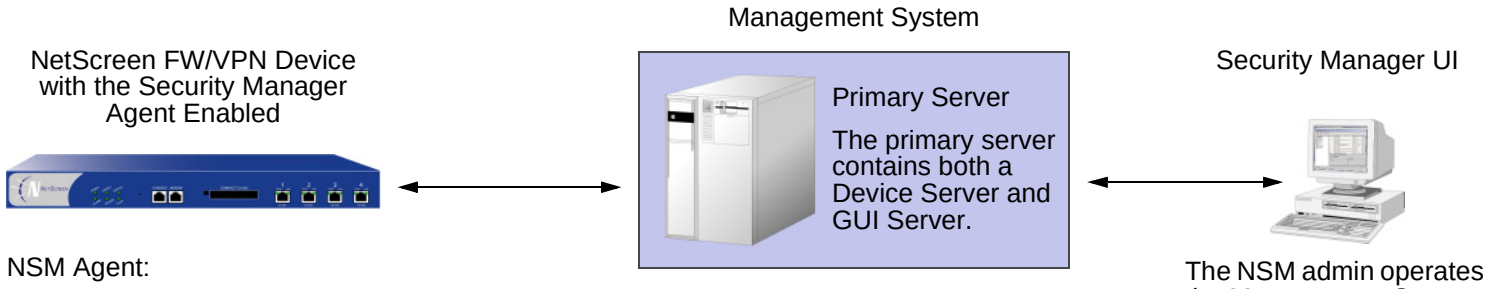

NSM Agent:

The NetScreen device uses its embedded NSM Agent to communicate with the Device Server.

Device and GUI Servers:

The Device Server pushes configuration changes to the NetScreen device and receives operational and statistical reports from it.

The GUI Server processes the configuration changes that it receives from one or more Security Manager clients.

For more information about these and other Security Manager components, refer to the *NetScreen-Security Manager 2004 Administrator's Guide*.

Initiating Connectivity Between Agent and Management System

Before the Security Manager can access and manage the NetScreen device, it is necessary to initiate communication between the Agent (which resides on the device) and the Management System (which resides on an external host).

Initialization may require up to two users at two different sites, depending upon the current availability of the NetScreen device. These users may include the *Security Manager administrator*, who uses the Security Manager UI on a client host, and the *on-site user*, who executes CLI commands on the NetScreen device via a console session. Possible initialization cases are as follows.

- Case 1: The device already has a known IP address, and is reachable over your network infrastructure.

  In this case, the Security Manager administrator adds the device using the Security Manager UI on the client host. (No on-site user is necessary.) The NetScreen device automatically connects back to the Management System, and is ready to send configuration information to the NSM database that resides there.

- Case 2: The IP address is unreachable.

  In this case, both users perform initialization tasks. The administrator adds the device through the Security Manager UI. The administrator also determines which CLI commands the on-site user needs, and delivers them to the user, who then executes them through the console. The device then automatically connects with the Management System, and is ready to send configuration information to the NSM database.

  Note: *If the device runs ScreenOS version 4.x, the on-site user must manually enable NetScreen-Global PRO, NACN, or both before the device can establish a connection to the Management System.*

- Case 3: The device is new and contains factory default settings.

  In this case, both users perform initialization tasks. The on-site user can use an encrypted configuration script called *Configlet*, which the Security Manager administrator generates. The process is as follows.

  a. The Security Manager administrator selects the device platform and ScreenOS version, using the Add Device wizard in the Security Manager UI.

  b. The administrator edits the device and enters any desired configuration.

c. The administrator Activates the device.

d. The administrator generates and delivers the Configlet file (or the necessary CLI commands, as with Case 2) to the on-site user.

e. The on-site user executes Configlet (or the CLI commands).

For more information, refer to “Adding Devices” in *NetScreen-Security Manager 2004 Administrator’s Guide*.

Enabling and Disabling the Agent

Before the NetScreen device can communicate with the Management System, you must enable the Agent that resides on the device.

Example: Enabling the Security Manager Agent

In the following example you enable the Agent.

WebUI

Configuration > Admin > NSM: Enter the following, and then click **Apply**:

Enable Communication with NetScreen Security Manager (NSM): (Select)

CLI

```
set nsmgmt enable
save
```

Changing Management System Server Address

The IP address by which the Agent identifies the external Management System servers is a configurable parameter.

Example: Setting the Primary Server IP Address

In the following example you set the primary server IP address to 1.1.1.1.

WebUI

Configuration > Admin > NSM: Enter the following, and then click **Apply**:

Primary IP Address/Name: 1.1.1.1

CLI

```
set nsmgmt server primary 1.1.1.1
save
```

Setting Report Parameters

The Agent monitors the NetScreen device events and transmits reports back to the Management System. This allows the Security Manager administrator to view the events from the Security Management UI.

The categories of events tracked by the Agent are as follows.

- *Alarms* report potentially dangerous attacks or traffic anomalies, including attacks detected through deep inspection.
- *Log events* report changes in device configuration and non-severe changes that occur on the device.
- *Protocol distribution* events report messages generated by the following services:
  - AH (Authentication Header)
  - ESP (Encapsulating Security Payload)
  - GRE (Generic Routing Encapsulation)

  - ICMP (Internet Control Message Protocol)
  - OSPF (Open Shortest Path First)
  - TCP (Transmission Control Protocol)
  - UDP (User Datagram Protocol)
- *Statistics* messages report the following statistical information.
  - Attack statistics
  - Ethernet statistics
  - Traffic flow statistics
  - Policy statistics

Example: Enabling Alarm and Statistics Reporting

In the following example you enable transmission of all Alarm and Statistics messages to the Management System.

WebUI

Configuration > Admin > NSM: Enter the following, and then click **Apply**:

Attack Statistics: (Select)

Policy Statistics: (Select)

Attack Alarms: (Select)

Traffic Alarms: (Select)

Flow Statistics: (Select)

Ethernet Statistics: (Select)

Deep Inspection Alarms: (Select)

Event Alarms: (Select)

CLI

```
set nsmgmt report statistics attack enable
set nsmgmt report statistics policy enable
set nsmgmt report alarm attack enable
set nsmgmt report alarm traffic enable
set nsmgmt report statistics flow enable
set nsmgmt report statistics ethernet enable
set nsmgmt report alarm idp enable
set nsmgmt report alarm other enable
save
```

CONTROLLING ADMINISTRATIVE TRAFFIC

ScreenOS provides the following options for configuring and managing the NetScreen device:

- **WebUI:** Selecting this option allows the interface to receive HTTP traffic for management via the Web user interface (WebUI).
- **Telnet:** A terminal emulation program for TCP/IP networks such as the Internet, Telnet is a common way to remotely control network devices. Selecting this option enables Telnet manageability.
- **SSH:** You can administer the NetScreen device from an Ethernet connection or a dial-in modem using Secure Command Shell (SSH). You must have an SSH client that is compatible with Version 1.5 of the SSH protocol. These clients are available for Windows 95 and later, Windows NT, Linux, and UNIX. The NetScreen device communicates with the SSH client through its built-in SSH server, which provides device configuration and management services. Selecting this option enables SSH manageability.
- **SNMP:** The NetScreen device supports both SNMPv1 and SNMPv2c, and all relevant Management Information Base II (MIB II) groups, as defined in RFC-1213. Selecting this option enables SNMP manageability.
- **SSL:** Selecting this option allows the interface to receive HTTPS traffic for secure management of the NetScreen device via the WebUI.
- **NS Security Manager:** Selecting this option allows the interface to receive NetScreen-Security Manager traffic.
- **Ping:** Selecting this option allows the NetScreen device to respond to an ICMP echo request, or ping, which determines whether a specific IP address is accessible over the network.
- **Ident-Reset:** Services like Mail and FTP send identification requests. If they receive no acknowledgement, they send the request again. While the request is processing, there is no user access. By enabling the Ident-reset option, the NetScreen device sends a TCP reset announcement in response to an IDENT request to port 113 and restores access that has been blocked by an unacknowledged identification request.

To use these options, you enable them on one or more interfaces, depending on your security and administrative needs.

MGT and VLAN1 Interfaces

Some NetScreen devices have a physical interface—Management (MGT)—dedicated exclusively for management traffic. Use this interface for management traffic when running the NetScreen device in NAT or Route mode.

In Transparent mode, you can configure all NetScreen devices to allow administration through the logical interface, VLAN1. To enable management traffic to reach the VLAN1 interface, you must enable the management options you want both on VLAN1 and on the Layer 2 zones—V1-Trust, V1-Untrust, V1-DMZ, user-defined Layer 2 zone—through which the management traffic passes to reach VLAN1.

To maintain the highest level of security, NetScreen recommends that you limit administrative traffic exclusively to the VLAN1 or MGT interface and user traffic to the security zone interfaces. Separating administrative traffic from network user traffic greatly increases administrative security and assures constant management bandwidth.

Example: Administration through the MGT Interface

In this example, you set the IP address of the MGT interface to 10.1.1.2/24 and enable the MGT interface to receive Web and SSH administrative traffic.

WebUI

Network > Interfaces > Edit (for mgt): Enter the following, and then click **OK**:

IP Address/Netmask: 10.1.1.2/24

Management Services: WebUI, SSH: (select)

CLI

```
set interface mgt ip 10.1.1.2/24
set interface mgt manage web
set interface mgt manage ssh
save
```

Example: Administration through the VLAN1 Interface

In this example, you set the IP address of the VLAN1 interface to 10.1.1.1/24 and enable the VLAN1 interface to receive Telnet and Web administrative traffic through the V1-Trust zone.

WebUI

Network > Interfaces > Edit (for VLAN1): Enter the following, and then click **OK**:

IP Address/Netmask: 10.1.1.1/24

Management Services: WebUI, Telnet: (select)

Network > Zones > Edit (for V1-Trust): Select the following, and then click **OK**:

Management Services: WebUI, Telnet: (select)

CLI

```
set interface vlan1 ip 10.1.1.1/24
set interface vlan1 manage web
set interface vlan1 manage telnet
set zone v1-trust manage web
set zone v1-trust manage telnet
save
```

Administrative Interface

On NetScreen devices that have multiple physical interfaces for network traffic, but no physical MGT interface, you might dedicate one physical interface exclusively for administration, separating management traffic completely from network user traffic. For example, you might have local management access the device through an interface bound to the Trust zone and remote management through an interface bound to the Untrust zone.

Example: Setting Administrative Interface Options

In this example, you bind ethernet1 to the Trust zone and ethernet3 to the Untrust zone. You assign ethernet1 the IP address 10.1.1.1/24 and give it the Manage IP address 10.1.1.2. (Note that the Manage IP address must be in the same subnet as the security zone interface IP address.) You also allow ethernet 1 to receive Web and Telnet traffic. You then assign ethernet3 the IP address 1.1.1.1/32. You do not allow administrative traffic through ethernet3.

WebUI

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Manage IP: 10.1.1.2

Management Services: WebUI, Telnet

Enter the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/32

Manageable: (clear)

CLI

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 manage-ip 10.1.1.2
set interface ethernet1 telnet
set interface ethernet1 web
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/32
save
```

Manage IP

Any physical, redundant, or aggregate interface or subinterface you bind to a security zone can have at least two IP addresses:

- An interface IP address, which connects to a network.
- A logical manage IP address for receiving administrative traffic.

When a NetScreen device is a backup unit in a redundant group for High Availability (HA), you can access and configure the unit through its manage IP address (or addresses)

Note: *The manage IP address differs from the VLAN1 address in the following two ways:*

- *When the NetScreen device is in Transparent mode, the VLAN1 IP address can be the endpoint of a VPN tunnel, but the manage IP address cannot.*
- *You can define multiple manage IP addresses—one for each network interface—but you can only define one VLAN1 IP address—for the entire system.*

If you select the Manageable option on the interface configuration page in the WebUI, you can manage the NetScreen device either through the interface IP address or the Manage IP address associated with that interface.

Example: Setting Manage IPs for Multiple Interfaces

In this example, ethernet2 is bound to the DMZ zone and ethernet3 is bound to the Untrust zone. You set the management options on each interface to provide access for the specific kinds of administrative traffic using each interface. You allow HTTP and Telnet access on ethernet2 for a group of local administrators in the DMZ zone, and SNMP access on ethernet3 for central management from a remote site. Ethernet2 and ethernet3 each have a manage IP address, to which the various kinds of administrative traffic are directed.

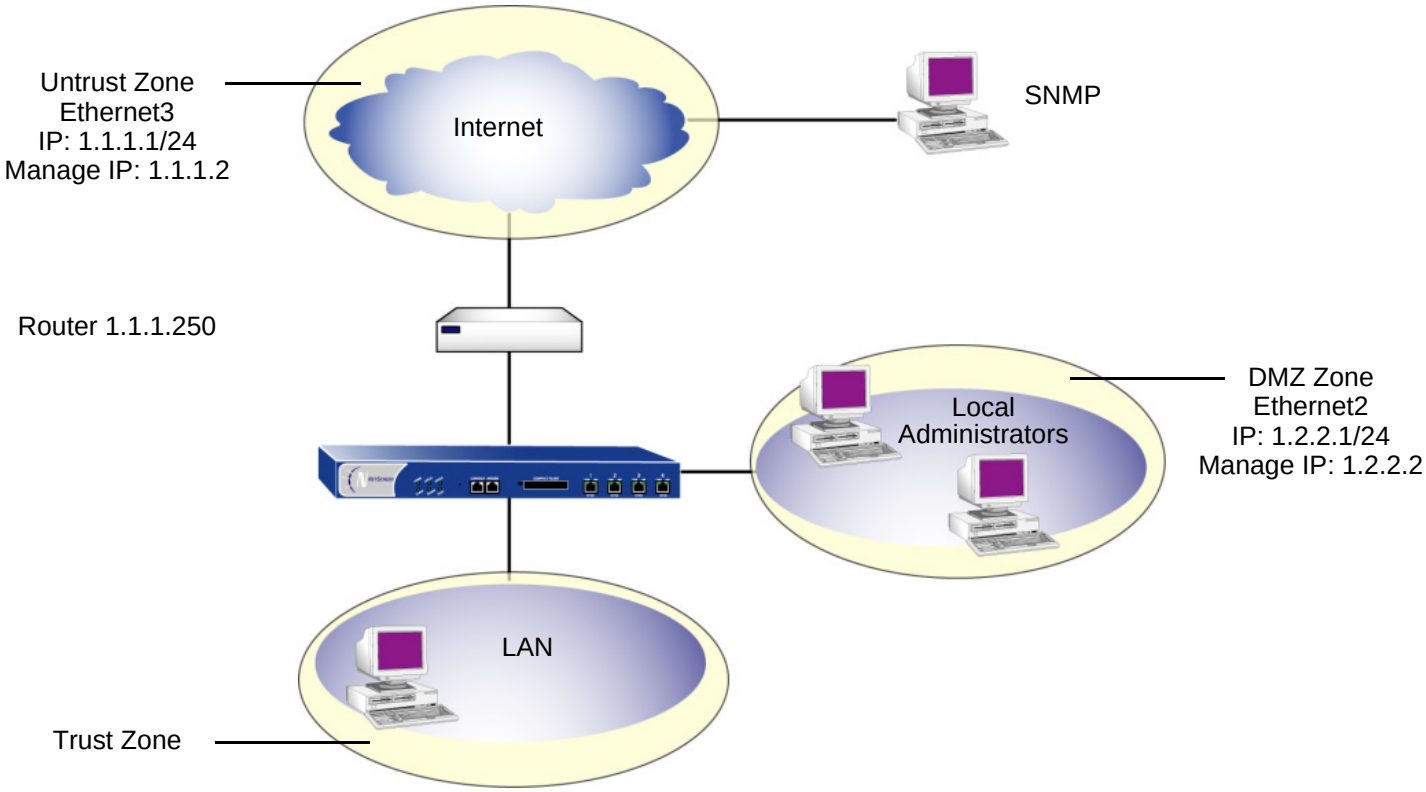

Note: *You also need to set a route directing self-generated SNMP traffic to use ethernet3 to reach the external router at IP address 1.1.1.250.*

WebUI

Network > Interfaces > Edit (ethernet2): Enter the following, and then click **OK**:

Zone Name: DMZ

Static IP: (select this option when present)

IP Address/Netmask: 1.2.2.1/24

Manage IP: 1.2.2.2

Management Services: WebUI, Telnet: (select)

Network > Interfaces > Edit (ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

Manage IP: 1.1.1.2

Management Services: SNMP

CLI

```
set interface ethernet2 zone dmz
set interface ethernet2 ip 1.2.2.1/24
set interface ethernet2 manage-ip 1.2.2.2
set interface ethernet2 manage web
set interface ethernet2 manage telnet

set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
set interface ethernet3 manage-ip 1.1.1.2
set interface ethernet3 manage snmp
save
```

LEVELS OF ADMINISTRATION

NetScreen devices support multiple administrative users. For any configuration changes made by an administrator, the NetScreen device logs the following information:

- The name of the administrator making the change
- The IP address from which the change was made
- The time of the change

There are several levels of administrative user. The availability of some of these levels depends on the model of your NetScreen device. The following sections list all the admin levels and the privileges for each level. These privileges are only accessible to an admin after he or she successfully logs in with a valid user name and password.

Root Administrator

The root administrator has complete administrative privileges. There is only one root administrator per NetScreen device. The root administrator has the following privileges:

- Manages the root system of the NetScreen device
- Adds, removes, and manages all other administrators
- Establishes and manages virtual systems, and assigns physical or logical interfaces to them
- Creates, removes, and manages virtual routers (VRs)
- Adds, removes, and manages security zones
- Assigns interfaces to security zones
- Performs asset recovery
- Sets the device to FIPS mode
- Resets the device to its default settings
- Updates the firmware
- Loads configuration files
- Clears all active sessions of a specified admin or of all active admins

Read/Write Administrator

The read/write administrator has the same privileges as the root administrator, but cannot create, modify, or remove other admin users. The read/write administrator has the following privileges:

- Creates virtual systems and assigns a virtual system administrator for each one
- Monitors any virtual system
- Tracks statistics (a privilege that cannot be delegated to a virtual system administrator)

Read-Only Administrator

The read-only administrator has only viewing privileges using the WebUI, and can only issue the **get** and **ping** CLI commands. The read-only administrator has the following privileges:

- Read-only privileges in the root system, using the following four commands: **enter**, **exit**, **get**, and **ping**
- Read-only privileges in virtual systems

Virtual System Administrator

Some NetScreen devices support virtual systems. Each virtual system (vsys) is a unique security domain, which can be managed by virtual system administrators with privileges that apply only to that vsys. Virtual system administrators independently manage virtual systems through the CLI or WebUI. On each vsys, the virtual system administrator has the following privileges:

- Creates and edits auth, IKE, L2TP, XAuth, and Manual Key users
- Creates and edits services
- Creates and edits policies
- Creates and edits addresses
- Creates and edits VPNs
- Modifies the virtual system administrator login password
- Creates and manages security zones
- Adds and removes virtual system read-only administrators

Virtual System Read-Only Administrator

A virtual system read-only administrator has the same set of privileges as a read-only administrator, but only within a specific virtual system. A virtual system read-only administrator has viewing privileges for his particular vsys through the WebUI, and can only issue the **enter**, **exit**, **get**, and **ping** CLI commands within his vsys.

> ***Note:*** *For more information on virtual systems, see "Virtual Systems" on page* ***7****-1.*

Defining Admin Users

The root administrator is the only one who can create, modify, and remove admin users. In the following example, the one performing the procedure must be a root administrator.

Example: Adding a Read-Only Admin

In this example, you—as the root admin—add a read-only administrator named Roger with password 2bd21wG7.

WebUI

Configuration > Admin > Administrators > New: Enter the following, and then click **OK**:

Name: Roger

New Password: 2bd21wG7[7]

Confirm New Password: 2bd21wG7

Privileges: Read-Only (Select)

CLI

```
set admin user Roger password 2bd21wG7 privilege read-only
save
```

7. The password can be up to 31 characters long and is case sensitive.

Example: Modifying an Admin

In this example, you—as the root admin—change Roger's privileges from read-only to read/write.

WebUI

Configuration > Admin > Administrators > Edit (for Roger): Enter the following, and then click **OK**:

Name: Roger

New Password: 2bd21wG7

Confirm New Password: 2bd21wG7

Privileges: Read-Write (Select)

CLI

```
unset admin user Roger
set admin user Roger password 2bd21wG7 privilege all
save
```

Example: Deleting an Admin

In this example, you—as the root admin—delete the admin user Roger.

WebUI

Configuration > Admin > Administrators: Click **Remove** in the Configure column for Roger.

CLI

```
unset admin user Roger
save
```

Example: Clearing an Admin's Sessions

In this example, you—as the root admin—terminate all active sessions of the admin user Roger. When you execute the following command, the NetScreen device closes all active sessions and automatically logs off Roger from the system.

WebUI

> ***Note:*** *You must use the CLI to clear an admin's sessions.*

CLI

```
clear admin name Roger
save
```

SECURING ADMINISTRATIVE TRAFFIC

To secure the NetScreen device during setup, perform the following steps:

1. On the Web interface, change the administrative port.

   See "Changing the Port Number" on page 43.

2. Change the user name and password for administration access.

   See "Changing the Admin Login Name and Password" on page 44.

3. Define the management client IP addresses for the admin users.

   See "Restricting Administrative Access" on page 49.

4. Turn off any unnecessary interface management service options.

   See "Controlling Administrative Traffic" on page 29.

5. Disable the ping and ident-reset service options on the interfaces, both of which respond to requests initiated by unknown parties and can reveal information about your network:

WebUI

Network > Interfaces > Edit (for the interface you want to edit): Disable the following service options, and then click **OK**:

Ping: Selecting this option allows the NetScreen device to respond to an ICMP echo request, or "ping," which determines whether a specific IP address is accessible from the device.

Ident-Reset: When a service such as Mail or FTP sends an identification request and receives no acknowledgment, it sends the request again. While the request is in progress, user access is disabled. With the Ident-Reset check box enabled, the NetScreen device automatically restores user access.

CLI

```
unset interface interface manage ping
unset interface interface manage ident-reset
```

Changing the Port Number

Changing the port number to which the NetScreen device listens for HTTP management traffic improves security. The default setting is port 80, the standard port number for HTTP traffic. After you change the port number, you must then type the new port number in the URL field in your Web browser when you next attempt to contact the NetScreen device. (In the following example, the administrator needs to enter http://188.30.12.2:15522.)

Example: Changing the Port Number

In this example, the IP address of the interface bound to the Trust zone is 10.1.1.1/24. To manage the NetScreen device via the WebUI on this interface, you must use HTTP. To increase the security of the HTTP connection, you change the HTTP port number from 80 (the default) to 15522.

WebUI

Configuration > Admin > Management: In the HTTP Port field, type **15522**, and then click **Apply**.

CLI

```
set admin port 15522
save
```

Changing the Admin Login Name and Password

By default, the initial login name for NetScreen devices is **netscreen**. The initial password is also **netscreen**. Because these have been widely published, NetScreen recommends you change the login name and password immediately. The login name and password are both case-sensitive. They can contain any character that can be entered from the keyboard except for ? and “. Record the new admin login name and password in a secure manner.

Warning: *Be sure to record your new password. If you forget it, you must reset the NetScreen device to its factory settings, and all your configurations will be lost. For more information, see* *“Resetting the Device to the Factory Default Settings” on page 48.*

Admin users for the NetScreen device can be authenticated using the internal database or an external auth server[8]. When the admin user logs on to the NetScreen device, it first checks the local internal database for authentication. If there is no entry present and an external auth server is connected, it then checks for a matching entry in the external auth server database. After an admin user successfully logs on to an external auth server, the NetScreen device maintains the admin’s login status locally.

Note: *For more information about admin user levels, see* *“Levels of Administration” on page 37.* *For more about using external auth servers, see “External Auth Servers” on page* ***2****-376.*

When the root admin changes any attribute of an admin user’s profile—user name, password, or privilege—any administrative session that that admin currently has open automatically terminates. If the root admin changes any of these attributes for himself, or if a root-level read/write admin or vsys read/write admin changes his own password, all of that user’s currently open admin sessions[9] terminate, other than the one in which he made the change.

8. NetScreen supports RADIUS, SecurID, and LDAP servers for admin user authentication. (For more information, see “Admin Users” on page **2**-465.) Although the root admin account must be stored on the local database, you can store root-level read/write and root-level read-only admin users on an external auth server. To store root-level and vsys-level admin users on an external auth server and query their privileges, the server must be RADIUS and you must load the netscreen.dct file on it. (See “NetScreen Dictionary File” on page **2**-381.)

9. The behavior of an HTTP or HTTPS session using the WebUI is different. Because HTTP does not support a persistent connection, any change that you make to your own user profile automatically logs you out of that and all other open sessions.

Example: Changing an Admin User's Login Name and Password

In this example, you—as the root admin—change a read/write administrator's login name from "John" to "Smith" and his password from xL7s62a1 to 3MAb99j2[10].

Note: *For information on the different levels of administrators, see "Levels of Administration" on page 37.*

WebUI

Configuration > Admin > Administrators > Edit (for John): Enter the following, and then click **OK**:

Name: Smith

New Password: 3MAb99j2

Confirm New Password: 3MAb99j2

CLI

```
unset admin user John
set admin user Smith password 3MAb99j2 privilege all
save
```

10. Instead of using actual words for passwords, which might be guessed or discovered through a dictionary attack, you can use an apparently random string of letters and numbers. To create such a string that you can easily remember, compose a sentence and use the first letter from each word. For example, "Charles will be 6 years old on November 21" becomes "Cwb6yooN21."

Example: Changing One's Own Password

Admin users with read/write privileges can change their own administrator password, but not their login name. In this example, an administrator with read/write privileges and the login name "Smith" changes his password from 3MAb99j2 to ru494Vq5.

WebUI

Configuration > Admin > Administrators > Edit (for first entry): Enter the following, and then click **OK**:

Name: Smith

New Password: ru494Vq5

Confirm New Password: ru494Vq5

CLI

```
set admin password ru494Vq5
save
```

Setting the Minimum Length of the Root Admin Password

In some corporations, one person might initially configure the device as the root admin, but another person later assumes the role of root admin and manages the device. To prevent the subsequent root admin from using short passwords that are potentially easier to decode, the initial root admin can set a minimum length requirement for the root admin's password to any number from 1 to 31.

Note that you can set the minimum password length only if you are the root admin and your own password meets the minimum length requirement you are attempting to set. Otherwise, the NetScreen device displays an error message.

To specify a minimum length for the root admin's password, enter the following command:

```
set admin password restrict length number
```

Note: *You must use the CLI to set this restriction .*

Resetting the Device to the Factory Default Settings

If the admin password is lost, you can use the following procedure to reset the NetScreen device to its default settings. The configurations will be lost, but access to the device will be restored. To perform this operation, you need to make a console connection, which is described in detail in the *NetScreen CLI Reference Guide* and the installer's guides.

Note: *By default the device recovery feature is enabled. You can disable it by entering the* ***unset admin device-reset*** *command. Also, if the NetScreen device is in FIPS mode, the recovery feature is automatically disabled.*

1. At the login prompt, type the serial number of the device.
2. At the password prompt, type the serial number again.

   The following message appears:

   !!!! Lost Password Reset !!!! You have initiated a command to reset the device to factory defaults, clearing all current configuration, keys and settings. Would you like to continue? y/n
3. Press the **Y** key.

   The following message appears:

   !! Reconfirm Lost Password Reset !! If you continue, the entire configuration of the device will be erased. In addition, a permanent counter will be incremented to signify that this device has been reset. This is your last chance to cancel this command. If you proceed, the device will return to factory default configuration, which is: System IP: 192.168.1.1; username: netscreen; password: netscreen. Would you like to continue? y/n
4. Press the **Y** key to reset the device.

   You can now log on using *netscreen* as the default username and password.

Restricting Administrative Access

You can administer NetScreen devices from one or multiple addresses of a subnet. By default, any host on the trusted interface can administer a NetScreen device. To restrict this ability to specific workstations, you must configure management client IP addresses.

Note: *The assignment of a management client IP address takes effect immediately. If you are managing the device via a network connection and your workstation is not included in the assignment, the NetScreen device immediately terminates your current session and you are no longer able to manage the device from that workstation.*

Example: Restricting Administration to a Single Workstation

In this example, the administrator at the workstation with the IP address 172.16.40.42 is the only administrator specified to manage the NetScreen device.

WebUI

Configuration > Admin > Permitted IPs: Enter the following, and then click **Add**:

IP Address / Netmask: 172.16.40.42/32

CLI

```
set admin manager-ip 172.16.40.42/32
save
```

Example: Restricting Administration to a Subnet

In this example, the group of administrators with workstations in the 172.16.40.0/24 subnet are specified to manage a NetScreen device.

WebUI

Configuration > Admin > Permitted IPs: Enter the following, and then click **Add**:

IP Address / Netmask: 172.16.40.0/24

CLI

```
set admin manager-ip 172.16.40.0 255.255.255.0
save
```

Restricting the Root Admin to Console Access

You can also require the root admin to log in to the NetScreen device through the console only. This restriction requires the root admin to have physical access to the device to log in, thus preventing unauthorized users from logging in remotely as the root admin. After you have set this restriction, the device denies access if anyone tries to log in as the root admin through other means, such as the WebUI, Telnet, or SSH, even if these management options are enabled on the ingress interface.

To restrict the access of the root admin to the console only, enter the following command:

```
set admin root access console
```

Note: *You must use the CLI to set this restriction.*

VPN Tunnels for Administrative Traffic

You can use virtual private network (VPN) tunnels to secure remote management of a NetScreen device from either a dynamically assigned or fixed IP address. Using a VPN tunnel, you can protect any kind of traffic, such as NetScreen-Security Manager, HTTP, Telnet, or SSH. (For information about creating a VPN tunnel to secure self-initiated traffic such as Security Manager reports, syslog reports or SNMP traps, see "VPN Tunnels for Self-Initiated Traffic" on page 97.)

NetScreen supports two types of VPN tunnel configurations:

- **Route-Based VPNs**: The NetScreen device uses route table entries to direct traffic to tunnel interfaces, which are bound to VPN tunnels. (For details, see Volume 5, "VPNs".)
- **Policy-Based VPNs**: The NetScreen device uses the VPN tunnel names specifically referenced in policies to direct traffic through VPN tunnels. (For details, see Volume 5, "VPNs".)

For each VPN tunnel configuration type, there are the following types of VPN tunnel:

- **Manual Key**: You manually set the three elements that define a Security Association (SA) at both ends of the tunnel: a Security Parameters Index (SPI), an encryption key, and an authentication key. To change any element in the SA, you must manually enter it at both ends of the tunnel.
- **AutoKey IKE with Preshared Key**: One or two preshared secrets—one for authentication and one for encryption—function as seed values. Using them, the IKE protocol generates a set of symmetrical keys at both ends of the tunnel; that is, the same key is used to encrypt and decrypt. At predetermined intervals, these keys are automatically regenerated.
- **AutoKey IKE with Certificates**: Using the Public Key Infrastructure (PKI), the participants at both ends of the tunnel use a digital certificate (for authentication) and an RSA public/private key pair (for encryption). The encryption is asymmetrical; that is, one key in a pair is used to encrypt and the other to decrypt.

> ***Note:*** *For a complete description of VPN tunnels, see Volume 5, "VPNs". For more information on NetScreen-Remote, refer to the NetScreen-Remote User's Guide.*

If you use a policy-based VPN configuration, you must create an address book entry with the IP address of an interface in any zone other than the one to which the outgoing interface is bound. You can then use that as the source address in policies referencing the VPN tunnel. This address also serves as the end entity address for the remote IPSec peer. If you are using a route-based VPN configuration, such an address book entry is unnecessary.

Example: Administration through a Route-Based Manual Key VPN Tunnel

In this example, you set up a route-based Manual Key VPN tunnel to provide confidentiality for administrative traffic. The tunnel extends from the NetScreen-Remote VPN client running on an admin's workstation at 10.1.1.56 to ethernet1 (10.1.1.1/24). The admin's workstation and ethernet1 are both in the Trust zone. You name the tunnel "tunnel-adm". You create an unnumbered tunnel interface, name it tunnel.1, and bind it to the Trust zone and to the VPN tunnel "tunnel-adm".

The NetScreen device uses the internal IP address configured on the NetScreen-Remote—10.10.10.1—as the destination address to target beyond the peer gateway address of 10.1.1.56. You define a route to 10.10.10.1/32 through tunnel.1. A policy is unnecessary because of the following two reasons:

- The VPN tunnel protects administrative traffic that terminates at the NetScreen device itself instead of passing through the device to another security zone.
- This is a route-based VPN, meaning that the route lookup—not a policy lookup—links the destination address to the tunnel interface, which is bound to the appropriate VPN tunnel.

Note: *Compare this example with* *"Example: Administration through a Policy-Based Manual Key VPN Tunnel" on page 58**.*

The NetScreen-Remote uses the IP address of ethernet3—1.1.1.1—as the destination address to target beyond the remote gateway at 10.1.1.1. The NetScreen-Remote configuration specifies the remote party ID type as "IP address" and the protocol as "All".

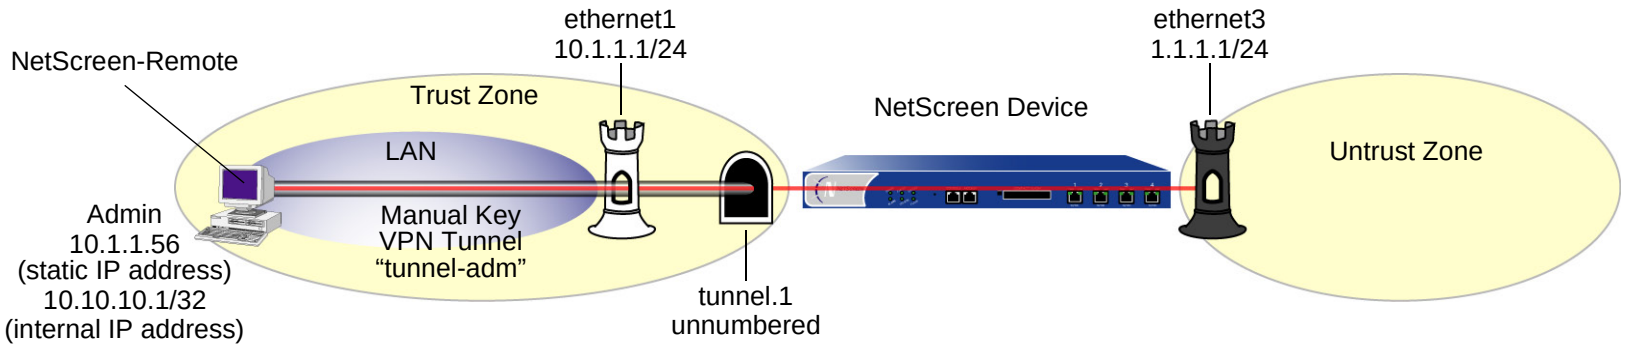

WebUI

1. Interfaces

Network > Interfaces > Edit (ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

Network > Interfaces > New Tunnel IF: Enter the following, and then click **OK**:

Tunnel Interface Name: Tunnel.1

Zone (VR): Trust (trust-vr)

Unnumbered: (select)

Interface: ethernet1(trust-vr)[11]

11. The unnumbered tunnel interface borrows the IP address of the specified security zone interface.

2. VPN

VPNs > Manual Key > New: Enter the following, and then click **OK**:

VPN Tunnel Name: tunnel-adm

Gateway IP: 10.1.1.56

Security Index (HEX number): 5555 (Local) 5555 (Remote)

Outgoing Interface: ethernet1

ESP-CBC: (select)

Encryption Algorithm: DES-CBC

Generate Key by Password[12]: netscreen1

Authentication Algorithm: MD5

Generate Key by Password: netscreen2

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

Bind to Tunnel Interface: (select), Tunnel.1

3. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 10.10.10.1/32

Gateway: (select)

Interface: Tunnel.1

Gateway IP Address: 0.0.0.0

12. Because NetScreen-Remote processes passwords into keys differently than other NetScreen products do, after you configure the tunnel do the following: (1) Return to the Manual Key Configuration dialog box (click **Edit** in the Configure column for “tunnel-adm”); (2) copy the generated hexadecimal keys; and (3) use those hexadecimal keys when configuring the NetScreen-Remote end of the tunnel.

CLI

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
set interface tunnel.1 zone trust

set interface tunnel.1 ip unnumbered interface ethernet1[13]
```

2. VPN

```
set vpn tunnel-adm manual 5555 5555 gateway 10.1.1.56 outgoing ethernet1 esp
    des password netscreen1 auth md5 password netscreen2[14]
set vpn tunnel-adm bind interface tunnel.1
```

3. Route

```
set vrouter trust-vr route 10.10.10.1/32 interface tunnel.1
save
```

13. The unnumbered tunnel interface borrows the IP address of the specified security zone interface.

14. Because NetScreen-Remote processes passwords into keys differently than other NetScreen products do, after you configure the tunnel do the following: (1) Type **get vpn admin-tun**; (2) copy the hexadecimal keys generated by “netscreen1” and “netscreen2”; and (3) use those hexadecimal keys when configuring the NetScreen-Remote end of the tunnel.

NetScreen-Remote Security Policy Editor

1. Click **Options** > **Global Policy Settings**, and select the **Allow to Specify Internal Network Address** check box.
2. Click **Options > Secure > Specified Connections**.
3. Click **Add a new connection**, and type **Admin** next to the new connection icon that appears.
4. Configure the connection options:

   Connection Security: Secure

   Remote Party Identity and Addressing:

   ID Type: IP Address, 1.1.1.1

   Protocol: All

   Connect using Secure Gateway Tunnel: (select)

   ID Type: IP Address, 10.1.1.1
5. Click the **PLUS** symbol, located to the left of the unix icon, to expand the connection policy.
6. Click **My Identity**, in the Select Certificate drop-down list, choose **None**, and in the Internal Network IP Address, type **10.10.10.1**.
7. Click **Security Policy**, and select **Use Manual Keys**.
8. Click the **PLUS** symbol, located to the left of the Security Policy icon, and then the **PLUS** symbol to the left of Key Exchange (Phase 2) to expand the policy further.
9. Click **Proposal 1**, and select the following IPSec Protocols:

   Encapsulation Protocol (ESP): (select)

   Encrypt Alg: DES

   Hash Alg: MD5

   Encapsulation: Tunnel

10. Click **Inbound Keys**, and in the Security Parameters Index field, type **5555**.
11. Click **Enter Key**, enter the following[15], and then click **OK**:

    Choose key format: Binary

    ESP Encryption Key: dccbee96c7e546bc

    ESP Authentication Key: dccbe9e6c7e546bcb0b667794ab7290c
12. Click **Outbound Keys**, and in the Security Parameters Index field, type **5555**.
13. Click **Enter Key**, enter the following[15], and then click **OK**:

    Choose key format: Binary

    ESP Encryption Key: dccbee96c7e546bc

    ESP Authentication Key: dccbe9e6c7e546bcb0b667794ab7290c
14. Click **Save**.

15. These are the two generated keys that you copied after configuring the NetScreen device.

Example: Administration through a Policy-Based Manual Key VPN Tunnel

In this example, you set up a policy-based Manual Key VPN tunnel for administrative traffic. The tunnel extends from the NetScreen-Remote VPN client running on an admin's workstation at 10.1.1.56 to ethernet1 (10.1.1.1/24). The admin's workstation and ethernet1 are both in the Trust zone. You name the tunnel "tunnel-adm" and bind it to the Trust zone.

The NetScreen device uses the internal IP address configured on the NetScreen-Remote—10.10.10.1—as the destination address to target beyond the peer gateway address of 10.1.1.56. You define a Trust zone address book entry specifying 10.10.10.1/32, and an Untrust zone address book entry specifying the IP address of ethernet3. Although the address of the ethernet3 interface is 1.1.1.1/24, the address you create has a 32-bit netmask: 1.1.1.1/32. You use this address and the internal address of the admin's workstation in the policy you create referencing the tunnel "tunnel-adm". A policy is necessary because this is a policy-based VPN, meaning that the policy lookup—not a route lookup—links the destination address to the appropriate VPN tunnel.

You must also define a route to 10.10.10.1/32 through ethernet1.

> ***Note:*** *Compare this example with "Example: Administration through a Route-Based Manual Key VPN Tunnel" on page 52.*

The NetScreen-Remote uses the IP address 1.1.1.1 as the destination address to target beyond the remote gateway at 10.1.1.1. The NetScreen-Remote tunnel configuration specifies the remote party ID type as IP address and the protocol as "All".

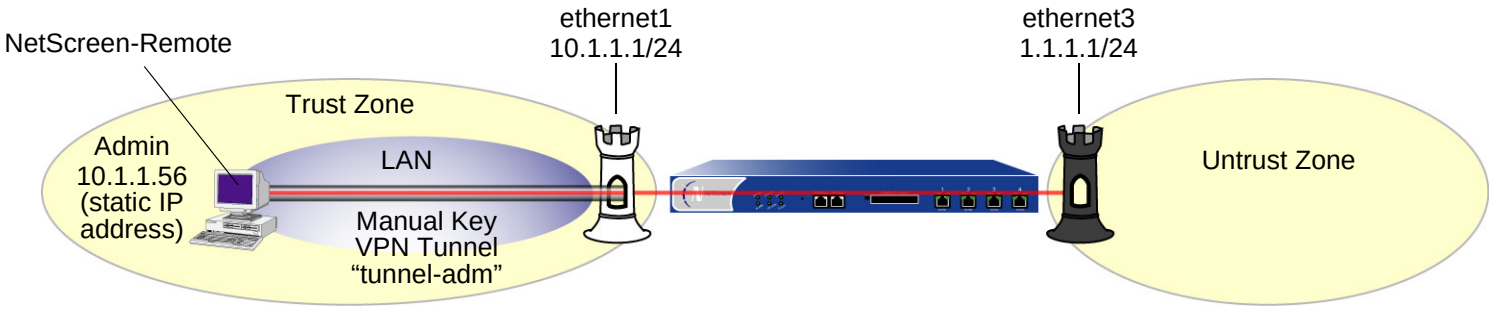

WebUI

1. Interfaces

Network > Interfaces > Edit (ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust
Static IP: (select this option when present)
IP Address/Netmask: 10.1.1.1/24

Select the following, and then click **OK**:
Interface Mode: NAT

Network > Interfaces > Edit (ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust
Static IP: (select this option when present)
IP Address/Netmask: 1.1.1.1/24

2. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Untrust-IF
IP Address/Domain Name:
IP/Netmask: (select), 1.1.1.1/32
Zone: Untrust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: admin
IP Address/Domain Name:
IP/Netmask: (select), 10.10.10.1/32
Zone: Trust

3. VPN

VPNs > Manual Key > New: Enter the following, and then click **OK**:

VPN Tunnel Name: tunnel-adm

Gateway IP: 10.1.1.56

Security Index (HEX Number): 5555 (Local) 5555 (Remote)

Outgoing Interface: ethernet1

ESP-CBC: (select)

Encryption Algorithm: DES-CBC

Generate Key by Password[16]: netscreen1

Authentication Algorithm: MD5

Generate Key by Password: netscreen2

4. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 10.10.10.1/32

Gateway: (select)

Interface: ethernet1

Gateway IP Address: 0.0.0.0

16. Because NetScreen-Remote processes passwords into keys differently than other NetScreen products do, after you configure the tunnel do the following: (1) Return to the Manual Key Configuration dialog box (click **Edit** in the Configure column for “tunnel-adm”); (2) copy the generated hexadecimal keys; and (3) use those hexadecimal keys when configuring the NetScreen-Remote end of the tunnel.

5. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), admin

Destination Address:

Address Book Entry: (select), Untrust-IF

Service: Any

Action: Tunnel

Tunnel:

VPN: tunnel-adm

Modify matching bidirectional VPN policy: (select)

Position at Top: (select)

CLI

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
```

2. Addresses

```
set address trust admin 10.10.10.1/32
set address untrust Untrust-IF 1.1.1.1/32
```

3. VPN

```
set vpn tunnel-adm manual 5555 5555 gateway 10.1.1.56 outgoing ethernet1 esp
    des password netscreen1 auth md5 password netscreen2
```
[17]

4. Route

```
set vrouter trust-vr route 10.10.10.1/32 interface ethernet1
```

5. Policies

```
set policy top from trust to untrust admin Untrust-IF any tunnel vpn tunnel-adm
set policy top from untrust to trust Untrust-IF admin any tunnel vpn tunnel-adm
save
```

17. Because NetScreen-Remote processes passwords into keys differently than other NetScreen products do, after you configure the tunnel do the following: (1) Type **get vpn admin-tun**; (2) copy the hexadecimal keys generated by “netscreen1” and “netscreen2”; and (3) use those hexadecimal keys when configuring the NetScreen-Remote end of the tunnel.

NetScreen-Remote Security Policy Editor

1. Click **Options > Secure > Specified Connections**.
2. Click **Add a new connection**, and type **Admin** next to the new connection icon that appears.
3. Configure the connection options:

   Connection Security: Secure

   Remote Party Identity and Addressing:

   ID Type: IP Address, 1.1.1.1

   Protocol: All

   Connect using Secure Gateway Tunnel: (select)

   ID Type: IP Address, 10.1.1.1
4. Click the **PLUS** symbol, located to the left of the unix icon, to expand the connection policy.
5. Click **My Identity**, and in the Select Certificate drop-down list, choose **None**.
6. Click **Security Policy**, and select **Use Manual Keys**.
7. Click the **PLUS** symbol, located to the left of the Security Policy icon, and then the **PLUS** symbol to the left of Key Exchange (Phase 2) to expand the policy further.
8. Click **Proposal 1**, and select the following IPSec Protocols:

   Encapsulation Protocol (ESP): (select)

   Encrypt Alg: DES

   Hash Alg: MD5

   Encapsulation: Tunnel

9. Click **Inbound Keys**, and in the Security Parameters Index field, type **5555**.
10. Click **Enter Key**, enter the following[18], and then click **OK**:

    Choose key format: Binary

    ESP Encryption Key: dccbee96c7e546bc

    ESP Authentication Key: dccbe9e6c7e546bcb0b667794ab7290c
11. Click **Outbound Keys**, and in the Security Parameters Index field, type **5555**.
12. Click **Enter Key**, enter the following[15], and then click **OK**:

    Choose key format: Binary

    ESP Encryption Key: dccbee96c7e546bc

    ESP Authentication Key: dccbe9e6c7e546bcb0b667794ab7290c
13. Click **Save**.

18. These are the two generated keys that you copied after configuring the NetScreen device.

2

Monitoring NetScreen Devices

This chapter discusses the following topics about monitoring NetScreen devices:

- "Storing Log Information" on page 66
- "Event Log" on page 67
  - "Viewing the Event Log" on page 68
  - "Sorting and Filtering the Event Log" on page 70
  - "Downloading the Event Log" on page 71
- "Traffic Log" on page 72
  - "Viewing the Traffic Log" on page 74
  - "Downloading the Traffic Log" on page 76
- "Self Log" on page 77
  - "Viewing the Self Log" on page 77
  - "Downloading the Self Log" on page 80
- "Asset Recovery Log" on page 81
- "Traffic Alarms" on page 82
- "Syslog" on page 87
  - "WebTrends" on page 89
- "SNMP" on page 91
  - "Implementation Overview" on page 94
- "VPN Tunnels for Self-Initiated Traffic" on page 97
- "Counters" on page 120

STORING LOG INFORMATION

All NetScreen devices allow you to store event and traffic log data internally (in flash storage) and externally (in a number of locations). Although storing log information internally is convenient, the amount of memory is limited. When the internal storage space completely fills up, the NetScreen device begins overwriting the oldest log entries with the latest ones. If this first-in-first-out (FIFO) mechanism occurs before you save the logged information, you can lose data. To mitigate such data loss, you can store event and traffic logs externally in a syslog or WebTrends server, or in the NetScreen-Global PRO database. The NetScreen device sends new event and traffic log entries to an external storage location every second.

The following list provides the possible destinations for logged data:

- **Console:** A useful destination for all log entries to appear when you are troubleshooting a NetScreen device through the console. Optionally, you might elect to have only alarm messages (critical, alert, emergency) appear here to alert you immediately if you happen to be using the console at the time an alarm is triggered.
- **Internal:** The internal database on a NetScreen device is a convenient destination for log entries, but of limited space.
- **Email:** A convenient method for sending event and traffic logs to remote administrators.
- **SNMP:** In addition to the transmission of SNMP traps, a NetScreen device can also send alarm messages (critical, alert, emergency) from its event log to an SNMP community.
- **Syslog:** All event and traffic log entries that a NetScreen device can store internally, it can also send to a syslog server. Because syslog servers have a much greater storage capacity than the internal flash storage on a NetScreen device, sending data to a syslog server can mitigate data loss that might occur when log entries exceed the maximum internal storage space. Syslog stores alert- and emergency-level events in the security facility that you specify, and all other events (including traffic data) in the facility you specify.
- **WebTrends:** Allows you to view log data for critical-, alert-, and emergency-level events in a more graphical format than syslog, which is a text-based tool.
- **CompactFlash (PCMCIA):** The advantage of this destination is portability. After storing data on a CompactFlash card, you can physically remove the card from the NetScreen device and store it or load it on another device.

EVENT LOG

NetScreen provides an event log for monitoring system events such as admin-generated configuration changes, and self-generated messages and alarms regarding operational behavior and attacks. The NetScreen device categorizes system events by the following severity levels:

- **Emergency:** Generates messages on SYN attacks, Tear Drop attacks, and Ping of Death attacks. For more information on these types of attacks, see Volume 4, “Attack Detection and Defense Mechanisms”.
- **Alert:** Generates messages for multiple user authentication failures and other firewall attacks not included in the emergency category.
- **Critical:** Generates messages for URL blocks, traffic alarms, high availability (HA) status changes, and global communications.
- **Error:** Generates messages for admin log on failures.
- **Warning:** Generates messages for admin logins and logouts, failures to log on and log out, and user authentication failures, successes, and timeouts.
- **Notification:** Generates messages for link status changes, traffic logs, and configuration changes.
- **Information:** Generates any kind of message not specified in other categories.
- **Debugging:** Generates all messages.

The event log displays the date, time, level and description of each system event. You can view system events for each category stored in flash storage on the NetScreen device through the WebUI or the CLI. You can also open or save the file to the location you specify, and then use an ASCII text editor (such as Notepad or WordPad) to view the file. Alternatively, you can send them to an external storage space (see “Storing Log Information” on page 66).

Note: *For detailed information about the messages that appear in the event log, refer to the* NetScreen Message Log Reference Guide.

Viewing the Event Log

You can view the event log stored in the device by using the CLI or the WebUI. You can display log entries by severity level and search the event log by keyword in both the WebUI and CLI.

To display the event log by severity level, do either of the following:

WebUI

Reports > System Log > Event: Select a severity level from the Log Level drop-down list.

CLI

```
get event level { emergency | alert | critical | error | warning | notification
    | information | debugging }
```

To search the event log by keyword, do either of the following:

WebUI

Reports > System Log > Event: Type a word or word phrase up to 15 characters in length in the search field, and then click **Search**.

CLI

```
get event include word_string
```

Example: Viewing the Event Log by Severity Level and Keyword

In this example, you view event log entries with a “warning” severity level and do a search for the keyword AV.

WebUI

Reports > System Log > Event:

Log Level: Warning (select)

Search: AV Click **Search**.

CLI

```
get event level warning include av
```

```
Date       Time     Module Level Type Description
2003-05-16 15:56:20 system warn 00547 AV scanman is removed.
2003-05-16 09:45:52 system warn 00547 AV test1 is removed.
Total entries matched = 2
```

Sorting and Filtering the Event Log

Additionally, you can use the CLI to sort or filter the event log based on the following criteria:

- **Source or Destination IP Address:** Only certain events contain a source or destination IP address, such as land attacks or ping flood attacks. When you sort event logs by source or destination IP address, the device sorts and displays only the event logs that contain source or destination IP addresses. It ignores all event logs with no source or destination IP address.

  When you filter the event log by by specifying a source or destination IP address or range of addresses, the device displays the log entries for the specified source or destination IP address, or range of addresses.

- **Date:** You can sort the event log by date only, or by date and time. When you sort log entries by date and time, the device lists the log entries in descending order by date and time.

  You can also filter event log entries by specifying a start date, an end date, or a date range. When you specify a start date, the device displays log entries with date/time stamps after the start date. When you specify an end date, the device displays log entries with date/time stamps before the end date.

- **Time:** When you sort logs by time, the device displays the log entries in descending order by time, regardless of the date. When you specify a start time, the device displays log entries with time stamps after the specified start time, regardless of the date. When you specify an end time, the device displays log entries with time stamps before the specified end time, regardless of the date. When you specify both a start and end time, the device displays log entries with time stamps within the specified time period.
- **Message Type ID Number:** You can display event log entries for a specific message type ID number, or you can display log entries with message type ID numbers within a specified range. The device displays log entries with the message type ID number(s) you specified, in descending order by date and time.

Example: Sorting Event Log Entries by IP Address

In this example you view event log entries that contain source IP addresses within the range 10.100.0.0 to 10.200.0.0. The log entries are also sorted by source IP address.

CLI

```
get event sort-by src-ip 10.100.0.0-10.200.0.0
```

Downloading the Event Log

You can open or save the event log to the location you specify, and then use an ASCII text editor (such as Notepad or WordPad) to view the file. Alternatively, you can send the log entries to an external storage space (see "Storing Log Information" on page 66). You can download the entire event log through the WebUI. You can download the event log by severity level through the CLI.

Example: Downloading the Event Log

In this example, you download the event log to the local directory "C:\netscreen\logs". You name the file "evnt07-02.txt".

WebUI

1. Reports > System Log > Event: Click **Save**.

   The File Download dialog box prompts you to open the file (using an ASCII editor) or save it to disk.

2. Select the **Save** option, and then click **OK**.

   The File Download dialog box prompts you to choose a directory.

3. Specify **C:\netscreen\logs**, name the file "evnt07-02.txt", and then click **Save**.

Example: Downloading the Event Log for Critical Events

In this example, you download the critical events entered in the event log to the root directory of a TFTP server at the IP address 10.10.20.200 (CLI). You name the file "crt_evnt07-02.txt".

CLI

```
get event level critical > tftp 10.10.20.200 crt_evnt07-02.txt
```

TRAFFIC LOG

The NetScreen device can monitor and record traffic that it permits or denies based on previously configured policies. You can enable the logging option for each policy that you configure. When you enable the logging option for a policy that permits traffic, the device records the traffic allowed by that policy. When you enable the logging option for a policy that denies traffic, the device records traffic that attempted to pass through the device, but was dropped because of that policy.

A traffic log notes the following elements for each session:

- Date and time that the connection started
- Source address and port number
- Translated source address and port number
- Destination address and port number
- The duration of the session
- The service used in the session

To log all traffic that a NetScreen device receives, you must enable the logging option for all policies. To log specific traffic, enable logging only on policies that apply to that traffic. To enable the logging option on a policy, do either of the following:

WebUI

Policies > (From *src_zone*, To *dst_zone*) New : Select **Logging** and then click **OK**.

CLI

```
set policy from src_zone to dst_zone src_addr dst_addr service action log
```

In addition to logging traffic for a policy, the device can also maintain a count in bytes of all network traffic to which the policy was applied. When you enable the counting option, the device includes the following information when it displays traffic log entries

- The number of bytes transmitted from a source to a destination
- The number of bytes transmitted from a destination to a source

To enable the counting option on a policy, do either of the following:

WebUI

Policies > (From *src_zone*, To *dst_zone*) New > Advanced: Select **Counting**, click **Return**, and then click **OK**.

CLI

```
set policy from src_zone to dst_zone src_addr dst_addr service action log count
```

Viewing the Traffic Log

You can view traffic log entries stored in flash storage on the NetScreen device through either the CLI or WebUI:

WebUI

Policies >  (for policy ID *number*)

or

Reports > Policies >  (for policy ID *number*)

CLI

```
get log traffic policy number
```

Example: Viewing Traffic Log Entries

In this example, you view the traffic log details of a policy with ID number 3, and for which you have previously enabled logging:

WebUI

Policies: Click the  icon for the policy with ID number 3.

The following information appears:

Date/Time	Source Address/Port	Destination Address/Port	Translated Source Address/Port	Translated Destination Address/Port	Service	Duration	Bytes Sent	Bytes Received
2003-01-09 21:33:43	1.1.1.1:1046	10.1.1.5:80	1.1.1.1:1046	10.1.1.5:80	HTTP	1800 sec.	326452	289207

Sorting and Filtering the Traffic Log

Similar to the event log, when you use the CLI to view the traffic log, you can sort or filter the log entries according to the following criteria:

- **Source or Destination IP Address:** You can sort the traffic log by source or destination IP address. YOu can also filter the traffic log by specifying a source or destination IP address or range of addresses.
- **Date:** You can sort the traffic log by date only, or by date and time. The device lists the log entries in descending order by date and time.

  You can also filter event log entries by specifying a start date, an end date, or a date range. When you specify a start date, the device displays log entries with date/time stamps after the start date. When you specify an end date, the device displays log entries with date/time stamps before the end date.
- **Time:** When you sort the traffic log by time, the device displays the log entries in descending order by time, regardless of the date. When you specify a start time, the device displays log entries with time stamps after the specified start time, regardless of the date. When you specify an end time, the device displays log entries with time stamps before the specified end time, regardless of the date. When you specify both a start and end time, the device displays log entries with time stamps within the specified time period.

Example: Sorting the Traffic Log by Time

In this example you view traffic logs sorted by time with a time stamp after 1:00 a.m.

CLI

```
get log traffic sort-by time start-time 01:00:00
```

Downloading the Traffic Log

You can also open or save the log to the location you specify, and then use an ASCII text editor (such as Notepad or WordPad) to view the file.

Alternatively, you can send traffic log entries to an external storage space (see "Storing Log Information" on page 66). The NetScreen device makes an entry in the traffic log when a session terminates. When you enable the NetScreen device to send traffic log entries to an external storage location, it sends new entries every second. Because the NetScreen device makes a traffic log entry when a session closes, the NetScreen device sends traffic log entries for all sessions that have closed within the past second. You can also include traffic log entries with event log entries sent by e-mail to an admin.

Example: Downloading a Traffic Log

In this example, you download the traffic log for a policy with ID number 12. For the WebUI, you download it to the local directory "C:\netscreen\logs". For the CLI, you download it to the root directory of a TFTP server at the IP address 10.10.20.200. You name the file "traf_log11-21-02.txt".

WebUI

1. Reports > Policies >  (for policy ID 12): Click **Save**.

   The File Download dialog box prompts you to open the file (using an ASCII editor) or save it to disk.
2. Select the **Save** option, and then click **OK**.

   The File Download dialog box prompts you to choose a directory.
3. Specify C:\netscreen\logs, name the file traf_log11-21-02.txt, and then click **Save**.

CLI

```
get log traffic policy 12 > tftp 10.10.20.200 traf_log11-21-02.txt
```

SELF LOG

NetScreen provides a self log to monitor and record all packets terminated at the NetScreen device. For example, if you disable some management options on an interface—such as WebUI, SNMP, and ping—and HTTP, SNMP, or ICMP traffic is sent to that interface, entries appear in the self log for each dropped packet.

To activate the self log, do one of the following:

WebUI

Configuration > Report Settings > Log Settings: Select the **Log Packets Terminated to Self** check box, and then click **Apply**.

CLI

```
set firewall log-self
```

Similar to the traffic log, the self log displays the date, time, source address/port, destination address/port, duration, and service for each dropped packet terminating at the NetScreen device.

Viewing the Self Log

You can view the self log, which is stored in flash storage on the NetScreen device, through either the CLI or WebUI.

WebUI

Reports > System Log > Self

CLI

```
get log self
```

Sorting and Filtering the Self Log

Similar to the event and traffic logs, when you use the CLI to view the self log, you can sort or filter the log entries according to the following criteria:

- **Source or Destination IP Address:** You can sort the self log by source or destination IP address. YOu can also filter the self log by specifying a source or destination IP address or range of addresses.
- **Date:** You can sort the self log by date only, or by date and time. The device lists the log entries in descending order by date and time.

  You can also filter self log entries by specifying a start date, an end date, or a date range. When you specify a start date, the device displays log entries with date/time stamps after the start date. When you specify an end date, the device displays log entries with date/time stamps before the end date.
- **Time:** When you sort the self log by time, the NetScreen device displays the log entries in descending order by time, regardless of the date. When you specify a start time, the device displays log entries with time stamps after the specified start time, regardless of the date. When you specify an end time, the device displays log entries with time stamps before the specified end time, regardless of the date. When you specify both a start and end time, the device displays log entries with time stamps within the specified time period.

Example: Filtering the Self Log by Time

In this example, you filter self log entries by the end time. The NetScreen device displays log entries with time stamps before the specified end time:

CLI

```
get log self end-time 16:32:57

================================================================================
Date       Time       Duration Source IP        Port Destination IP   Port Serv
================================================================================
2003-08-21 16:32:57    0:00:00 10.100.25.1         0 224.0.0.5           0 OSPF
2003-08-21 16:32:47    0:00:00 10.100.25.1         0 224.0.0.5           0 OSPF

           Total entries matched = 2
```

Downloading the Self Log

You can also save the log as a text file to a location you specify, and then use an ASCII text editor (such as Notepad or WordPad) to view it.

Example: Downloading the Self Log

In this example, you download a self log to the local directory “C:\netscreen\logs” (WebUI) or to the root directory of a TFTP server at the IP address 10.10.20.200 (CLI). You name the file “self_log07-03-02.txt”.

WebUI

1. Reports > System Log > Self: Click **Save**.

   The File Download dialog box prompts you to open the file (using an ASCII editor) or save it to disk.
2. Select the **Save** option, and then click **OK**.

   The File Download dialog box prompts you to choose a directory.
3. Specify C:\netscreen\logs, name the file self_log07-03-02.txt, and then click **Save**.

CLI

```
get log self > tftp 10.10.20.200 self_log07-03-02.txt
```

ASSET RECOVERY LOG

NetScreen provides an asset recovery log to display information about each time the device is returned to its default settings using the asset recovery procedure (see "Resetting the Device to the Factory Default Settings" on page 48). In addition to viewing the asset recovery log through the WebUI or CLI, you can also open or save the file to the location you specify. Use an ASCII text editor (such as Notepad) to view the file.

Example: Downloading the Asset Recovery Log

In this example, you download the asset recovery log to the local directory "C:\netscreen\logs" (WebUI) or to the root directory of a TFTP server at the IP address 10.10.20.200 (CLI). You name the file "sys_rst.txt".

WebUI

1. Reports > System Log > Asset Recovery: Click **Save**.

   The File Download dialog box prompts you to open the file (using an ASCII editor) or save it to disk.
2. Select the **Save** option, and then click **OK**.

   The File Download dialog box prompts you to choose a directory.
3. Specify C:\netscreen\logs, name the file sys_rst.txt, and then click **Save**.

CLI

```
get log asset-recovery > tftp 10.10.20.200 sys_rst.txt
```

TRAFFIC ALARMS

The NetScreen device supports traffic alarms when traffic exceeds thresholds that you have defined in policies. You can configure the NetScreen device to alert you through one or more of the following methods whenever the NetScreen device generates a traffic alarm:

- Console
- Internal (Event Log)
- E-mail
- SNMP
- Syslog
- WebTrends
- NetScreen-Global PRO

You set alarm thresholds to detect anomalous activity. To know what constitutes anomalous activity, you must first establish a baseline of normal activity. To create such a baseline for network traffic, you must observe traffic patterns over a period of time. Then, after you have determined the amount of traffic that you consider as normal, you can set alarm thresholds above that amount. Traffic exceeding that threshold triggers an alarm to call your attention to a deviation from the baseline. You can then evaluate the situation to determine what caused the deviation and whether you need to take action in response.

You can also use traffic alarms to provide policy-based intrusion detection and notification of a compromised system. Examples of the use of traffic alarms for these purposes are provided below.

Example: Policy-Based Intrusion Detection

In this example, there is a Web server with IP address 211.20.1.5 (and name “web1”) in the DMZ zone. You want to detect any attempts from the Untrust zone to access this Web server via Telnet. To accomplish this, you create a policy denying Telnet traffic from any address in the Untrust zone destined to the Web server named web1 in the DMZ zone, and you set a traffic alarm threshold at 64 bytes. Because the smallest size of IP packet is 64 bytes, even one Telnet packet attempting to reach the Web server from the Untrust zone will trigger an alarm.

WebUI

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: web1

IP Address/Domain Name:

IP/Netmask: (select), 211.20.1.5/32

Zone: DMZ

Policies > (From: Untrust, To: DMZ) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), web1

Service: Telnet

Action: Deny

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

Counting: (select)

Alarm Threshold: 64 Bytes/Sec, 0 Kbytes/Min

CLI

```
set address dmz web1 211.20.1.5/32
set policy from untrust to dmz any web1 telnet deny count alarm 64 0
save
```

Example: Compromised System Notification

In this example, you use traffic alarms to provide notification of a compromised system. You have an FTP server with IP address 211.20.1.10 (and name ftp1) in the DMZ zone. You want to allow FTP-get traffic to reach this server. You don't want traffic of any kind to originate from the FTP server. The occurrence of such traffic would indicate that the system has been compromised, perhaps by a virus similar to the NIMDA virus. You define an address for the FTP server in the Global zone, so that you can then create two global policies.

WebUI

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: ftp1

IP Address/Domain Name:

IP/Netmask: (select), 211.20.1.10/32

Zone: Global

Policies > (From: Global, To: Global) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), ftp1

Service: FTP-Get

Action: Permit

Policies > (From: Global, To: Global) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), ftp1

Destination Address:

Address Book Entry: (select), Any

Service: ANY

Action: Deny

> Advanced: Enter the following, and then click **Return** to set the advanced options and return to the basic configuration page:

Counting: (select)

Alarm Threshold: 64 Bytes/Sec, 0 Kbytes/Min

CLI

```
set address global ftp1 211.20.1.10/32
set policy global any ftp1 ftp-get permit
set policy global ftp1 any any deny count alarm 64 0
save
```

Example: Sending E-mail Alerts

In this example, you set up notification by e-mail alerts when there is an alarm. The mail server is at 172.16.10.254, the first e-mail address to be notified is jharker@netscreen.com, and the second address is driggs@netscreen.com. The NetScreen device includes traffic logs with event logs sent via e-mail.

WebUI

Configuration > Report Settings > Email: Enter the following information, then click **Apply**:

Enable E-Mail Notification for Alarms: (select)

Include Traffic Log: (select)

SMTP Server Name: 172.16.10.254[1]

E-Mail Address 1: jharker@netscreen.com

E-Mail Address 2: driggs@netscreen.com

CLI

```
set admin mail alert
set admin mail mail-addr1 jharker@netscreen.com
set admin mail mail-addr2 driggs@netscreen.com
set admin mail server-name 172.16.10.254
set admin mail traffic-log
save
```

1. If you have DNS enabled, you can also use a host name for the mail server, such as mail.netscreen.com.

SYSLOG

A NetScreen device can generate syslog messages for system events at predefined severity levels (see the list of severity levels in “Event Log” on page 67), and optionally for traffic that policies permit across a firewall. It sends these messages to up to four designated syslog hosts running on UNIX/Linux systems. For each syslog host, you can specify the following:

- Whether the NetScreen device includes traffic log entries, event log entries, or both traffic and event log entries
- Whether to send traffic through a VPN tunnel to the syslog server and—if through a VPN tunnel—which interface to use as the source interface (for examples, see “Example: Self-Generated Traffic through a Route-Based Tunnel” on page 99 and “Example: Self-Generated Traffic through a Policy-Based Tunnel” on page 109)
- The port to which the NetScreen device sends syslog messages
- The security facility, which classifies and sends emergency and alert level messages to the Syslog host; and the regular facility, which classifies and sends all other messages for events unrelated to security

By default, the NetScreen device sends messages to syslog hosts via UDP (port 514). To increase the reliability of the message delivery, you can change the transport protocol for each syslog host to TCP.

You can use syslog messages to create e-mail alerts for the system administrator, or to display messages on the console of the designated host using UNIX syslog conventions.

Note: *On UNIX/Linux platforms, modify the /etc/rc.d/init.d/syslog file so that syslog retrieves information from the remote source (syslog -r).*

Example: Enabling Multiple Syslog Servers

In this example, you configure the NetScreen device to send event and traffic logs via TCP to three syslog servers at the following IP addresses/port numbers: 1.1.1.1/1514, 2.2.2.1/2514, and 3.3.3.1/3514. You set both the security and facility levels to Local0.

WebUI

Configuration > Report Settings > Syslog: Enter the following, and then click **Apply**:

Enable syslog messages: Select this option to send logs to the specified syslog servers.

No.: Select 1, 2, and 3 to indicate you are adding 3 syslog servers.

IP/ Hostname: 1.1.1.1, 2.2.2.1, 3.3.3.1

Port: 1514, 2514, 3514

Security Facility: Local0, Local0, Local0

Facility: Local0, Local0, Local0

Event Log: (select)

Traffic Log: (select)

TCP: (select)

CLI

```
set syslog config 1.1.1.1 port 1514
set syslog config 1.1.1.1 log all
set syslog config 1.1.1.1 facilities local0 local0
set syslog config 1.1.1.1 transport tcp
set syslog config 2.2.2.1 port 2514
set syslog config 2.2.2.1 log all
set syslog config 2.2.2.1 facilities local0 local0
set syslog config 2.2.2.1 transport tcp
set syslog config 3.3.3.1 port 3514
```

```
set syslog config 3.3.3.1 log all
set syslog config 3.3.3.1 facilities local0 local0
set syslog config 2.2.2.1 transport tcp
set syslog enable
save
```

WebTrends

WebTrends offers a product called the WebTrends Firewall Suite that allows you to customize syslog reports to display the information you want in a graphical format. You can create reports on all events and severity levels or focus on an area such as firewall attacks (emergency-level events).

Note: *The WebTrends Syslog Server and the WebTrends Firewall Suite must run on the same Windows NT system. You must have administrator rights to configure it.*

You can also send WebTrends messages through a VPN tunnel. In the WebUI, use the **Use Trust Zone Interface as Source IP for VPN** option. In the CLI, use the **set webtrends vpn** command.

Example: Enabling Syslog and WebTrends for Notification Events

In the following example, you set up the syslog facility to send notification messages to port 514 on a WebTrends syslog server at 172.10.16.25. The security and facility levels are set to Local0. Traffic logs are included with the system event messages.

WebUI

1. Syslog Settings

   Configuration > Report Settings > Syslog: Enter the following, and then click **Apply**:

   Enable syslog messages: (select)

   Include Traffic Log: (select)

   Syslog Host Name/Port: 172.10.16.25/514[2]

2. The syslog host port number must match the WebTrends port number.

Security Facility: Local0

Facility: Local0

2. WebTrends Settings

Configuration > Report Settings > WebTrends: Enter the following, and then click **Apply**:

Enable WebTrends Messages: (select)

WebTrends Host Name/Port: 172.10.16.25/514

3. Severity Levels

Configuration > Report Settings > Log Settings: Enter the following, then click **Apply**:

WebTrends Notification: (select)

Syslog Notification: (select)

Note: *When you enable syslog and WebTrends on a NetScreen device running in Transparent mode, you must set up a static route. See "Routing Tables and Static Routing" on page* ***2****-29.*

CLI

1. Syslog Settings

```
set syslog config 172.10.16.25 port 514
set syslog config 172.10.16.25 local0 local0
set syslog config 172.10.16.25 log all
set syslog enable
```

2. WebTrends Settings

```
set webtrends host-name 172.10.16.25
set webtrends port 514
set webtrends enable
```

3. Severity Levels

```
set log module system level notification destination syslog
set log module system level notification destination webtrends
save
```

SNMP

The Simple Network Management Protocol (SNMP) agent for the NetScreen device provides network administrators with a way to view statistical data about the network and the devices on it, and to receive notification of system events of interest.

NetScreen supports the SNMPv1 protocol, described in RFC-1157, “A Simple Network Management Protocol” and the SNMPv2c protocol, described in the following RFCs:

- RFC-1901, “Introduction to Community-based SNMPv2”
- RFC-1905, “Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2)”
- RFC-1906, “Transport Mappings for Version 2 of the Simple Network Management Protocol (SNMPv2)

NetScreen also supports all relevant Management Information Base II (MIB II) groups defined in RFC-1213, “Management Information Base for Network Management of TCP/IP-based internets: MIB-II”. NetScreen also has private enterprise MIB files, which you can load into an SNMP MIB browser. A list of the NetScreen MIBs is included in the appendix. (See Appendix A, “SNMP MIB Files”.)

Accordingly, the NetScreen SNMP agent generates the following traps, or notifications, when specified events or conditions occur:

- **Cold Start Trap:** The NetScreen device generates a cold start trap when it becomes operational after you power it on.
- **Trap for SNMP Authentication Failure:** The SNMP agent in the NetScreen device triggers the authentication failure trap if someone attempts to connect to it using an incorrect SNMP community string or if the IP address of the host attempting the connection is not defined in an SNMP community. (This option is enabled by default.)
- **Traps for System Alarms:** NetScreen device error conditions and firewall conditions trigger system alarms. Three NetScreen enterprise traps are defined to cover alarms related to hardware, security, and software. (For more information on firewall settings and alarms, see “ICMP Fragments” on page **4**-2, and “Traffic Alarms” on page 82.)
- **Traps for Traffic Alarms:** Traffic alarms are triggered when traffic exceeds the alarm thresholds set in policies. (For more information on configuring policies, see “Policies” on page **2**-197.)

The following table lists possible alarm types and their associated trap number:

Trap Enterprise ID	Description
100	Hardware problems
200	Firewall problems
300	Software problems
400	Traffic problems
500	VPN problems
600	NSRP problems
700	Global Pro problems
800	DRP problems
900	Interface failover problems

Note: *The network administrator must have an SNMP manager application such as HP OpenView® or SunNet Manager™ to browse the SNMP MIB II data and to receive traps from either the trusted or untrusted interface. There are also several shareware and freeware SNMP manager applications available from the Internet.*

NetScreen devices do not ship with a default configuration for the SNMP manager. To configure your NetScreen device for SNMP, you must first create communities, define their associated hosts, and assign permissions (read/write or read-only).

When you create an SNMP community, you can specify whether the community supports SNMPv1, SNMPv2c, or both SNMP versions, as required by the SNMP management stations. (For backward compatibility with earlier ScreenOS releases that only support SNMPv1, NetScreen devices support SNMPv1 by default.) If an SNMP community supports both SNMP versions, you must specify a trap version for each community member.

For security reasons, an SNMP community member with read/write privileges can change only the following variables on a NetScreen device:

- **sysContact** - Contact information for the admin of the NetScreen device in case the SNMP admin needs to contact him or her. This can be the NetScreen admin's name, e-mail address, telephone number, location in an office, or a combination of such information.
- **sysLocation** - The physical location of the NetScreen device. This can be anything from the name of a country, city, or building, to its exact location on a rack in a network operation center (NOC).
- **sysName** - The name that SNMP administrators use for the NetScreen device. By convention, this is a fully-qualified domain name (FQDN), but it can also be any name that is meaningful to the SNMP admins.
- **snmpEnableAuthenTraps** - This enables or disables the SNMP agent in the NetScreen device to generate a trap whenever someone attempts to contact the SNMP agent with an incorrect SNMP community name.
- **ipDefaultTTL** - The default value inserted into the time-to-live (TTL) field in the IP header of datagrams originating from the NetScreen device whenever the transport layer protocol does not supply a TTL value.
- **ipForwarding** - This indicates whether or not the NetScreen device forwards traffic—other than that destined for the NetScreen device itself. By default, the NetScreen device indicates that it does not forward traffic (a deceit to disguise its true nature).

Implementation Overview

The following points summarize how NetScreen has implemented SNMP in its devices:

- SNMP administrators are grouped in SNMP communities. A NetScreen device can support up to three communities, with up to eight members in each community.
- A community member can be either a single host or a subnet of hosts, depending on the netmask you use when defining the member. By default, the NetScreen device assigns an SNMP community member with a 32-bit netmask (255.255.255.255), which defines it as a single host.
- If you define an SNMP community member as a subnet, any device on that subnet can poll the NetScreen device for SNMP MIB information. However, the NetScreen device cannot send an SNMP trap to a subnet, only to an individual host.
- Each community has either read-only or read-write permission for the MIB II data.
- Each community can support SNMPv1, SNMPv2c, or both. If a community supports both versions of SNMP, you must specify a trap version for each community member.
- You can allow or deny each community from receiving traps.
- You can access the MIB II data and traps through any physical interface.
- Each system alarm (a system event classified with a severity level of critical, alert, or emergency) generates a single NetScreen enterprise SNMP trap to each of the hosts in each community that is set to receive traps.
- The NetScreen device sends Cold Start / Link Up / Link Down traps to all hosts in communities that you set to receive traps.
- If you specify trap-on for a community, you also have the option to allow traffic alarms.
- You can send SNMP messages through a route-based or policy-based VPN tunnel. For more information, see “VPN Tunnels for Self-Initiated Traffic” on page 97.

Example: Defining a Read/Write SNMP Community

In this example, you create an SNMP community, named *MAge11.* You assign it read/write privileges and enable its members to receive MIB II data and traps. It has the following two members 1.1.1.5/32 and 1.1.1.6/32. Each of these members has an SNMP manager application running a different version of SNMP: SNMPv1 and SNMPv2c.

Note: *Because the community name functions as a password, protect its secrecy with caution.*

You provide contact information for the local admin of the NetScreen device in case an SNMP community member needs to contact him—name: al_baker@mage.com. You also provide the location of the NetScreen device—location: 3-15-2. These numbers indicate that the device is on the third floor, in the fifteenth row, and in the second position in that row.

You also enable the SNMP agent to generate traps whenever someone illegally attempts an SNMP connection to the NetScreen device. Authentication failure traps is a global setting that applies to all SNMP communities and is disabled by default.

Finally, you enable SNMP manageability on ethernet1, an interface that you have previously bound to the Trust zone. This is the interface through which the SNMP manager application communicates with the SNMP agent in the NetScreen device.

WebUI

Configuration > Report Settings > SNMP: Enter the following settings, and then click **Apply**:

System Contact: al_baker@mage.com

Location: 3-15-2

Enable Authentication Fail Trap: (select)

Configuration > Report Settings > SNMP > New Community: Enter the following settings, and then click **OK**:

Community Name: MAge11

Permissions:

Write: (select)

Trap: (select)

Including Traffic Alarms: (clear)

Version: ANY (select)

Hosts IP Address/Netmask and Trap Version:
1.1.1.5/32 v1
1.1.1.6/32 v2c

Network > Interfaces > Edit (for ethernet1): Enter the following settings, and then click **OK**:

Service Options:

Management Services: SNMP

CLI

```
set snmp contact al_baker@mage.com
set snmp location 3-15-2
set snmp auth-trap enable
set snmp community MAge11 read-write trap-on version any
set snmp host Mage 1.1.1.5/32 trap v1
set snmp host Mage 1.1.1.6/32 trap v2
set interface ethernet1 manage snmp
save
```

VPN TUNNELS FOR SELF-INITIATED TRAFFIC

You can use virtual private network (VPN) tunnels to secure remote monitoring of a NetScreen device from a fixed IP address. Using a VPN tunnel, you can protect traffic addressed to and initiated from a NetScreen device. Types of traffic initiated from a NetScreen device can include NetScreen-Global PRO reports, event log entries sent to syslog and WebTrends servers, and SNMP MIB traps.

NetScreen supports two types of VPN tunnel configurations:

- **Route-Based VPNs**: The NetScreen device uses route table entries to direct traffic to tunnel interfaces, which are bound to VPN tunnels.

  To send traffic such as event log entries, NetScreen-Global PRO reports, or SNMP traps generated by the NetScreen device through a route-based VPN tunnel, you must manually enter a route to the proper destination. The route must point to the tunnel interface that is bound to the VPN tunnel through which you want the NetScreen device to direct the traffic. No policy is required.

- **Policy-Based VPNs**: The NetScreen device uses the VPN tunnel names specifically referenced in policies to direct traffic through VPN tunnels.

  To send self-initiated traffic through a policy-based VPN tunnel, you must include the source and destination addresses in the policy. For the source address, use the IP address of an interface on the NetScreen device. For the destination address, use the IP address of the storage server or SNMP community member's workstation, if it is located behind a remote NetScreen device. If the remote SNMP community member uses the NetScreen-Remote VPN client to make VPN connections to the local NetScreen device, use an internal IP address defined on the NetScreen-Remote as the destination address.

  Note: *In releases prior to ScreenOS 5.0.0, the source address had to be the default interface bound to the Trust zone, and the destination address had to be in the Untrust zone. In the current release, this restriction has been eliminated.*

If either the remote gateway or the end entity has a dynamically assigned IP address, then the NetScreen device cannot initiate the formation of a VPN tunnel because these addresses cannot be predetermined, and thus you cannot define routes to them. In such cases, the remote host must initiate the VPN connection. After either a policy-based or route-based VPN tunnel is established, both ends of the tunnel can initiate traffic if policies permit it. Also, for a route-based VPN, there must be a route to the end entity through a tunnel interface bound to the VPN

tunnel—either because you manually entered the route or because the local NetScreen device received the route through the exchange of dynamic routing messages after a tunnel was established. (For information about dynamic routing protocols, see Volume 6, “Dynamic Routing”.) You can also use VPN monitoring with the rekey option or IKE heartbeats to ensure that once the tunnel is established, it remains up regardless of VPN activity. (For more information about these options, see “VPN Monitoring” on page **5**-307, and “Monitoring Mechanisms” on page **5**-384.)

For each VPN tunnel configuration type, you can use any of the following types of VPN tunnel:

- **Manual Key**: You manually set the three elements that define a Security Association (SA) at both ends of the tunnel: a Security Parameters Index (SPI), an encryption key, and an authentication key. To change any element in the SA, you must manually enter it at both ends of the tunnel.
- **AutoKey IKE with Preshared Key**: One or two preshared secrets—one for authentication and one for encryption—function as seed values. Using them, the IKE protocol generates a set of symmetrical keys at both ends of the tunnel; that is, the same key is used to encrypt and decrypt. At predetermined intervals, these keys are automatically regenerated.
- **AutoKey IKE with Certificates**: Using the Public Key Infrastructure (PKI), the participants at both ends of the tunnel use a digital certificate (for authentication) and an RSA public/private key pair (for encryption). The encryption is asymmetrical; that is, one key in a pair is used to encrypt and the other to decrypt.

> ***Note:*** *For a complete description of VPN tunnels, see Volume 5, “VPNs”. For more information on NetScreen-Remote, refer to the* NetScreen-Remote User’s Guide.

Example: Self-Generated Traffic through a Route-Based Tunnel

In this example, you configure a local NetScreen device (NetScreen-A) to send SNMPv1 MIB traps and syslog reports through a route-based AutoKey IKE VPN tunnel to an SNMP community member behind a remote NetScreen device (NetScreen-B). The tunnel uses a preshared key (Ci5y0a1aAG) for data origin authentication and the security level predefined as “Compatible” for both Phase 1 and Phase 2 proposals. You, as the local admin for NetScreen-A, create the tunnel.1 interface and bind it to vpn1. You and the admin for NetScreen-B define the following proxy IDs:

NetScreen-A		NetScreen-B	
Local IP	10.1.1.1/32	Local IP	10.2.2.2/32
Remote IP	10.2.2.2/32	Remote IP	10.1.1.1/32
Service	Any	Service	Any

You bind ethernet1 to the Trust zone, and ethernet3 to the Untrust zone. The default gateway IP address is 1.1.1.250. All zones are in the trust-vr routing domain.

Note: *Compare this example with “Example: Self-Generated Traffic through a Policy-Based Tunnel” on page 109.*

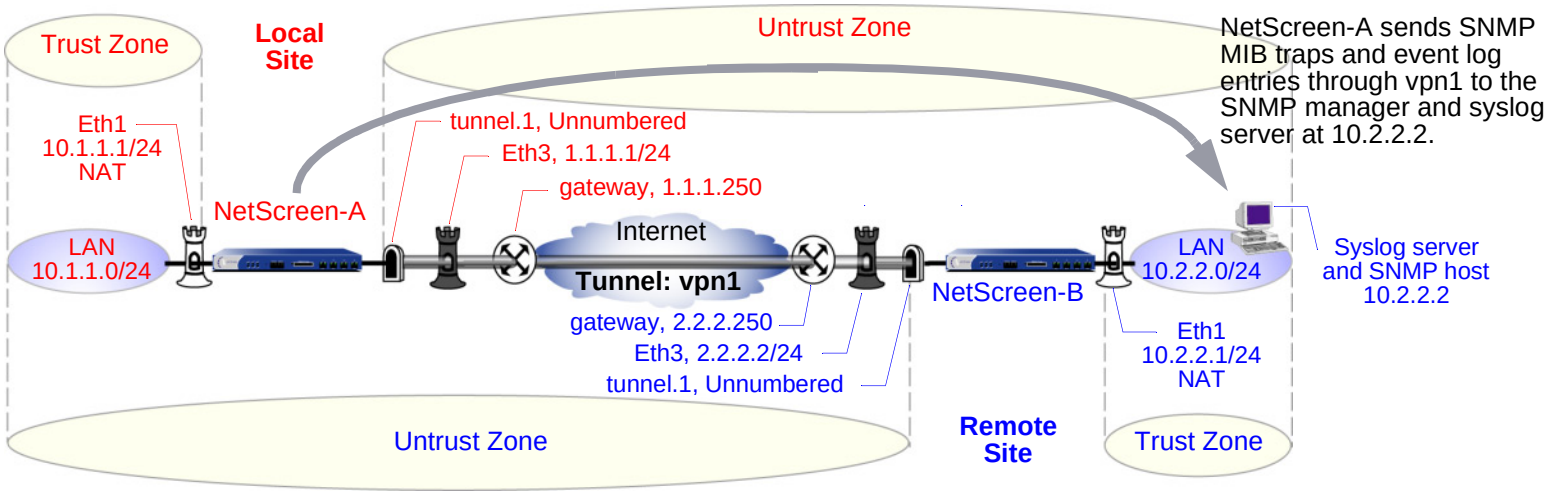

The remote admin for NetScreen-B uses similar settings to define that end of the AutoKey IKE VPN tunnel so that the preshared key, proposals, and proxy IDs match.

You also configure an SNMP community named “remote_admin” with read/write privileges, and you enable the community to receive traps. You define the host at 10.2.2.2/32 as a community member[3].

WebUI (NetScreen-A)

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24[4]

Select the following, and then click **OK**:

Interface Mode: NAT (select)[5]

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

Service Options:

Management Services: SNMP

3. This example assumes that the remote admin has already set up the syslog server and SNMP manager application that supports SNMPv1. When the remote admin sets up the VPN tunnel on his NetScreen device, he uses 1.1.1.1 as the remote gateway and 10.1.1.1 as the destination address.
4. When the remote admin configures the SNMP manager, he must enter **10.1.1.1** in the Remote SNMP Agent field. This is the address to which the SNMP manager sends queries.
5. By default, any interface that you bind to the Trust zone is in NAT mode. Consequently, this option is already enabled for interfaces bound to the Trust zone.

Network > Interfaces > New Tunnel IF: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.1

Zone (VR): Untrust (trust-vr)

Unnumbered: (select)

Interface: ethernet1(trust-vr)

2. Syslog and SNMP

Configuration > Report Settings > Syslog: Enter the following, and then click **Apply**:

Enable Syslog Messages: (select)

No.: Select 1 to indicate you are adding 1 syslog server.

IP/ Hostname: 10.2.2.2

Port: 514

Security Facility: auth/sec

Facility: Local0

Configuration > Report Settings > SNMP > New Community: Enter the following, and then click **OK**:

Community Name: remote_admin

Permissions:

Write: (select)

Trap: (select)

Including Traffic Alarms: (clear)

Version: V1

Hosts IP Address/Netmask:

10.2.2.2/32 V1

Trap Version:

V1

3. VPN

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: vpn1

Security Level: Compatible

Remote Gateway: Create a Simple Gateway: (select)

Gateway Name: to_admin

Type: Static IP, Address/Hostname: 2.2.2.2

Preshared Key: Ci5y0a1aAG

Security Level: Compatible

Outgoing interface ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Bind to: Tunnel Interface: (select), tunnel.1

Proxy-ID: (select)

Local IP/Netmask: 10.1.1.1/32

Remote IP/Netmask: 10.2.2.2/32

Service: ANY

4. Routes

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask:10.2.2.2/32

Gateway: (select)

Interface: tunnel.1

Gateway IP Address: 0.0.0.0

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: (select) 1.1.1.250

CLI (NetScreen-A)

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24[6]
set interface ethernet1 nat[7]
set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
set interface ethernet3 manage snmp
set interface tunnel.1 zone untrust
set interface tunnel.1 ip unnumbered interface ethernet1
```

2. VPN

```
set ike gateway to_admin address 2.2.2.2 outgoing-interface ethernet3 preshare
    Ci5y0a1aAG sec-level compatible
set vpn vpn1 gateway to_admin sec-level compatible
set vpn vpn1 bind interface tunnel.1
set vpn vpn1 proxy-id local-ip 10.1.1.1/32 remote-ip 10.2.2.2/32 any
```

6. When the remote admin configures the SNMP manager, he must enter **10.1.1.1** in the Remote SNMP Agent field. This is the address to which the SNMP manager sends queries.

7. By default, any interface that you bind to the Trust zone is in NAT mode. Consequently, this option is already enabled for interfaces bound to the Trust zone.

3. Syslog and SNMP

```
set syslog config 10.2.2.2 auth/sec local0
set syslog enable
set snmp community remote_admin read-write trap-on version v1
set snmp host remote_admin 10.2.2.2/32
```

4. Routes

```
set vrouter trust-vr route 10.2.2.2/32 interface tunnel.1
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
save
```

WebUI (NetScreen-B)

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.2.2.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 2.2.2.2/24

Network > Interfaces > New Tunnel IF: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.1

Zone (VR): Untrust (trust-vr)

Unnumbered: (select)

Interface: ethernet1(trust-vr)

2. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: addr1

IP Address/Domain Name: IP/Netmask: 10.2.2.2/32

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: ns-a

IP Address/Domain Name: IP/Netmask: 10.1.1.1/32

Zone: Untrust

3. Service Group

Objects > Services > Groups > New: Enter the following group name, move the following services, and then click **OK**:

Group Name: s-grp1

Select **Syslog** and use the **<<** button to move the service from the Available Members column to the Group Members column.

Select **SNMP** and use the **<<** button to move the service from the Available Members column to the Group Members column.

4. VPN

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: vpn1

Security Level: Compatible

Remote Gateway: Create a Simple Gateway: (select)

Gateway Name: to_admin

Type: Static IP, Address/Hostname: 1.1.1.1

Preshared Key: Ci5y0a1aAG

Security Level: Compatible

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Bind to: Tunnel Interface: (select), tunnel.1

Proxy-ID: (select)

Local IP / Netmask: 10.2.2.2/32

Remote IP / Netmask: 10.1.1.1/32

Service: Any

5. Routes

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask:10.1.1.1/32

Gateway: (select)

Interface: tunnel.1

Gateway IP Address: 0.0.0.0

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: (select) 2.2.2.250

6. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), addr1

Destination Address:

Address Book Entry: (select), ns-a

Service: s-grp1

Action: Permit

Position at Top: (select)

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), ns-a

Destination Address:

Address Book Entry: (select), addr1

Service: s-grp1

Action: Permit

Position at Top: (select)

CLI (NetScreen-B)

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.2.2.1/24
set interface ethernet1 nat
set interface ethernet3 zone untrust
set interface ethernet3 ip 2.2.2.2/24
set interface tunnel.1 zone untrust
set interface tunnel.1 ip unnumbered interface ethernet1
```

2. Addresses

```
set address trust addr1 10.2.2.2/32
set address untrust ns-a 10.1.1.1/32
```

3. Service Group

```
set group service s-grp1
set group service s-grp1 add syslog
set group service s-grp1 add snmp
```

4. VPN

```
set ike gateway to_admin address 1.1.1.1 outgoing-interface ethernet3 preshare
    Ci5y0a1aAG sec-level compatible
set vpn vpn1 gateway to_admin sec-level compatible
set vpn vpn1 bind interface tunnel.1
set vpn vpn1 proxy-id local-ip 10.2.2.2/32 remote-ip 10.1.1.1/32 any
```

5. Routes

```
set vrouter trust-vr route 10.1.1.1/32 interface tunnel.1
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 2.2.2.250
```

6. Policies

```
set policy top from trust to untrust addr1 ns-a s-grp1 permit
set policy top from untrust to trust ns-a addr1 s-grp1 permit
save
```

Example: Self-Generated Traffic through a Policy-Based Tunnel

In this example, you configure a local NetScreen device (NetScreen-A) to send SNMPv2c MIB traps and syslog reports[8] through a policy-based AutoKey IKE VPN tunnel (vpn1) to an SNMP community member behind a remote NetScreen device (NetScreen-B). The tunnel uses a preshared key (Ci5y0a1aAG) for data origin authentication and the security level predefined as “Compatible” for both Phase 1 and Phase 2 proposals.

Both you and the remote admin bind ethernet1 to the Trust zone, and ethernet3 to the Untrust zone on NetScreen-A and NetScreen-B. The default gateway IP address for NetScreen-A is 1.1.1.250. The default gateway IP address for NetScreen-B is 2.2.2.250. All zones are in the trust-vr routing domain.

Note: *Compare this example with “Example: Self-Generated Traffic through a Route-Based Tunnel” on page 99.*

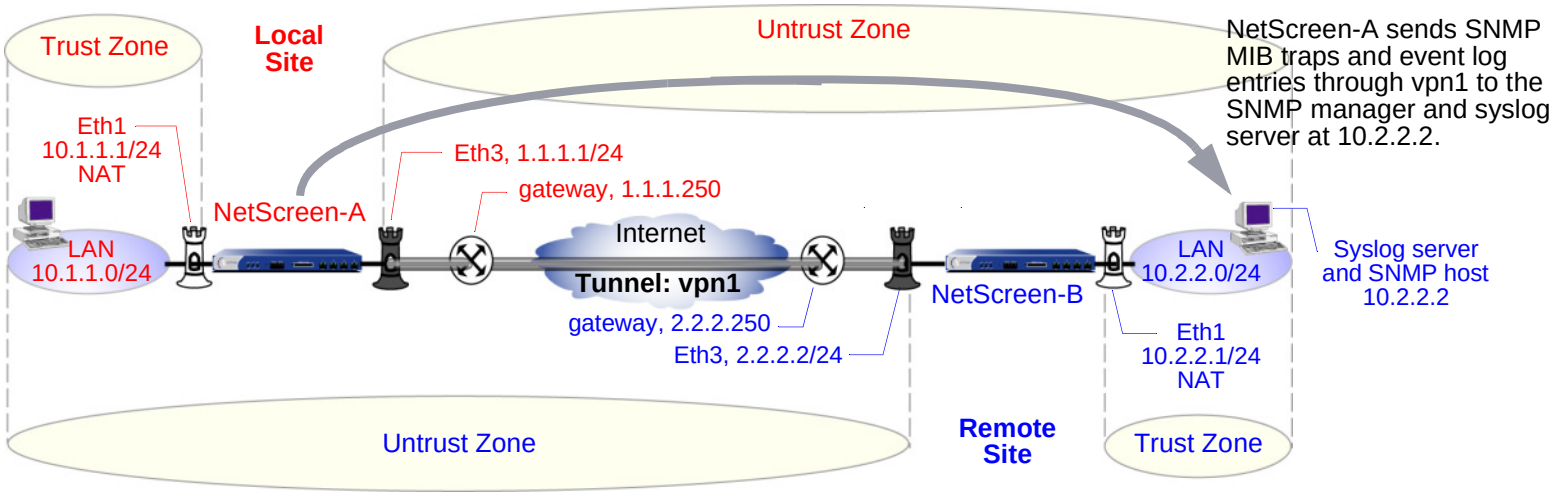

You also configure an SNMP community named “remote_admin” with read/write privileges, and you enable the community to receive traps. You define the host at 10.2.2.2/32 as a community member.

8. This example assumes that the remote admin has already set up the syslog server and an SNMP manager application that supports SNMPv2c. When the remote admin sets up the VPN tunnel on his NetScreen device, he uses 1.1.1.1 as the remote gateway and 10.1.1.1 as the destination address.

The inbound and outbound policies on NetScreen-A match the outbound and inbound policies on NetScreen-B. The addresses and service used in the policies are as follows:

- 10.1.1.1/32, the address of the Trust zone interface on NetScreen-A
- 10.2.2.2/32, the address of the host for the SNMP community member and syslog server
- Service group named “s-grp1”, which contains SNMP and syslog services

From the policies that you and the admin for NetScreen-B create, the two NetScreen devices derive the following proxy IDs for vpn1:

NetScreen-A			NetScreen-B	
Local IP	10.1.1.1/32		Local IP	10.2.2.2/32
Remote IP	10.2.2.2/32		Remote IP	10.1.1.1/32
Service	Any		Service	Any

Note: *NetScreen treats a service group as “any” in proxy IDs.*

WebUI (NetScreen-A)

1. Interfaces – Security Zones

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **OK**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24[9]

Select the following, and then click **OK**:

Interface Mode: NAT (select)[10]

9. When the remote admin configures the SNMP manager, he must enter **10.1.1.1** in the Remote SNMP Agent field. This is the address to which the SNMP manager sends queries.

10. By default, any interface that you bind to the Trust zone is in NAT mode. Consequently, this option is already enabled for interfaces bound to the Trust zone.

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

Service Options:

Management Services: SNMP

2. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: trust_int

IP Address/Domain Name: IP/Netmask: 10.1.1.1/32

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: remote_admin

IP Address/Domain Name: IP/Netmask: 10.2.2.2/32

Zone: Untrust

3. Service Group

Objects > Services > Groups > New: Enter the following group name, move the following services, and then click **OK**:

Group Name: s-grp1

Select **Syslog** and use the **<<** button to move the service from the Available Members column to the Group Members column.

Select **SNMP** and use the **<<** button to move the service from the Available Members column to the Group Members column.

4. VPN

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: vpn1

Security Level: Compatible

Remote Gateway: Create a Simple Gateway: (select)

Gateway Name: to_admin

Type: Static IP, Address/Hostname: 2.2.2.2

Preshared Key: Ci5y0a1aAG

Security Level: Compatible

Outgoing Interface: ethernet3

5. Syslog and SNMP

Configuration > Report Settings > Syslog: Enter the following, and then click **Apply**:

Enable Syslog Messages: (select)

Source Interface: ethernet1

No.: Select 1 to indicate you are adding 1 syslog server.

IP/ Hostname: 10.2.2.2

Port: 514

Security Facility: auth/sec

Facility: Local0

Configuration > Report Settings > SNMP > New Community: Enter the following, and then click **OK**:

Community Name: remote_admin

Permissions:

Write: (select)

Trap: (select)

Including Traffic Alarms: (clear)

Version: V2C

Hosts IP Address/Netmask:

10.2.2.2/32 V2C

Trap Version:

V2C

Source Interface

ethernet1 (select)

Configuration > Report Settings > SNMP: Enter the following, and then click **Apply**:

Enable Authentication Fail Trap: (select)

6. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 1.1.1.250

7. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), trust_int

Destination Address:

Address Book Entry: (select), remote_admin

Service: s-grp1

Action: Tunnel

Tunnel VPN: vpn1

Modify matching outgoing VPN policy: (select)

Position at Top: (select)

CLI (NetScreen-A)

1. Interfaces – Security Zones

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat[11]
set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
set interface ethernet3 manage snmp
```

2. Addresses

```
set address trust trust_int 10.1.1.1/32
set address untrust remote_admin 10.2.2.2/32
```

3. Service Group

```
set group service s-grp1
set group service s-grp1 add syslog
set group service s-grp1 add snmp
```

11. By default, any interface that you bind to the Trust zone is in NAT mode. Consequently, this option is already enabled for interfaces bound to the Trust zone.

4. VPN

```
set ike gateway to_admin address 2.2.2.2 outgoing-interface ethernet3 preshare
    Ci5y0a1aAG sec-level compatible
set vpn vpn1 gateway to_admin sec-level compatible
```

5. Syslog and SNMP

```
set syslog config 10.2.2.2 auth/sec local0
set syslog src-interface ethernet1
set syslog enable
set snmp community remote_admin read-write trap-on version v2c
set snmp host remote_admin 10.2.2.2/32 src-interface ethernet1
```

6. Route

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
```

7. Policies

```
set policy top from trust to untrust trust_int remote_admin s-grp1 tunnel vpn
    vpn1
set policy top from untrust to trust remote_admin trust_int s-grp1 tunnel vpn
    vpn1
save
```

WebUI (NetScreen-B)

1. Interfaces – Security Zones

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.2.2.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 2.2.2.2/24

2. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: addr1

IP Address/Domain Name:

IP/Netmask: 10.2.2.2/32

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: ns-a

IP Address/Domain Name:

IP/Netmask: 10.1.1.1/32

Zone: Untrust

3. Service Group

Objects > Services > Group: Enter the following group name, move the following services, and then click **OK**:

Group Name: s-grp1

Select **Syslog** and use the **<<** button to move the service from the Available Members column to the Group Members column.

Select **SNMP** and use the **<<** button to move the service from the Available Members column to the Group Members column.

4. VPN

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: vpn1

Security Level: Compatible

Remote Gateway: Create a Simple Gateway: (select)

Gateway Name: to_admin

Type: Static IP, IP Address: 1.1.1.1

Preshared Key: Ci5y0a1aAG

Security Level: Compatible

Outgoing interface ethernet3

5. Route

Network > Routing > Routing Table > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: (select) 2.2.2.250

6. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), addr1

Destination Address:

Address Book Entry: (select), ns-a

Service: s-grp1

Action: Tunnel

Tunnel VPN: vpn1

Modify matching outgoing VPN policy: (select)

Position at Top: (select)

CLI (NetScreen-B)

1. Interfaces – Security Zones

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.2.2.1/24
set interface ethernet1 nat
set interface ethernet3 zone untrust
set interface ethernet3 ip 2.2.2.2/24
```

2. Addresses

```
set address trust addr1 10.2.2.2/32
set address untrust ns-a 10.1.1.1/32
```

3. Service Group

```
set group service s-grp1
set group service s-grp1 add syslog
set group service s-grp1 add snmp
```

4. VPN

```
set ike gateway to_admin address 1.1.1.1 outgoing-interface ethernet3 preshare
    Ci5y0a1sec-level compatible
set vpn vpn1 gateway to_admin sec-level compatible
```

5. Route

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 2.2.2.250
```

6. Policies

```
set policy top from trust to untrust addr1 ns-a s-grp1 tunnel vpn vpn1
set policy top from untrust to trust ns-a addr1 s-grp1 tunnel vpn vpn1
save
```

COUNTERS

NetScreen provides screen, hardware, and flow counters for monitoring traffic. Counters give processing information for specified zones and interfaces, and help you to verify configurations for desired policies.

NetScreen provides the following screen counters for monitoring general firewall behavior and for viewing the amount of traffic affected by specified policies:

- **Bad IP Option Protection** – the number of frames discarded due to malformed or incomplete IP options
- **Dst IP-based session limiting** – the number of sessions dropped after the session threshold was reached
- **FIN bit with no ACK bit** – the number of packets detected and dropped with an illegal combination of flags
- **Fragmented packet protection** – the number of blocked IP packet fragments
- **HTTP Component Blocked** – the number of blocked packets with HTTP components
- **HTTP Component Blocking for ActiveX controls** – the number of ActiveX components blocked
- **HTTP Component Blocking for .exe files** – the number of blocked HTTP packets with .exe files
- **HTTP Component Blocking for Java applets** – the number of blocked Java components
- **HTTP Component Blocking for .zip files** – the number of blocked HTTP packets with .zip files
- **ICMP Flood Protection** – the number of ICMP packets blocked as part of an ICMP flood
- **ICMP Fragment** – the number of ICMP frames with the More Fragments flag set, or with offset indicated in the offset field
- **IP Spoofing Attack Protection** – the number of IP addresses blocked as part of an IP spoofing attack
- **IP Sweep Protection** – the number of IP sweep attack packets detected and blocked
- **Land Attack Protection** – the number of packets blocked as part of a suspected land attack
- **Large ICMP Packet** – the number of ICMP frames detected with an IP length greater than 1024
- **limit session** – the number of undeliverable packets because the session limit had been reached
- **Loose Src Route IP Option** – the number of IP packets detected with the Loose Source Route option enabled
- **Malicious URL Protection** – the number of suspected malicious URLs blocked

- **Ping-of-Death Protection** – the number of suspected and rejected ICMP packets that are oversized or of an irregular size
- **Port Scan Protection** – the number of port scans detected and blocked
- **Record Route IP Option** – the number of frames detected with the Record Route option enabled
- **Security IP Option** – the number of frames discarded with the IP Security option set
- **Src IP-based session limiting** – the number of sessions dropped after the session threshold was reached
- **Source Route IP Option Filter** – the number of IP source routes filtered
- **Stream IP Option** – the number of packets discarded with the IP Stream identifier set
- **Strict Src Route IP Option** – the number of packets detected with the Strict Source Route option enabled
- **SYN-ACK-ACK-Proxy DoS** – the number of blocked packets because of the SYN-ACK-ACK-proxy DoS SCREEN option
- **SYN and FIN bits set** – the number of packets detected with an illegal combination of flags
- **SYN Flood Protection** – the number of SYN packets detected as part of a suspected SYN flood
- **SYN Fragment Detection** – the number of packet fragments dropped as part of a suspected SYN fragments attack
- **Timestamp IP Option** – the number of IP packets discarded with the Internet Timestamp option set
- **TCP Packet without Flag** – the number of illegal packets dropped with missing or malformed flags field
- **Teardrop Attack Protection** – the number of packets blocked as part of a Teardrop attack
- **UDP Flood Protection** – the number of UDP packets dropped as part of a suspected UDP flood
- **Unknown Protocol Protection** – the number of packets blocked as part of an unknown protocol
- **WinNuke Attack Protection** – the number of packets detected as part of a suspected WinNuke attack

NetScreen provides the following hardware counters for monitoring hardware performance and packets with errors:

- **drop vlan** – the number of dropped packets because of missing VLAN tags, an undefined subinterface, or because VLAN trunking was not enabled when the NetScreen device was in Transparent mode
- **early frame** – counters used in an ethernet driver buffer descriptor management
- **in align err** – the number of incoming packets with an alignment error in the bit stream
- **in bytes** – the number of bytes received

- **in coll err** – the number of incoming collision packets
- **in crc err** – the number of incoming packets with a cyclic redundancy check (CRC) error
- **in dma err** – the number of incoming packets with a dma error
- **in misc err** – the number of incoming packets with a miscellaneous error
- **in no buffer** – the number of unreceivable packets because of unavailable buffers
- **in overrun** – the number of transmitted overrun packets
- **in packets** – the number of packets received
- **in short frame** – the number of incoming packets with an ethernet frame shorter than 64 bytes (including the frame checksum)
- **in underrun** – the number of transmitted underrun packets
- **late frame** – counters used in an ethernet driver buffer descriptor management
- **out bs pak** – the number of packets held in back store while searching for an unknown MAC address
- **out bytes** – the number of bytes sent
- **out coll err** – the number of outgoing collision packets
- **out cs lost** – the number of dropped outgoing packets because the Carrier Sense Multiple Access/Collision Detect (CSMA/CD) protocol lost the signal[12]
- **out defer** – the number of deferred outgoing packets
- **out discard** – the number of discarded outgoing packets
- **out heartbeat** – the number of outgoing heartbeat packets
- **out misc err** – the number of outgoing packets with a miscellaneous error
- **out no buffer** – the number of unsent packets because of unavailable buffers
- **out packets** – the number of packets sent
- **re xmt limit** – the number of dropped packets when the retransmission limit was exceeded while an interface was operating at half duplex

12. For more information about the Carrier Sense Multiple Access/Collision Detect (CSMA/CD) protocol, see the IEEE 802.3 standard available at http://standards.ieee.org.

NetScreen also provides the following flow counters[13] for monitoring the number of packets inspected at the flow level:

- **address spoof** – the number of suspected address spoofing attack packets received
- **auth deny** – the number of times user authentication was denied.
- **auth fail** – the number of times user authentication failed
- **big bkstr** – the number of packets that are too big to buffer in the ARP backstore while waiting for MAC-to-IP address resolution
- **connections** – the number of sessions established since the last boot
- **encrypt fail** – the number of failed Point-to-Point Tunneling Protocol (PPTP) packets
- ***icmp broadcast** – the number of ICMP broadcasts received
- **icmp flood** – the number of ICMP packets that are counted toward the ICMP flood threshold
- **illegal pak**– the number of packets dropped because they do not conform to the protocol standards
- **in arp req** – the number of incoming arp request packets
- **in arp resp** – the number of outgoing arp request packets
- **in bytes** – the number of bytes received
- **in icmp** – the number of Internet Control Message Protocol (ICMP) packets received
- **in other** – the number of incoming packets that are of a different Ethernet type
- **in packets** – the number of packets received
- **in self** – the number of packets addressed to the NetScreen Management IP address
- ***in un auth** – the number of unauthorized incoming TCP, UDP, and ICMP packets
- ***in unk prot** – the number of incoming packets using an unknown ethernet protocol
- **in vlan** – the number of incoming vlan packets
- **in vpn** – the number of IPSec packets received
- **invalid zone** – the number of packets destined for an invalid security zone
- **ip sweep** – the number of packets received and discarded beyond the specified ip sweep threshold

13. Counters preceded by an asterisk are not yet operational at the time of this writing and always display a value of 0.

- **land attack** – the number of suspected land attack packets received
- **loopback drop** – the number of packets dropped because they cannot be looped back through the NetScreen device. An example of a loopback session is when a host in the Trust zone sends traffic to a MIP or VIP address that is mapped to a server that is also in the Trust zone. The NetScreen device creates a loopback session that directs such traffic from the host to the MIP or VIP server.
- **mac relearn** – the number of times that the MAC address learning table had to relearn the interface associated with a MAC address because the location of the MAC address changed.
- **mac tbl full** – the number of times that the MAC address learning table completely filled up.
- **mal url** – the number of blocked packets destined for a URL determined to be malicious
- ***misc prot** – the number of packets using a protocol other than TCP, UDP, or ICMP
- **mp fail** – the number of times a problem occurred when sending a PCI message between the master processor module and the processor module
- **no conn** – the number of packets dropped because of unavailable Network Address Translation (NAT) connections
- **no dip** – the number of packets dropped because of unavailable Dynamic IP (DIP) addresses
- **no frag netpak** – the number of times that the available space in the netpak buffer fell below 70%
- ***no frag sess** – the number of times that fragmented sessions were greater than half of the maximum number of NAT sessions
- **no g-parent** – the number of packets dropped because the parent connection could not be found
- **no gate** – the number of packets dropped because no gate was available
- **no gate sess** – the number of terminated sessions because there were no gates in the firewall for them
- **no map** – the number of packets dropped because there was no map to the trusted side
- **no nat vector** – the number of packets dropped because the Network Address Translation (NAT) connection was unavailable for the gate
- ***no nsp tunnel** – the number of dropped packets sent to a tunnel interface to which no VPN tunnel is bound
- **no route** – the number of unroutable packets received
- **no sa** – the number of packets dropped because no Security Associations (SA) was defined
- **no sa policy** – the number of packets dropped because no policy was associated with an SA

- ***no xmit vpnf** – the number of dropped VPN packets due to fragmentation
- **null zone** – the number of dropped packets erroneously sent to an interface bound to the Null zone
- **nvec err** – the number of packets dropped because of NAT vector error
- **out bytes** – the number of bytes sent
- **out packets** – the number of packets sent
- **out vlan** – the number of outgoing vlan packets
- **ping of death** – the number of suspected Ping of Death attack packets received
- **policy deny** – the number of packets denied by a defined policy
- **port scan** – the number of packets that are counted as a port scan attempt
- **proc sess** – the number of times that the total number of sessions on a processor module exceeded the maximum threshold
- **sa inactive** – the number of packets dropped because of an inactive SA
- **sa policy deny** – the number of packets denied by an SA policy
- **sessn thresh** – the threshold for the maximum number of sessions
- ***slow mac** – the number of frames whose MAC addresses were slow to resolve
- **src route** – the number of packets dropped because of the filter source route option
- **syn frag** – the number of dropped SYN packets because of a fragmentation
- **tcp out of seq** – the number of TCP segments received whose sequence number is outside the acceptable range
- **tcp proxy** – the number of packets dropped from using a TCP proxy such as the SYN flood protection option or user authentication
- **tear drop** – the number of packets blocked as part of a suspected Tear Drop attack
- **tiny frag** – the number of tiny fragmented packets received
- **trmn drop** – the number of packets dropped by traffic management
- **trmng queue** – the number of packets waiting in the queue
- **udp flood** – the number of UDP packets that are counted toward the UDP flood threshold
- **url block** – the number of HTTP requests that were blocked

- **winnuke** – the number of WinNuke attack packets received
- **wrong intf** – the number of session creation messages sent from a processor module to the master processor module
- **wrong slot** – the number of packets erroneously sent to the wrong processor module

Example: Viewing Screen Counters

In this example, you view the NetScreen screen counters for the Trust zone.

WebUI

Reports > Counters > Zone Screen: Select **Trust** from the Zone drop-down list.

CLI

```
get counter screen zone trust
```

A

SNMP MIB Files

NetScreen provides MIB files to support SNMP communication between your organization's applications and the SNMP agent in the NetScreen device. To obtain the latest MIB files, open a Web browser and visit www.netscreen.com/services/tac_online/index.jsp. Select a NetScreen product, and then select the MIB files for the ScreenOS version loaded on the NetScreen device.

The MIB files for the current ScreenOS version are fully compatible with SNMP agents in previous versions of ScreenOS. The NetScreen MIB files are organized in a multi-tier hierarchical structure and are described as follows:

- "The Primary-Level MIB File Folders" on page II
- "Secondary-Level MIB Folders" on page IV
  - "netscreenProducts" on page IV
  - "netScreenIds" on page V
  - "netscreenVpn" on page V
  - "netscreenQos" on page V
  - "netscreenSetting" on page VI
  - "netscreenZone" on page VI
  - "netscreenPolicy" on page VI
  - "netscreenNAT" on page VII
  - "netscreenAddr" on page VII
  - "netscreenService" on page VII
  - "netscreenSchedule" on page VII
  - "netscreenVsys" on page VII
  - "netscreenResource" on page VIII
  - "netscreenIp" on page VIII
  - "netscreenVR" on page VIII

The Primary-Level MIB File Folders

The MIB files are arranged in a hierarchical folder structure. The primary-level MIB folders are as follows:

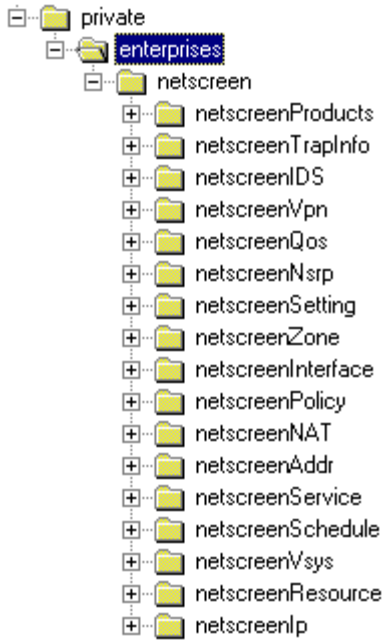

Each folder contains a category of MIB files.

netscreenProducts	Assigns Object Identifiers (OIDs) to different NetScreen product series.
netscreenTrapInfo	Defines enterprise traps sent by the NetScreen device.
netscreenIDS	Defines the NetScreen device intrusion detection service (IDS) configuration.
netscreenVpn	Defines NetScreen device VPN configuration and runtime information.
netscreenQos	Defines NetScreen device Quality of Service configuration.

netscreenNsrp	Defines NetScreen device NSRP configuration.
netscreenSetting	Defines miscellaneous NetScreen device configuration settings, such as DHCP, e-mail, authentication, and administrator.
netscreenZone	Defines zone information residing in the NetScreen Device.
netscreenInterface	Defines the NetScreen device's interface configuration, including the virtual interface.
netscreenPolicy	Defines the outgoing and incoming policy configuration for the NetScreen device.
netscreenNAT	Defines NAT configuration, including Map IP, Dynamic IP and Virtual IP.
netscreenAddr	Represents the address table on a NetScreen interface.
netscreenService	Describes services (including user-defined) recognized by the NetScreen device.
netscreenSchedule	Defines NetScreen device task schedule information, configured by the user.
netscreenVsys	Defines NetScreen device virtual system (VSYS) configuration.
netscreenResource	Accesses information regarding the NetScreen device's resource utilization.
netscreenIp	Accesses NetScreen device private IP-related information.
netScreen Chassis	Empty placeholder folder for future MIB support folders
netscreenVR	Defines NetScreen device virtual router (VR) configuration.

Secondary-Level MIB Folders

This section describes the secondary-level MIB files for NetScreen devices. Each secondary-level folder contains subsequent-level folders or MIB files.

netscreenProducts

netscreenGeneric	Generic object identifiers (OIDs) for NetScreen products
netscreenNs5	NetScreen-5XP OIDs
netscreenNs10	NetScreen-10XP OIDs
netscreenNs100	NetScreen-100 OIDs
netscreenNs1000	NetScreen-1000 OIDs
netscreenNs500	NetScreen-500 OIDs
netscreenNs50	NetScreen-50 OIDs
netscreenNs25	NetScreen-25 OIDs
netscreenNs204	NetScreen-204 OIDs
netscreenNs208	NetScreen-208 OIDs

netScreenIds

nsIdsProtect		IDS service on NetScreen device
	nsIdsProtectSetTable	IDS service enabled on NetScreen device
	nsIdsProtectThreshTable	IDS service threshold configuration
nsIdsAttkMonTable		Statistical Information about intrusion attempt

netscreenVpn

netscreenVpnMon	Show SA information of vpn tunnel
nsVpnManualKey	Manual key configuration
nsVpnIke	IKE configuration
nsVpnGateway	VPN tunnel gateway configuration
nsVpnPhaseOneCfg	IPSec Phase One configuration
nsVpnPhaseTwoCfg	IPSec Phase Two configuration
nsVpnCert	Certification configuration
nsVpnL2TP	L2TP configuration
nsVpnPool	IP pool configuration
nsVpnUser	VPN user configuration

netscreenQos

nsQosPly	QoS configuration on policy

netscreenSetting

nsSetGeneral	General configuration of NS device
nsSetAuth	Authentication method configuration
nsSetDNS	DNS server setting
nsSetURLFilter	URL filter setting
nsSetDHCP	DHCP server setting
nsSetSysTime	System time setting
nsSetEmail	E-mail setting
nsSetLog	Syslog setting
nsSetSNMP	SNMP agent configuration
nsSetGlbMng	Global management configuration
nsSetAdminUser	Administration user configuration
nsSetWebUI	Web UI configuration

netscreenZone

nsZoneCfg	Zone configuration for the device

netscreenPolicy

NsPlyTable	Policy configuration
NsPlyMonTable	Statistical Information about each policy

netscreenNAT

nsNatMipTable	Mapped IP configuration
nsNatDipTable	Dynamic IP configuration
nsNatVip	Virtual IP Configuration

netscreenAddr

nsAddrTable	Address table on a NetScreen interface

netscreenService

nsServiceTable	Service Information
nsServiceGroupTable	Service Group Information
nsServiceGrpMemberTable	Service Group Member Info

netscreenSchedule

nschOnceTable	One-time schedule information
nschRecurTable	Re-occur schedule information

netscreenVsys

nsVsysCfg	NetScreen device virtual system (VSYS) configuration

netscreenResource

nsresCPU	CPU utilization
nsresMem	Memory utilization
nsresSession	Session utilization

Note: *NetScreen no longer supports the failedSession counter.*

netscreenIp

nsIpArp	ARP table

netscreenVR[1]

nsOSPF	Open Shortest Path First (OSPF) protocol information
nsBGP	Border Gateway Protocol (BGP4) protocol information
nsRIP	Routing Information Protocol (RIP) protocol information

1. The netscreenVR MIBs are based on Structure of Management Information version 2 (SMIv2). All the other MIBs are based on SMIv1. You can access all the MIB II data, regardless of whether you are running SNMPv1 or SNMPv2c.

Index

A

administration
  CLI (Command Line Interface) 9
  restricting 49, 50
  WebUI 3
administrative traffic 30
alarms
  E-mail alert 82
  reporting to NSM 26
  thresholds 82
  traffic 82–86
asset recovery log 81
AutoKey IKE VPN 51, 98

B

back store 122
bit stream 121
browser requirements 3

C

cables, serial 21
character types, ScreenOS supported x
CLI 9, 30, 31
  conventions vi
command line interface
  See CLI
CompactFlash 66
configuration settings
  browser requirements 3
console 66
conventions
  CLI vi
  illustration ix
  names x
  WebUI vii
creating
  keys 7

D

DIP 124
Dynamic IP
  See DIP

E

e-mail alert notification 86, 89, 90
event log 67

F

filter source route 125

H

HTTP 5
  session ID 5
Hypertext Transfer Protocol
  See HTTP

I

Ident-Reset 29
illustration
  conventions ix
inactive SA 125
in-short error 122
interfaces
  manageable 34
  management options 29
internal flash storage 66
IP addresses
  manage IP 34
  NSM servers 26

K

keys
  creating 7

L

logging 66–81
  asset recovery log 81
  CompactFlash (PCMCIA) 66
  console 66
  e-mail 66
  event log 67
  internal 66
  NSM reporting 26
  self log 77
  SNMP 66, 91
  syslog 66, 87
  WebTrends 66, 89
logging in
  root admin 50
  Telnet 10

M

manage IP 34
management client IP addresses 49
Management information base II
  See MIB II
management methods
  CLI 9
  console 21
  SSL 7
  Telnet 9
  WebUI 3
management options 29
  manageable 34
  NSM 29
  ping 29
  SCS 29
  SNMP 29
  SSL 29
  Telnet 29
  Transparent mode 30
  WebUI 29
Manual Key
  VPNs 51, 98

messages
alert 67
critical 67
debug 67
emergency 67
error 67
info 67
notice 67
warning 67
WebTrends 90
MGT interface
management options 30
MIB files A-I
MIB folders
primary A-II
MIB II 29, 91
modem port 22

N

names
conventions x
NAT vector error 125
NetScreen Security Manager
See NSM
Network Address Translation (NAT) 124
NSM
Agent 23, 26
definition 23
enabling the Agent 25
initial connectivity setup 24
management options 29
Management System 23, 26
reporting events 26, 27
UI 23

O

operating system 9

P

packets 125
address spoofing attack 123
collision 122
denied 125
dropped 124, 125
fragmented 125
incoming 122
Internet Control Message Protocol (ICMP) 120, 123
IPSec 123
land attack 124
Network Address Translation (NAT) 124
Point to Point Tunneling Protocol (PPTP) 123
received 121, 122, 123, 125
transmitted underrun 122
unreceivable 122
unroutable 124
parent connection 124
password
forgetting 44
root admin 47
PCMCIA 66
ping
management options 29
PKI
key 7
Point-to-Point Tunneling Protocol (PPTP) 123
ports
modem 22
protocol distribution
reporting to NSM 26

R

RADIUS 44
reset to factory defaults 48

S

SA policy 125
SCS 29
Secure Sockets Layer
See SSL
Security Associations (SA) 124
self log 77
serial cables 21
session ID 5
SMTP server IP 86
SNMP 29, 91
cold start trap 91
community, private 95
community, public 95
configuration 95
encryption 94, 97
implementation 94
management options 29
MIB files A-I
MIB folders, primary A-II
system alarm traps 91
traffic alarm traps 91
trap types 92
traps 91
SNMP traps
100, hardware problems 92
200, firewall problems 92
300, software problems 92
400, traffic problems 92
500, VPN problems 92
allow or deny 94
source route 125
SSH 11–17
authentication method priority 17
automated logins 19
connection procedure 12
forcing PKA authentication only 17
host key 12
loading public keys, CLI 16
loading public keys, TFTP 16, 19
loading public keys, WebUI 16
password authentication 15
PKA 15
PKA authentication 15
PKA key 12
server key 12
session key 12
SSL 7
management options 29
SSL Handshake Protocol
See SSLHP
SSLHP 7
statistics
reporting to NSM 27

syslog 66
encryption 97
facility 88, 90, 101, 112
host 87
host name 88, 89, 90, 101, 112
messages 87
port 88, 101, 112
security facility 88, 90, 101, 112

T

TCP
proxy 125
Telnet 9, 29
traffic
alarms 82–86
Transparent mode
management options 30

U

users
multiple administrative users 37

V

virtual private network
See VPNs
virtual system
administrators 38
read-only admins 38
VLAN1
management options 30
VPNs
AutoKey IKE 51, 98
for administrative traffic 97
Manual Key 51, 98

W

Web browser requirements 3
Web user interface
See WebUI
WebTrends 66, 89
encryption 89, 97
messages 90
WebUI 3, 30, 31
conventions vii

NetScreen Concepts & Examples

ScreenOS Reference Guide

Volume 4: Attack Detection and Defense Mechanisms

ScreenOS 5.0.0

P/N 093-0927-000

Rev. B

Copyright Notice

NetScreen Technologies, Inc.
Building #3
805 11th Avenue
Sunnyvale, CA 94089
www.netscreen.com

FCC Statement

The following information is for FCC compliance of Class A devices: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. The equipment generates, uses, and can radiate radio-frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case users will be required to correct the interference at their own expense.

The following information is for FCC compliance of Class B devices: The equipment described in this manual generates and may radiate radio-frequency energy. If it is not installed in accordance with NetScreen's installation instructions, it may cause interference with radio and television reception. This equipment has been tested and found to comply with the limits for a Class B digital device in accordance with the specifications in part 15 of the FCC rules. These specifications are designed to provide reasonable protection against such interference in a residential installation. However, there is no guarantee that interference will not occur in a particular installation.

If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Consult the dealer or an experienced radio/TV technician for help.
- Connect the equipment to an outlet on a circuit different from that to which the receiver is connected.

Caution: Changes or modifications to this product could void the user's warranty and authority to operate this device.

Disclaimer

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR NETSCREEN REPRESENTATIVE FOR A COPY.

Contents

Preface v
Conventions vi
CLI Conventions vi
WebUI Conventions vii
Illustration Conventions ix
Naming Conventions and Character Types x
NetScreen Documentation xi
Chapter 1 Protecting a Network 1
Stages of an Attack 2
Detection and Defense Mechanisms 3
Exploit Monitoring 5
Example: Monitoring Attacks from the Untrust Zone 6
Chapter 2 Reconnaissance Deterrence 7
IP Address Sweep 8
Port Scanning 10
Network Reconnaissance Using IP Options 12
Operating System Probes 16
SYN and FIN Flags Set 16
FIN Flag without ACK Flag 18
TCP Header without Flags Set 20
Evasion Techniques 22
FIN Scan 22
IP Spoofing 22
Example: L3 IP Spoof Protection 25
Example: L2 IP Spoof Protection 29
IP Source Route Options 31
Chapter 3 Denial-of-Service Attack Defenses 35
Firewall DoS Attacks 36
Session Table Flood 36
Source- and Destination-Based Session Limits 36
Example: Source-Based Session Limiting 39
Example: Destination-Based Session Limiting 40
Aggressive Aging 40
Example: Aggressively Aging Out Sessions 42
SYN-ACK-ACK Proxy Flood 43
Network DoS Attacks 45
SYN Flood 45
Example: SYN Flood Protection 52
ICMP Flood 59
UDP Flood 61
Land Attack 63
OS-Specific DoS Attacks 65
Ping of Death 65
Teardrop Attack 67
WinNuke 69

Chapter 4 Content Monitoring and Filtering 71
Fragment Reassembly 72
Malicious URL Protection 72
Application Layer Gateway 73
Example: Blocking Malicious URLs in Packet Fragments 74
Antivirus Scanning 76
Internal AV Scanning 77
Enabling Internal AV Scanning 81
Updating the Pattern File Automatically or Semi-Automatically 82
Example: Automatic Pattern Update 83
Example: Semi-Automatic Pattern Update 83
Configuring Content Processing 84
Example: Internal AV Scanning for SMTP 84
Example: Internal AV Scanning for SMTP and HTTP 85
Configuring Decompression and Maximum Content Size 85
Example: Dropping Large Files 86
Applying Internal AV Scanning 87
Example: Internal AV Scanning (POP3) 87
External AV Scanning 90
Defining AV Objects 93
Example: Defining Three AV Objects 99
Applying External AV Scanning 102
Example: Antivirus with One AV Object 103
Example: Antivirus with Two AV Objects 106
URL Filtering 113
Example: URL Filtering Configuration 119
Chapter 5 Deep Inspection 123
Deep Inspection Overview 124
Attack Object Database Server 128
Example: Immediate Update 129
Example: Automatic Updates 130
Example: Automatic Notification and Immediate Update 132
Example: Manual Update 134
Attack Objects and Groups 136
Stateful Signatures 138
TCP Stream Signatures 139
Protocol Anomalies 139
Attack Object Groups 140
Changing Severity Levels 140
Attack Actions 142
Example: Attack Actions – Close Server, Close, Close Client 143
Mapping Custom Services to Applications 152
Example: Mapping an Application to a Custom Service 153
Customized Attack Objects and Groups 156
User-Defined Stateful SIgnature Attack Objects 156
Contexts 156
Signatures 157
Example: User-Defined Stateful Signature Attack Objects 160
TCP Stream Signature Attack Objects 164
Example: User-Defined Stream Signature Attack Object 164

Granular Blocking of HTTP Components 167
ActiveX Controls.....167
Java Applets.....168
EXE Files 168
ZIP Files.....168
Example: Blocking Java Applets and .exe Files 169
Chapter 6 Suspicious Packet Attributes.....171
ICMP Fragments 172
Large ICMP Packets.....174
Bad IP Options 176
Unknown Protocols.....178
IP Packet Fragments 180
SYN Fragments.....182
Index.....IX--I

Preface

Volume 4, “Attack Detection and Defense Mechanisms” describes the network security options available in ScreenOS. Many of these are SCREEN options that you can enable at the security zone level. Screen options apply to traffic reaching the NetScreen device through any interface bound to a zone for which you have enabled such options. SCREEN options offer protection against IP address and port scans, denial-of-service (DoS) attacks, and other kinds of malicious activity. You can apply other network security options, such as URL filtering, antivirus checking, and intrusion detection and prevention (IDP), at the policy level. These options only apply to traffic under the jurisdiction of the policies in which they are enabled.

Note: *The subject of policies is only presented in this volume peripherally, as it applies to the network security options that you can enable at the policy level. For a complete examination of policies, see “Policies” on page* ***2*** *-197.*

The material within this volume is organized as follows:

- Chapter 1, “Protecting a Network” outlines the basic stages of an attack and the firewall options available to combat the attacker at each stage.
- Chapter 2, “Reconnaissance Deterrence” describes the options available for blocking IP address sweeps, port scans, and attempts to discover the type of operating system (OS) of a targeted system.
- Chapter 3, “Denial-of-Service Attack Defenses” explains firewall, network, and OS-specific DoS attacks and how NetScreen mitigates such attacks.
- Chapter 4, “Content Monitoring and Filtering” describes how to protect Hypertext Transfer Protocol (HTTP) and File Transfer Protocol (FTP) users from malicious uniform resource locators (URLs) and how to configure the NetScreen device to work with third party products to provide antivirus scanning and URL filtering.
- Chapter 5, “Deep Inspection” describes how to configure the NetScreen device to obtain IDP attack object updates, how to create user-defined attack objects and attack object groups, and how to apply IDP at the policy level.
- Chapter 1, “Suspicious Packet Attributes” presents several SCREEN options that protect network resources from potential attacks indicated by unusual IP and ICMP packet attributes.

CONVENTIONS

This document contains several types of conventions, which are introduced in the following sections:

- "CLI Conventions"
- "WebUI Conventions" on page vii
- "Illustration Conventions" on page ix
- "Naming Conventions and Character Types" on page x

CLI Conventions

The following conventions are used when presenting the syntax of a command line interface (CLI) command:

- Anything inside square brackets [ ] is optional.
- Anything inside braces { } is required.
- If there is more than one choice, each choice is separated by a pipe (|). For example,

  ```
  set interface { ethernet1 | ethernet2 | ethernet3 } manage
  ```

  means "set the management options for the ethernet1, ethernet2, or ethernet3 interface".
- Variables appear in *italic.* For example:

  ```
  set admin user name password
  ```

When a CLI command appears within the context of a sentence, it is in **bold** (except for variables, which are always in *italic*). For example: "Use the **get system** command to display the serial number of a NetScreen device."

Note: *When typing a keyword, you only have to type enough letters to identify the word uniquely. For example, typing* ***set adm u joe j12fmt54*** *is enough to enter the command* ***set admin user joe j12fmt54****. Although you can use this shortcut when entering commands, all the commands documented here are presented in their entirety.*

WebUI Conventions

Throughout this book, a chevron (>) is used to indicate navigation through the WebUI by clicking menu options and links. For example, the path to the address configuration dialog box is presented as **Objects > Addresses > List > New**. This navigational sequence is shown below.

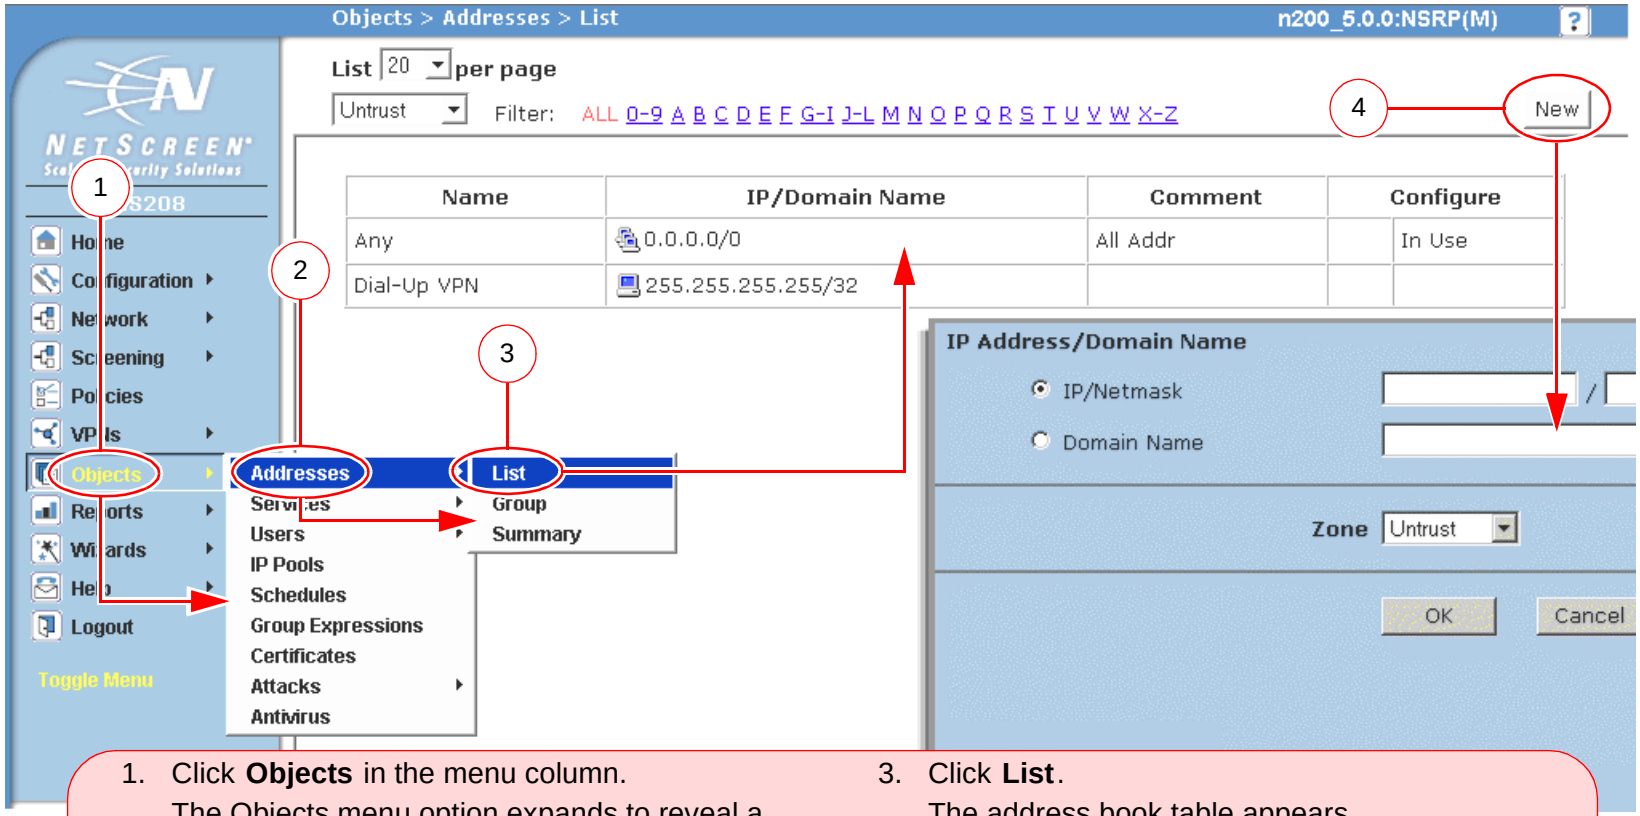

1. Click **Objects** in the menu column.
   The Objects menu option expands to reveal a subset of options for Objects.
2. (Applet menu) Hover the mouse over **Addresses**.
   (DHTML menu) Click **Addresses**.
   The Addresses option expands to reveal a subset of options for Addresses.
3. Click **List**.
   The address book table appears.
4. Click the **New** link.
   The new address configuration dialog box appears.

To perform a task with the WebUI, you must first navigate to the appropriate dialog box where you can then define objects and set parameters. The set of instructions for each task is divided into two parts: a navigational path and configuration details. For example, the following set of instructions includes the path to the address configuration dialog box and the settings for you to configure:

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: addr_1

IP Address/Domain Name:

IP/Netmask: (select), 10.2.2.5/32

Zone: Untrust

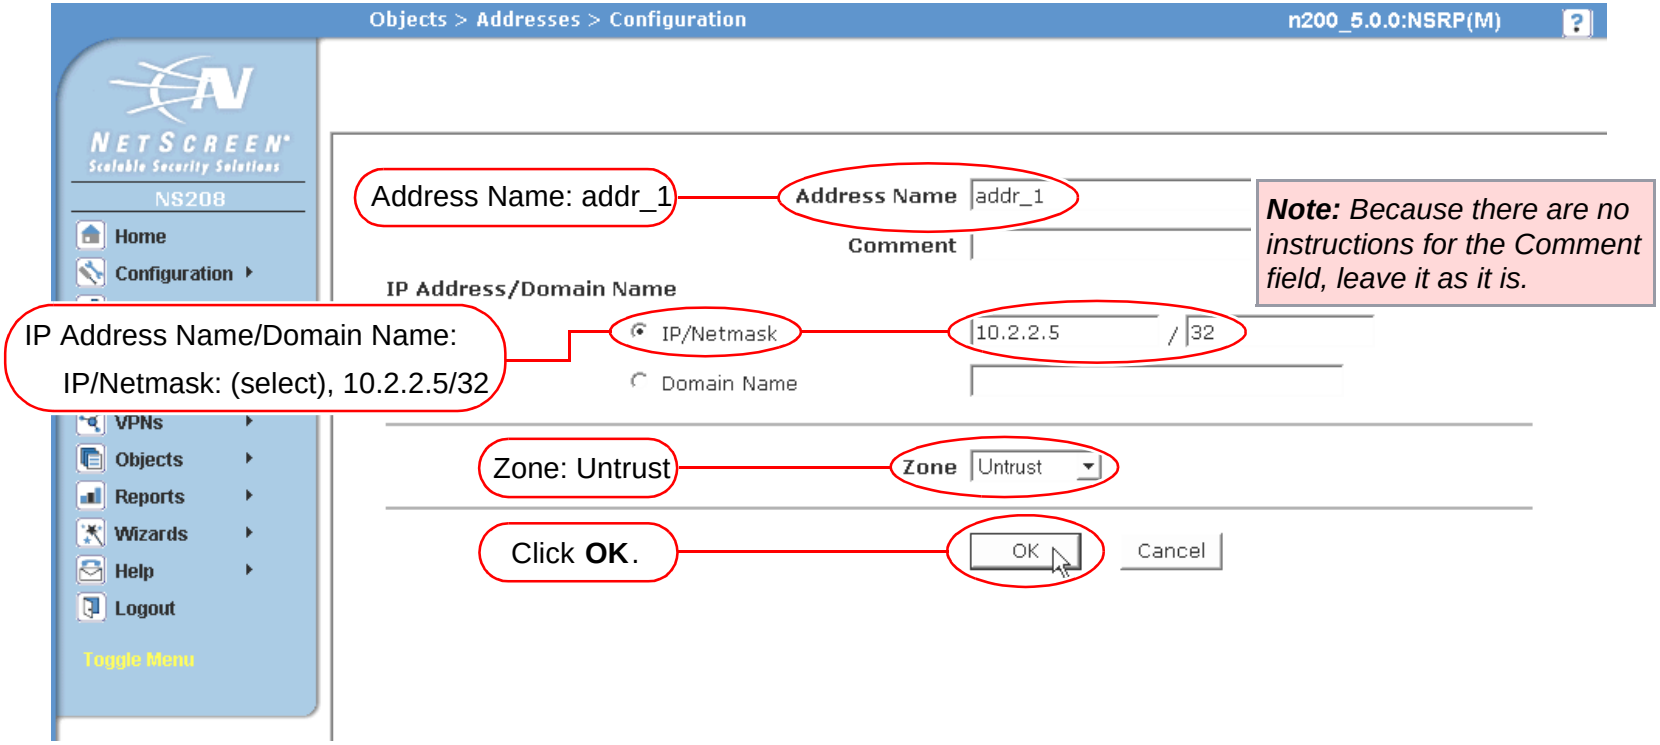

Illustration Conventions

The following graphics make up the basic set of images used in illustrations throughout this book:

Generic NetScreen Device

Virtual Routing Domain

Security Zone

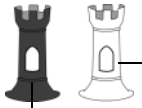

Security Zone Interfaces

White = Protected Zone Interface (example: Trust Zone)

Black = Outside Zone Interface (example: Untrust Zone)

Tunnel Interface

VPN Tunnel

Router Icon

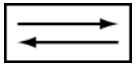

Switch Icon

Local Area Network (LAN) with a Single Subnet

(example: 10.1.1.0/24)

Internet

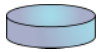

Dynamic IP (DIP) Pool

Desktop Computer

Laptop Computer

Generic Network Device

(examples: NAT server, Access Concentrator)

Server

Naming Conventions and Character Types

ScreenOS employs the following conventions regarding the names of objects—such as addresses, admin users, auth servers, IKE gateways, virtual systems, VPN tunnels, and zones—defined in ScreenOS configurations.

- If a name string includes one or more spaces, the entire string must be enclosed within double quotes ("); for example, **set address trust "local LAN" 10.1.1.0/24**.
- NetScreen trims any spaces leading or trailing text within a set of double quotes; for example, **" local LAN "** becomes **"local LAN"**.
- NetScreen treats multiple consecutive spaces as a single space.
- Name strings are case sensitive, although many CLI key words are case insensitive. For example, **"local LAN"** is different from **"local lan"**.

ScreenOS supports the following character types:

- Single-byte character sets (SBCS) and multiple-byte character sets (MBCS). Examples of SBCS are ASCII, European, and Hebrew. Examples of MBCS—also referred to as double-byte character sets (DBCS)—are Chinese, Korean, and Japanese.

  Note: *A console connection only supports SBCS. The WebUI supports both SBCS and MBCS, depending on the character sets that your Web browser supports.*

- ASCII characters from 32 (0x20 in hexidecimals) to 255 (0xff), except double quotes ("), which have special significance as an indicator of the beginning or end of a name string that includes spaces.

NETSCREEN DOCUMENTATION

To obtain technical documentation for any NetScreen product, visit www.netscreen.com/resources/manuals/.

To obtain the latest software version, visit www.netscreen.com/services/download_soft. Select a category of software product from the dropdown list, then follow the displayed instructions. (You must be a registered user to download NetScreen software.)

If you find any errors or omissions in the following content, please contact us at the e-mail address below:

techpubs@netscreen.com

1

Protecting a Network

There can be many reasons for invading a protected network. The following list contains some common objectives:

- Gathering the following kinds of information about the protected network:
  - The topology of the network
  - The IP addresses of active hosts
  - The numbers of active ports on active hosts
  - The operating systems of active hosts
- Overwhelming a host on a protected network with bogus traffic to induce a Denial-of-Service (DoS)
- Overwhelming the protected network with bogus traffic to induce a network-wide DoS
- Overwhelming a firewall with bogus traffic, and thereby inducing a DoS for the network behind it
- Causing damage to and stealing data from a host on the protected network
- Gaining access to a host on the protected network to obtain information
- Gaining control of a host to launch other exploits
- Gaining control of a firewall to control access to the network that it protects

ScreenOS provides detective and defensive tools to uncover and thwart the efforts of attackers to achieve the above objectives when they attempt to target a network protected by a NetScreen device.

This chapter first presents an overview of the main stages of an attack and of the various defense mechanisms that you can employ to thwart an attack at each stage:

- "Stages of an Attack" on page 2
- "Detection and Defense Mechanisms" on page 3
- "Exploit Monitoring" on page 5

STAGES OF AN ATTACK

Each attack typically progresses in two major stages. In the first stage, the attacker gathers information, and in the second stage he or she launches the attack.

1. Perform reconnaissance.
    1. Map the network and determine which hosts are active (IP address sweep).
    2. Discern which ports are active (port scans) on the hosts discovered by the IP address sweep.
    3. Determine the operating system (OS), which might expose a weakness in the OS or suggest an attack to which that particular OS is susceptible.
2. Launch the attack.
    1. Conceal the origin of the attack.
    2. Perform the attack.
    3. Remove or hide evidence.

DETECTION AND DEFENSE MECHANISMS

An exploit can be an information-gathering probe or an attack to compromise, disable, or harm a network or network resource. In some cases, the distinction between the two objectives of an exploit can be unclear. For example, a barrage of TCP SYN segments might be an IP address sweep with the intent of triggering responses from active hosts, or it might be a SYN flood attack with the intent of overwhelming a network so that it can no longer function properly. Furthermore, because an attacker usually precedes an attack by performing reconnaissance on the target, we can consider information-gathering efforts as a precursor to an impending attack—that is, they constitute the first stage of an attack. Thus, the term “exploit” encompasses both reconnaissance and attack activities, and the distinction between the two is not always clear cut.

NetScreen provides various detection methods and defense mechanisms at the zone and policy levels to combat exploits at all stages of their execution:

- SCREEN options at the zone level[1]
- Firewall policies at the inter-, intra-, and super-zone policy levels. (“super-zone” meaning global policies, where no security zones are referenced)

To secure all connection attempts, NetScreen devices use a dynamic packet filtering method known as stateful inspection. Using this method, the NetScreen device notes various components in the IP packet and TCP segment headers— source and destination IP addresses, source and destination port numbers, and packet sequence numbers—and maintains the state of each TCP session and pseudo UDP session traversing the firewall. (The NetScreen device also modifies session states based on changing elements such as dynamic port changes or session termination.) When a responding TCP packet arrives, the NetScreen device compares the information reported in its header with the state of its associated session stored in the inspection table. If they match, the responding packet is allowed to pass the firewall. If the two do not match, the packet is dropped.

NetScreen SCREEN options secure a zone by inspecting, and then allowing or denying, all connection attempts that require crossing an interface bound to that zone. The NetScreen device then applies firewall policies, which can contain content filtering and intrusion detection and prevention (IDP) components, to the traffic that passes the SCREEN filters.

1. Although the VLAN and MGT zones are function zones and not security zones, you can set SCREEN options for them. The VLAN zone supports the same set of SCREEN options as a Layer 3 security zone. (Layer 2 security zones support an additional SYN flood option that Layer 3 zones do not: Drop Unknown MAC). Because the following SCREEN options do not apply to the MGT zone, they are not available for that zone: SYN flood protection, SYN-ACK-ACK proxy flood protection, HTTP component blocking, and WinNuke attack protection.

The sets of defense mechanisms that a NetScreen firewall provides for network protection are outlined below:

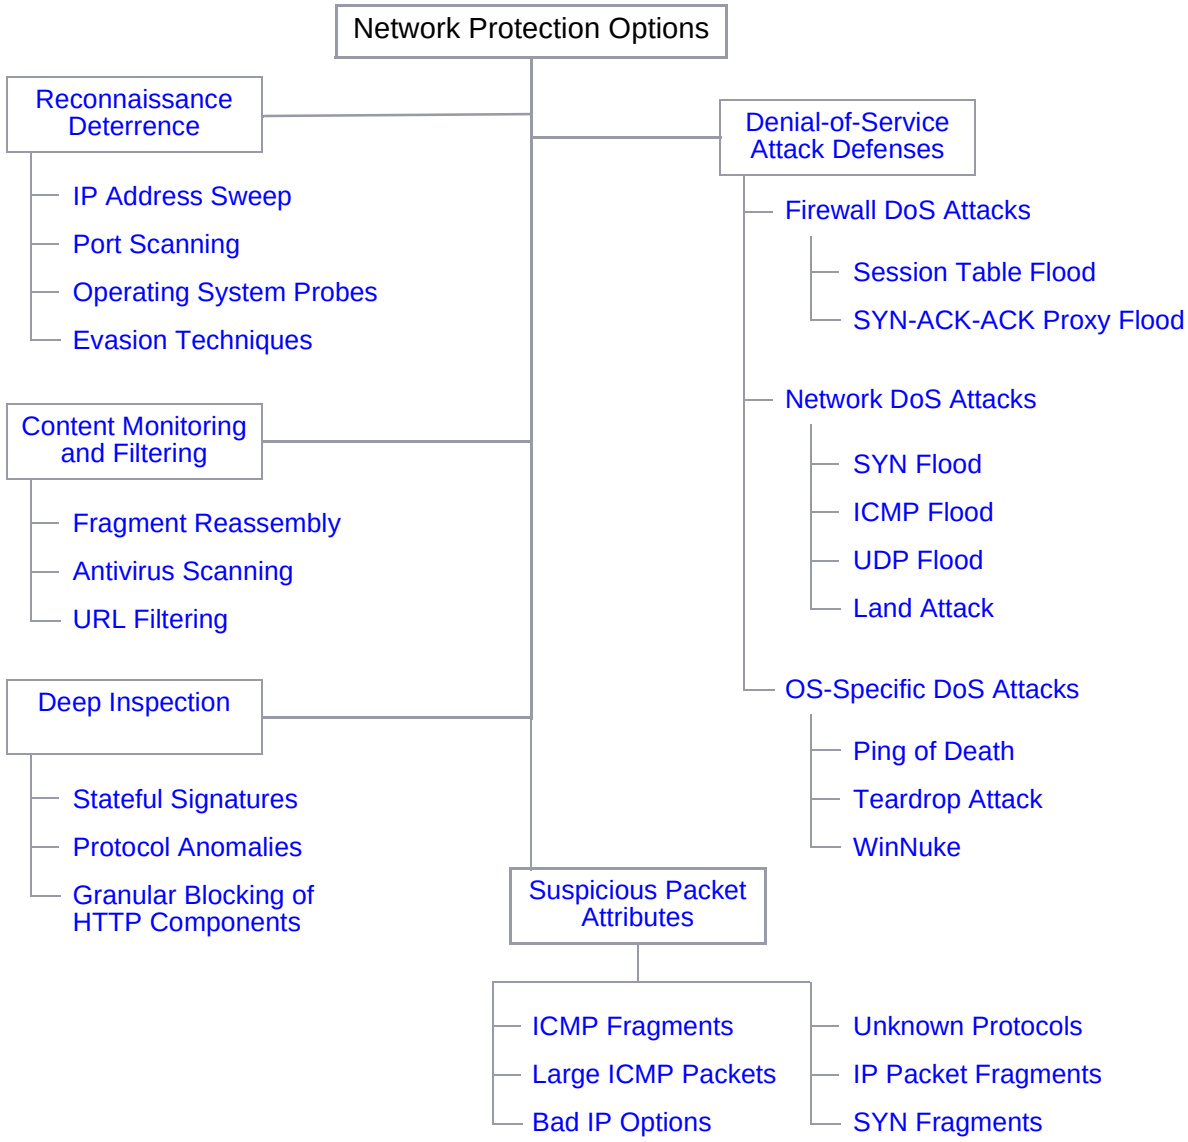

As previously stated, NetScreen network protection settings operate at two levels: security zone and policy. The NetScreen device performs reconnaissance deterrence and DoS attack defenses at the security zone level. In the area of content monitoring and filtering, the NetScreen device applies fragment reassembly at the zone level and antivirus (AV) scanning and uniform resource locator (URL) filtering at the policy level. The NetScreen device applies IDP at the policy level, except for the detection and blocking of HTTP components, which occurs at the zone level. Zone-level firewall settings are SCREEN options. A network protection option set in a policy is a component of that policy.

EXPLOIT MONITORING

Although you typically want the NetScreen device to block exploits, there might be times when you want to gather intelligence about them. You might want to learn specifically about a particular exploit—to discover its intention, its sophistication, and possibly (if the attacker is careless or unsophisticated) its source.

If you want to gather information about an exploit, you can let it occur, monitor it, analyze it, perform forensics, and then respond as delineated in a previously prepared incident response plan. You can instruct the NetScreen device to notify you of an exploit, but instead of taking action, the NetScreen device allows the exploit to transpire. You can then study what occurred, and try to understand the attacker's method, strategy, and objectives. Increased understanding of the threat to the network can then allow you to better fortify your defenses. Although a smart attacker can conceal his or her location and identity, you might be able to gather enough information to discern where the attack originated. You also might be able to estimate the attacker's capabilities. This kind of information allows you to gauge your response.

Example: Monitoring Attacks from the Untrust Zone

In this example, IP spoofing attacks from the Untrust zone have occurred on a daily basis, usually between 9:00 PM and 12:00 AM. Instead of dropping the packets with the spoofed source IP addresses, you want the NetScreen device to notify you of their arrival but allow them to pass, perhaps directing them to a honeypot[2] that you have connected on the DMZ interface connection. At 8:55 PM, you change the firewall behavior from notification and rejection of packets belonging to a detected attack to notification and acceptance. When the attack occurs, you can then use the honeypot to monitor the attacker's activity after crossing the firewall. You might also work in cooperation with the upstream ISP to begin tracking the source of the packets back to their source.

WebUI

Screening > Screen (Zone: Untrust): Enter the following, and then click **Apply**:

Generate Alarms without Dropping Packet: (select)

IP Address Spoof Protection: (select)

CLI

```
set zone untrust screen alarm-without-drop
set zone untrust screen ip-spoofing
save
```

2. A honeypot is a decoy network server that is designed to lure attackers and then record their actions during an attack.

2

Reconnaissance Deterrence

Attackers can better plan their attack when they first know the layout of the targeted network (which IP addresses have active hosts), the possible entry points (which port numbers are active on the active hosts), and the constitution of their victims (which operating system the active hosts are running). To gain this information, attackers must perform reconnaissance. NetScreen provides several SCREEN options to deter attackers' reconnaissance efforts and thereby hinder them from obtaining valuable information about the protected network and network resources.

- "IP Address Sweep" on page 8
- "Port Scanning" on page 10
- "Network Reconnaissance Using IP Options" on page 12
- "Operating System Probes" on page 16
  - "SYN and FIN Flags Set" on page 16
  - "FIN Flag without ACK Flag" on page 18
  - "TCP Header without Flags Set" on page 20
- "Evasion Techniques" on page 22
  - "FIN Scan" on page 22
  - "IP Spoofing" on page 22
  - "IP Source Route Options" on page 31

IP ADDRESS SWEEP

An address sweep occurs when one source IP address sends 10 ICMP packets to different hosts within a defined interval (5000 microseconds is the default). The purpose of this scheme is to send ICMP packets—typically echo requests—to various hosts in the hopes that at least one replies, thus uncovering an address to target. The NetScreen device internally logs the number of ICMP packets to different addresses from one remote source. Using the default settings, if a remote host sends ICMP traffic to 10 addresses in 0.005 seconds (5000 microseconds), NetScreen flags this as an address sweep attack, and rejects the 11th and all further ICMP packets from that host for the remainder of that second.

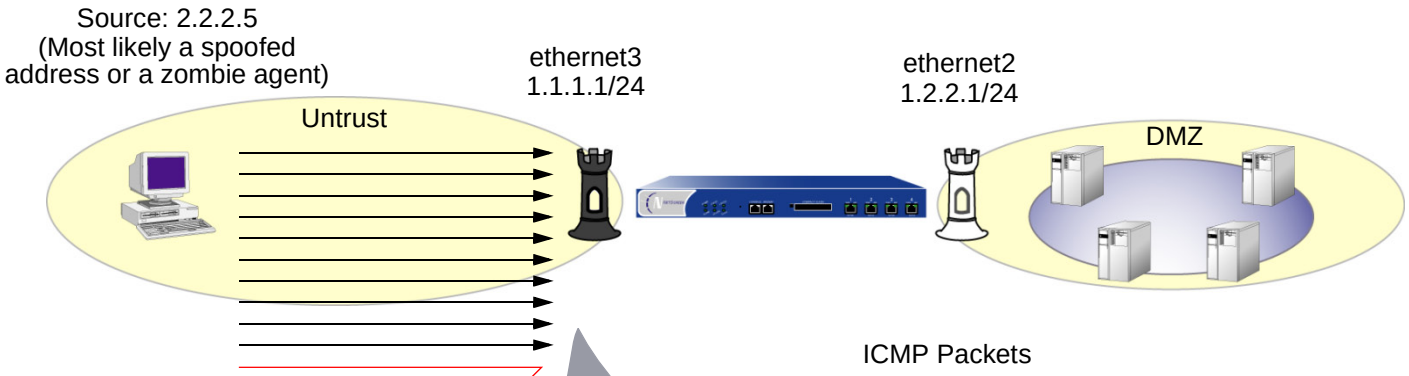

ICMP Packets

	Src addr	Dst addr
	2.2.2.5	1.2.2.5
	2.2.2.5	1.2.2.160
	2.2.2.5	1.2.2.84
	2.2.2.5	1.2.2.211
	2.2.2.5	1.2.2.10
	2.2.2.5	1.2.2.20
	2.2.2.5	1.2.2.21
	2.2.2.5	1.2.2.240
	2.2.2.5	1.2.2.17
	2.2.2.5	1.2.2.123
Rejected	2.2.2.5	1.2.2.6

The NetScreen device makes an entry in its session table for the first 10 ICMP packets from 2.2.2.5 and does a route lookup and policy lookup for these. If no policy permits these packets, the NetScreen device tags these as invalid and removes them from the session table in the next “garbage sweep”, which occurs every two seconds. After the tenth packet, the NetScreen device rejects all further ICMP traffic from 2.2.2.5.

Note: *A* ***zombie agent*** *is a compromised host under the covert control of an attacker.*

Consider enabling this SCREEN option for a security zone only if there is a policy permitting ICMP traffic from that zone. Otherwise, you do not need to enable it. The lack of such a policy denies all ICMP traffic from that zone, precluding an attacker from successfully performing an IP address sweep anyway.

To block IP address sweeps originating in a particular security zone, do either of the following:

WebUI

Screening > Screen (Zone: select a zone name): Enter the following, and then click **Apply**:

IP Address Sweep Protection: (select)

Threshold: (enter a value to trigger IP address sweep protection[1])

CLI

```
set zone zone screen ip-sweep threshold number
set zone zone screen ip-sweep
```

1. The value unit is microseconds. The default value is 5000 microseconds.

PORT SCANNING

A port scan occurs when one source IP address sends IP packets containing TCP SYN segments to 10 different ports at the same destination IP address within a defined interval (5000 microseconds is the default). The purpose of this scheme is to scan the available services in the hopes that at least one port will respond, thus identifying a service to target. The NetScreen device internally logs the number of different ports scanned from one remote source. Using the default settings, if a remote host scans 10 ports in 0.005 seconds (5000 microseconds), NetScreen flags this as a port scan attack, and rejects all further packets from the remote source (regardless of the destination IP address) for the remainder of that second.

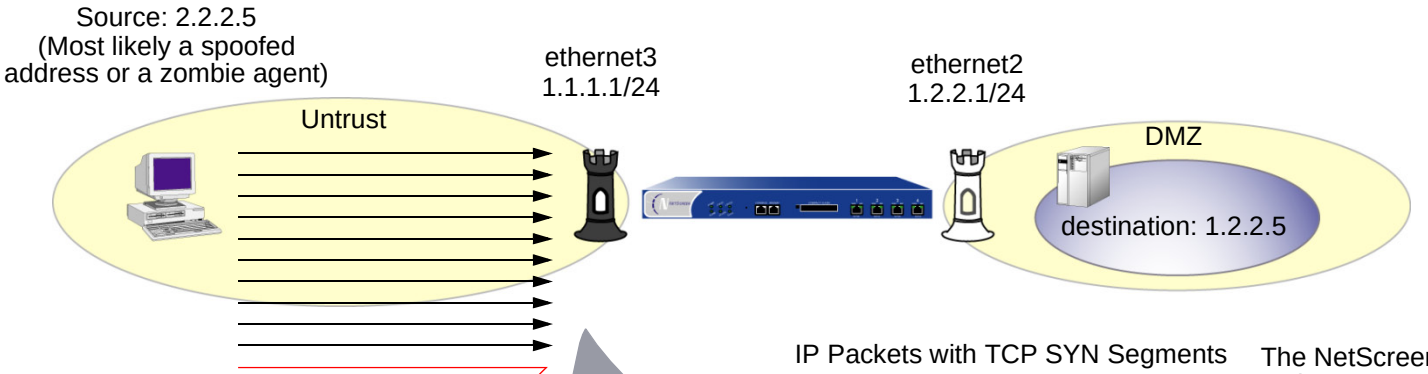

(Note that after 10 IP packets containing TCP SYN segments to different ports at the same destination IP address, the NetScreen device logs this as a port scan and rejects further packets from the source address.)

IP Packets with TCP SYN Segments

	Dst addr:port	Dst addr:port
	2.2.2.5:17820	1.2.2.5:21
	2.2.2.5:42288	1.2.2.5:23
	2.2.2.5:22814	1.2.2.5:53
	2.2.2.5:15401	1.2.2.5:80
	2.2.2.5:13373	1.2.2.5:111
	2.2.2.5:33811	1.2.2.5:113
	2.2.2.5:17821	1.2.2.5:123
	2.2.2.5:19003	1.2.2.5:129
	2.2.2.5:26450	1.2.2.5:137
	2.2.2.5:38087	1.2.2.5:138
Rejected	2.2.2.5:24111	1.2.2.5:139

The NetScreen device makes an entry in its session table for the first 10 connection attempts from 2.2.2.5 to 1.2.2.5 and does a route lookup and policy lookup for these. If no policy permits these connection attempts, the NetScreen device tags these as invalid and removes them from the session table in the next "garbage sweep", which occurs every two seconds. After the tenth attempt, the NetScreen device rejects all further connection attempts from 2.2.2.5.

To block port scans originating in a particular security zone, do either of the following:

WebUI

Screening > Screen (Zone: select a zone name): Enter the following, and then click **Apply**:

Port Scan Protection: (select)

Threshold: (enter a value to trigger protection against port scans[2])

CLI

```
set zone zone screen port-scan threshold number
set zone zone screen port-scan
```

2. The value unit is microseconds. The default value is 5000 microseconds.

Network Reconnaissance Using IP Options

The Internet Protocol standard “RFC 791, Internet Protocol” specifies a set of options to provide special routing controls, diagnostic tools, and security. These options appear after the destination address in an IP packet header.

IP Header

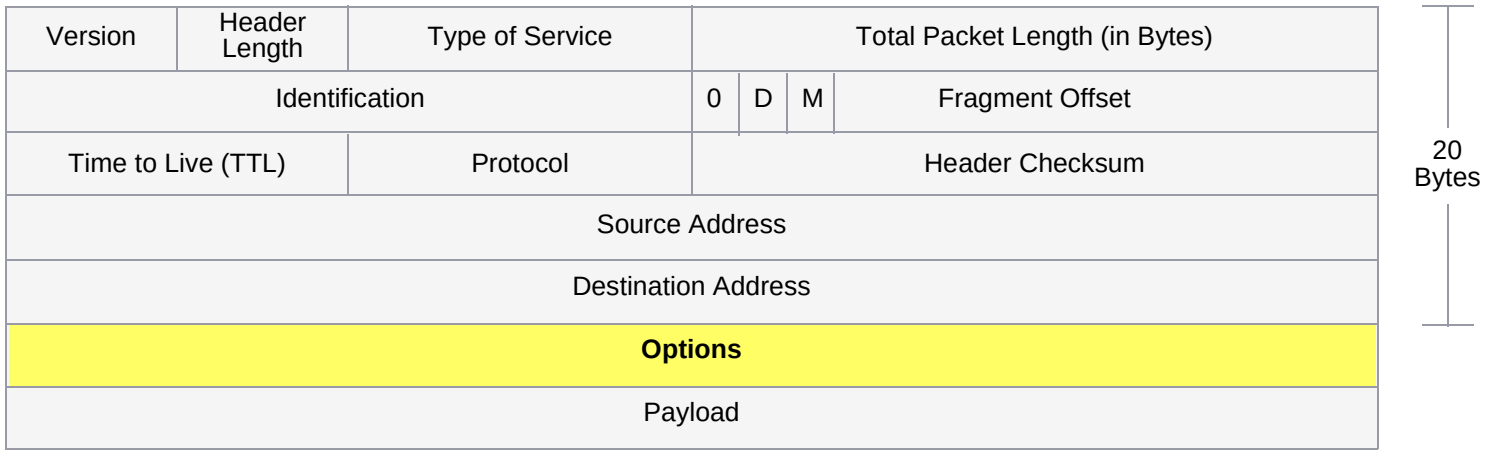

RFC 791 admits that these options are “unnecessary for the most common communications“ and, in reality, they rarely appear in IP packet headers. When they do appear, they are frequently being put to some nefarious use. The following is a list of all the IP options and their accompanying attributes:

Type	Class	Number	Length	Intended Use	Nefarious Use
End of Options	0*	0	0	Indicates the end of one or more IP options.	None
No Options	0	1	0	Indicates that there are no IP options in the header.	None

Type	Class	Number	Length	Intended Use	Nefarious Use
Security	0	2	11 bits	Provides a way for hosts to send security, compartmentation, TCC (closed user group) parameters, and Handling Restriction Codes compatible with Department of Defense (DoD) requirements. (This option, as specified in RFC 791 and RFC 1038, is obsolete.)	Unknown, but because it is obsolete, its presence in an IP header is suspect.
Loose Source Route	0	3	Varies	Specifies a partial route list for a packet to take on its journey from source to destination. The packet must proceed in the order of addresses specified, but it is allowed to pass through other routers in between those specified.	Evasion. The attacker can use the specified routes to hide the true source of a packet or gain access to a protected network. (See “IP Source Route Options” on page 31.)
Record Route	0	7	Varies	Records the IP addresses of the network devices along the path that the IP packet travels. The destination machine can then extract and process the route information. (Due to the size limitation of 40 bytes for both the option and storage space, this can only record up to 9 IP addresses.)	Reconnaissance. If the destination host is a compromised machine in the attacker’s control, he or she can glean information about the topology and addressing scheme of the network through which the packet passed.
Stream ID	0	8	4 bits	(Obsolete) Provided a way for the 16-bit SATNET stream identifier to be carried through networks that did not support the stream concept.	Unknown, but because it is obsolete, its presence in an IP header is suspect.

Type	Class	Number	Length	Intended Use	Nefarious Use
Strict Source Route	0	9	Varies	Specifies the complete route list for a packet to take on its journey from source to destination. The last address in the list replaces the address in the destination field.	Evasion. An attacker can use the specified routes to hide the true source of a packet or gain access to a protected network. (See "IP Source Route Options" on page 31.)
Timestamp	2[†]	4		Records the time (in Universal Time[‡]) when each network device receives the packet during its trip from the point of origin to its destination. The network devices are identified by IP number. This option develops a list of IP addresses of the routers along the path of the packet and the duration of transmission between each one.	Reconnaissance. If the destination host is a compromised machine in the attacker's control, he or she can glean information about the topology and addressing scheme of the network through which the packet passed.

[*] The class of options identified as "0" was designed to provide extra packet or network control.

[†] The class of options identified as "2" was designed diagnostics, debugging, and measurement

[‡] The timestamp uses the number of milliseconds since midnight Universal Time (UT). UT is also known as "Greenwich Mean Time" (GMT), which is the basis for the international time standard.

The following SCREEN options detect IP options that an attacker can use for reconnaissance or for some unknown but suspect purpose:

- **Record Route:** The NetScreen device detects packets where the IP option is 7 (Record Route) and records the event in the SCREEN counters list for the ingress interface.
- **Timestamp:** The NetScreen device detects packets where the IP option list includes option 4 (Internet Timestamp) and records the event in the SCREEN counters list for the ingress interface.
- **Security:** The NetScreen device detects packets where the IP option is 2 (security) and records the event in the SCREEN counters list for the ingress interface.
- **Stream ID:** The NetScreen device detects packets where the IP option is 8 (Stream ID) and records the event in the SCREEN counters list for the ingress interface.

To detect packets with the above IP options set, do either of the following, where the specified security zone is the one from which the packets originate:

WebUI

Screening > Screen (Zone: select a zone name): Enter the following, and then click **Apply**:

IP Record Route Option Detection: (select)

IP Timestamp Option Detection: (select)

IP Security Option Detection: (select)

IP Stream Option Detection: (select)

CLI

```
set zone zone screen ip-record-route
set zone zone screen ip-timestamp-opt
set zone zone screen ip-security-opt
set zone zone screen ip-stream-opt
```

Operating System Probes

Before launching an exploit, an attacker might try to probe the targeted host to learn its operating system (OS). With that knowledge, he can better decide which attack to launch and which vulnerabilities to exploit. A NetScreen device can block reconnaissance probes commonly used to gather information about OS types.

SYN and FIN Flags Set

Both the SYN and FIN control flags are not normally set in the same TCP segment header. The SYN flag synchronizes sequence numbers to initiate a TCP connection. The FIN flag indicates the end of data transmission to finish a TCP connection. Their purposes are mutually exclusive. A TCP header with the SYN and FIN flags set is anomalous TCP behavior, causing various responses from the recipient, depending on the OS.

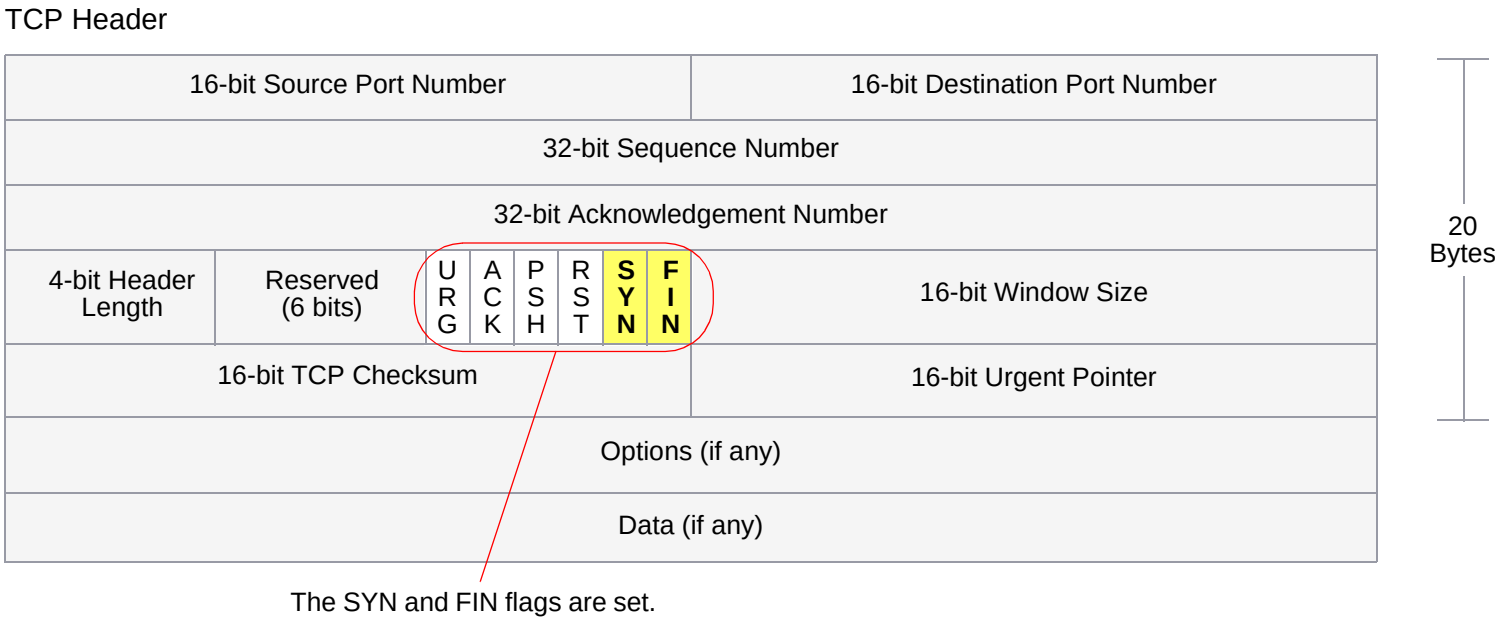

An attacker can send a segment with both flags set to see what kind of system reply is returned and thereby determine what kind of OS is on the receiving end. The attacker can then use any known system vulnerabilities for further attacks.

When you enable this SCREEN option, the NetScreen device checks if the SYN and FIN flags are set in TCP headers. If it discovers such a header, it drops the packet.

To block packets with both the SYN and FIN flags set, do either of the following, where the specified security zone is the one from which the packets originate:

WebUI

Screening > Screen (Zone: select a zone name): Select **SYN and FIN Bits Set Protection**, and then click **Apply**.

CLI

```
set zone zone screen syn-fin
```

FIN Flag without ACK Flag

TCP segments with the FIN control flag set (to signal the conclusion of a session and terminate the connection) normally also have the ACK flag set (to acknowledge the previous packet received). Because a TCP header with the FIN flag set but not the ACK flag is anomalous TCP behavior, there is no uniform response to this[3]. The OS might respond by sending a TCP segment with the RST flag set. Another might completely ignore it. The victim's response can provide the attacker with a clue as to its OS. (Other purposes for sending a TCP segment with the FIN flag set are to evade detection while performing address and port scans and to evade defenses on guard for a SYN flood by performing a FIN flood instead. For information about FIN scans, see "FIN Scan" on page 22.)

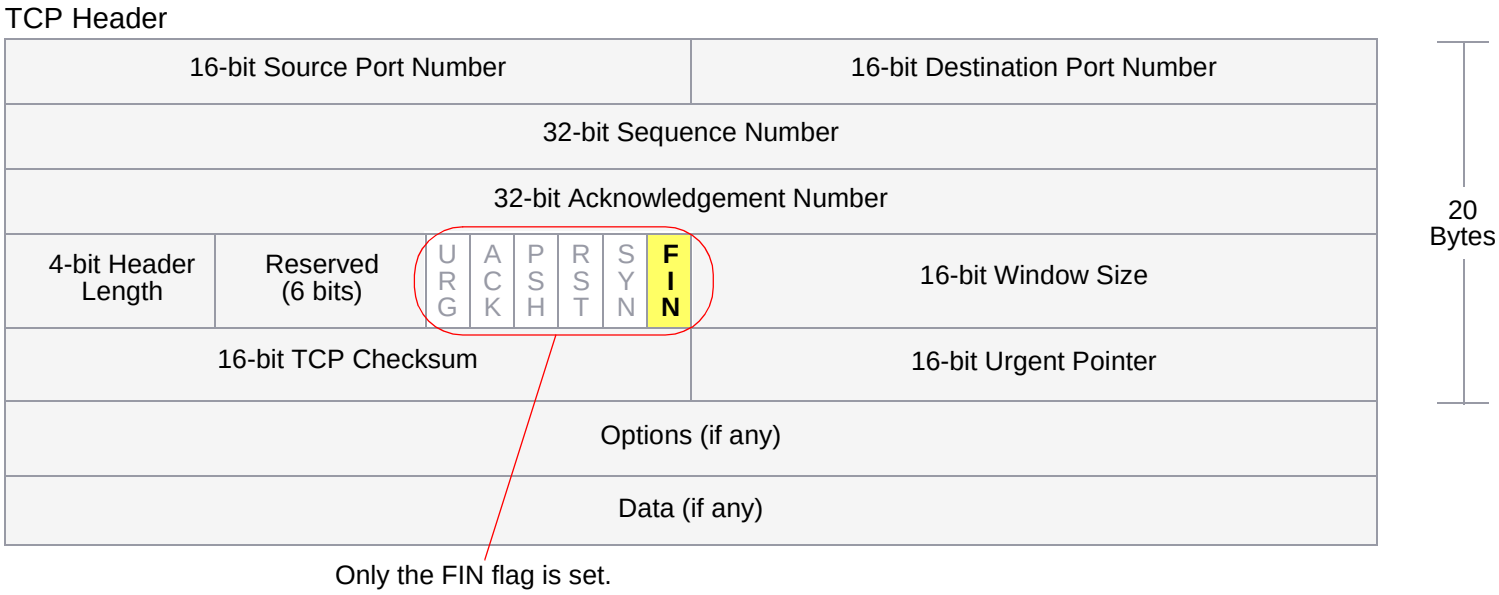

When you enable this SCREEN option, the NetScreen device checks if the FIN flag is set but not the ACK flag in TCP headers. If it discovers a packet with such a header, it drops the packet.

3. Vendors have interpreted RFC 793 "Transmission Control Protocol" variously when designing their TCP/IP implementations. When a TCP segment arrives with the FIN flag set but not the ACK flag, some implementations send RST segments. Some drop the packet without sending a RST.

To block packets with the FIN flag set but not the ACK flag, do either of the following, where the specified security zone is the one from which the packets originate:

WebUI

Screening > Screen (Zone: select a zone name): Select **FIN Bit with No ACK Bit in Flags Protection**, and then click **Apply**.

CLI

```
set zone zone screen fin-no-ack
```

TCP Header without Flags Set

A normal TCP segment header has at least one flag control set. A TCP segment with no control flags set is an anomalous event. Because different operating systems respond differently to such anomalies, the response (or lack of response) from the targeted device can provide a clue as to the type of OS it is running.

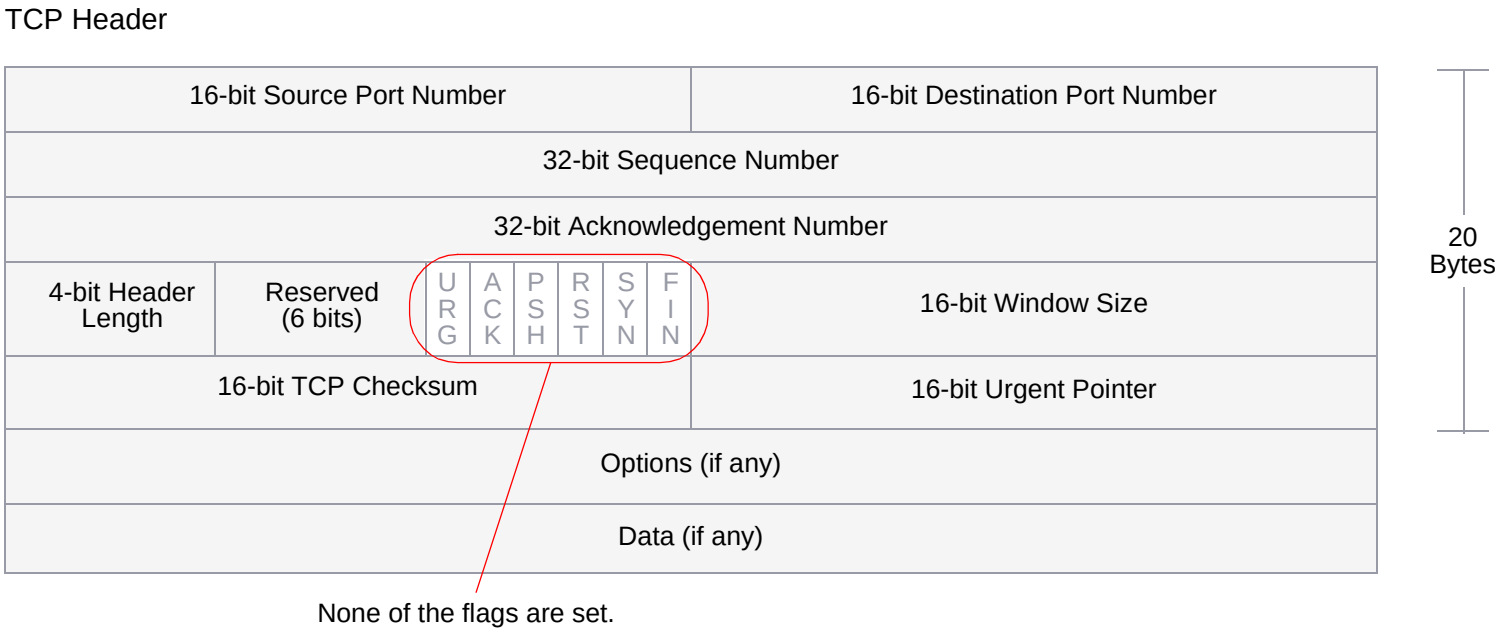

When you enable the NetScreen device to detect TCP segment headers with no flags set, the NetScreen device drops all TCP packets with a missing or malformed flags field.

To block packets with no flags set, do either of the following, where the specified security zone is the one from which the packets originate:

WebUI

Screening > Screen (Zone: select a zone name): Select **TCP Packet without Flag Protection**, and then click **Apply**.

CLI

```
set zone zone screen tcp-no-flag
```

Evasion Techniques

Whether gathering information or launching an attack, it generally behooves the attacker to avoid detection. Although some IP address and port scans are blatant and easily detectable, more wily attackers use a variety of means to conceal their activity. Such techniques as using FIN scans instead of SYN scans—which attackers know most firewalls and intrusion detection programs detect—indicate an evolution of reconnaissance and exploit techniques to evade detection and successfully accomplish their tasks. (to elicit a RST and thereby discover the IP address of an active host)

FIN Scan

A FIN scan sends TCP segments with the FIN flag set in the attempt to provoke a response (a TCP segment with the RST flag set) and thereby discover an active host or an active port on a host. An attacker might use this approach rather than perform an address sweep with ICMP echo requests or an address scan with SYN segments because he or she knows that many firewalls typically guard against the latter two approaches—but not necessarily against FIN segments. The use of TCP segments with the FIN flag set might evade detection and thereby help the attacker succeed in his or her reconnaissance efforts.

The packet flow behavior for a NetScreen device is to reject TCP segments with non-SYN flags set unless they belong to an established session. NetScreen devices never allow such unsolicited segments to reach a protected host. In addition, you can enable the SCREEN option that specifically blocks TCP segments with the FIN flag set but not the ACK flag, which is anomalous for a TCP segment.

IP Spoofing

One method of attempting to gain access to a restricted area of the network is to insert a bogus source address in the packet header to make the packet appear to come from a trusted source. This technique is called IP spoofing. NetScreen has two IP spoofing detection methods, both of which accomplish the same task: determining that the packet came from a location other than that indicated in its header. The method that a NetScreen device uses depends if it is operating at Layer 3 or Layer 2 in the OSI model.

- **Layer 3** – When interfaces on the NetScreen device are operating in Route or NAT mode, the mechanism to detect IP spoofing relies on route table entries. If, for example, a packet with source IP address 10.1.1.6 arrives at ethernet3, but the NetScreen device has a route to 10.1.1.0/24 through ethernet1, IP spoof checking notes that this address arrived at an invalid interface—as defined in the route table, a valid packet from 10.1.1.6 can only arrive via ethernet1, not ethernet3. Therefore, the device concludes that the packet has a spoofed source IP address and discards it.

  If the source IP address in a packet does not appear in the route table, by default the NetScreen device allows that packet to pass (assuming that a policy exists permitting it). Using the following CLI command—where the specified security zone is the one from which the packets originate—you can instruct the NetScreen device to drop any packet whose source IP address is not in the route table:

  set zone *zone* screen ip-spoofing drop-no-rpf-route

1. An IP packet arrives at ethernet3. Its source IP address is 10.1.1.6.

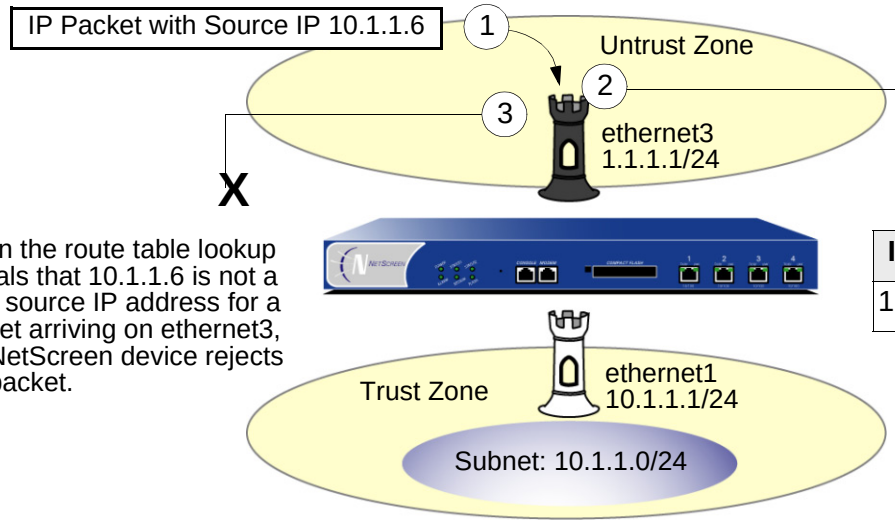

2. Because IP spoof protection is enabled in the Untrust zone, the NetScreen device checks if 10.1.1.6 is a valid source IP address for a packet arriving on ethernet3.

Route Table

ID	IP-Prefix	Interface	Gateway	P
1	10.1.10/24	eth1	0.0.0.0	C

3. When the route table lookup reveals that 10.1.1.6 is not a valid source IP address for a packet arriving on ethernet3, the NetScreen device rejects the packet.

- **Layer 2** – When interfaces on the NetScreen device are operating in Transparent mode, the IP spoof checking mechanism makes use of the address book entries. For example, you define an address for "serv A" as 1.2.2.5/32 in the V1-DMZ zone. If a packet with source IP address 1.2.2.5 arrives at a V1-Untrust zone interface (ethernet3), IP spoof checking notes that this address arrived at an invalid interface. The address belongs to the V1-DMZ zone, not to the V1-Untrust zone, and is accepted only at ethernet2, which is bound to V1-DMZ. The device concludes that packet has a spoofed source IP address and discards it.

1. An IP packet arrives from the V1-Untrust zone. Its source IP address is 1.2.2.5.

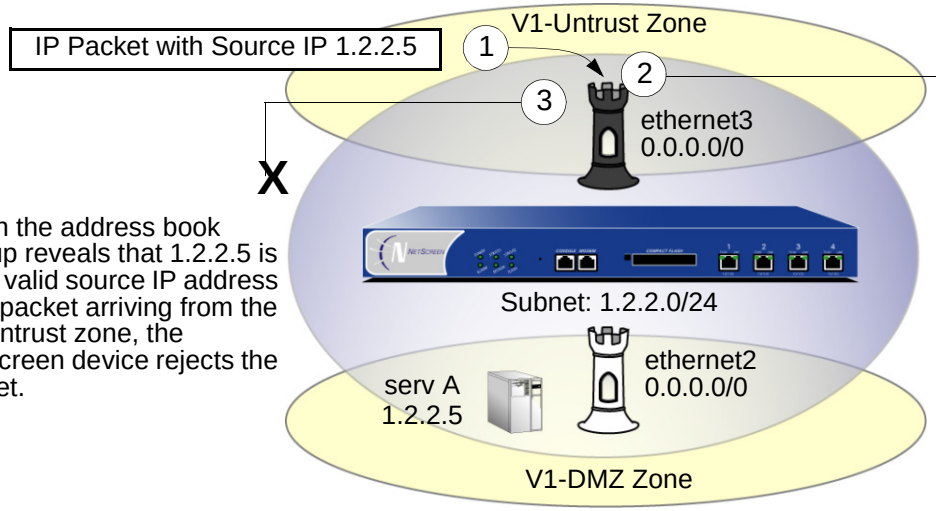

2. Because IP spoof protection is enabled in the V1-Untrust zone, the NetScreen device checks if 1.2.2.5 is a valid source IP address for a packet arriving from the V1-Untrust zone.

Address Zone Name: V1-DMZ

Name	Address	Netmask
serv A	1.2.2.5	255.255.255.255

3. When the address book lookup reveals that 1.2.2.5 is not a valid source IP address for a packet arriving from the V1-Untrust zone, the NetScreen device rejects the packet.

Be careful when defining addresses for the subnet that straddles multiple security zones. In the above illustration, 1.2.2.0/24 belongs to both the V1-Untrust and V1-DMZ zones. If you configure the NetScreen device as follows, the device will block traffic from the V1-DMZ zone that you want it to permit:

- You define an address for 1.2.2.0/24 in the V1-Untrust zone.
- You have a policy permitting traffic from any address in the V1-DMZ zone to any address in the V1-Untrust zone (**set policy from v1-dmz to v1-untrust any any any permit**).
- You enable IP spoof checking.

Because addresses in the V1-DMZ zone are also in the 1.2.2.0/24 subnet, when traffic from these addresses reaches ethernet2, the IP spoof check refers to the address book and finds 1.2.2.0/24 in the V1-Untrust zone. Consequently, the NetScreen device blocks the traffic.

Example: L3 IP Spoof Protection

In this example, you enable IP spoof protection for the Trust, DMZ, and Untrust zones for a NetScreen device operating at Layer 3. By default, the NetScreen device automatically makes entries in the route table for the subnets specified in interface IP addresses. In addition to these automatic route table entries you manually enter the following three routes:

Destination:	Egress Interface:	Next Gateway:
10.1.2.0/24	ethernet1	10.1.1.250
1.2.3.0/24	ethernet2	1.2.2.250
0.0.0.0/0	ethernet3	1.1.1.250

If you enable the IP spoof protection SCREEN option but do not enter the above three routes, the NetScreen device will drop all traffic from the addresses in the “Destination” column and enter alarms in the event log. For example, if a packet with the source address 10.1.2.5 arrives at ethernet1 and there is no route to the 10.1.2.0/24 subnet via ethernet1, the NetScreen device will determine that that packet has arrived at an invalid interface and drop it.

All the security zones in this example are in the trust-vr routing domain.

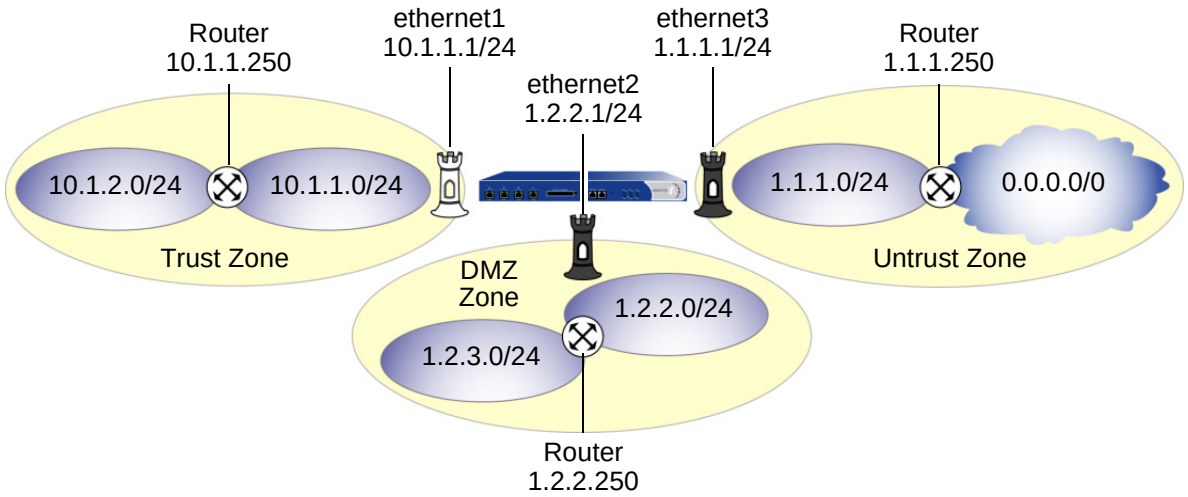

WebUI

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Enter the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet2): Enter the following, and then click **OK**:

Zone Name: DMZ

Static IP: (select this option when present)

IP Address/Netmask: 1.2.2.1/24

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

2. Routes

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 10.1.2.0/24

Gateway: (select)

Interface: ethernet1

Gateway IP Address: 10.1.1.250

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 1.2.3.0/24

Gateway: (select)

Interface: ethernet2

Gateway IP Address: 1.2.2.250

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 1.1.1.250

3. IP Spoof Protection

Screening > Screen (Zone: Trust): Select **IP Address Spoof Protection**, and then click **Apply**.

Screening > Screen (Zone: DMZ): Select **IP Address Spoof Protection**, and then click **Apply**.

Screening > Screen (Zone: Untrust): Select **IP Address Spoof Protection**, and then click **Apply**.

CLI

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet2 zone dmz
set interface ethernet2 ip 1.2.2.1/24

set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
```

2. Routes

```
set vrouter trust-vr route 10.1.2.0/24 interface ethernet1 gateway 10.1.1.250
set vrouter trust-vr route 1.2.3.0/24 interface ethernet2 gateway 1.2.2.250
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
```

3. IP Spoof Protection

```
set zone trust screen ip-spoofing
set zone dmz screen ip-spoofing
set zone untrust screen ip-spoofing
save
```

Example: L2 IP Spoof Protection

In this example, you protect the V1-DMZ zone from IP spoofing on traffic originating in the V1-Untrust zone. First, you define the following addresses for three Web servers in the V1-DMZ zone:

- servA: 1.2.2.10
- servB: 1.2.2.20
- servC: 1.2.2.30

You then enable IP spoofing in the V1-Untrust zone.

If an attacker in the V1-Untrust zone attempts to spoof the source IP address using any of the three addresses in the V1-DMZ zone, the NetScreen device checks the address against those in the address books. When it finds that the source IP address on a packet coming from the V1-Untrust zone belongs to a defined address in the V1-DMZ zone, the NetScreen device rejects the packet.

WebUI

1. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: servA

IP Address/Domain Name:

IP/Netmask: (select), 1.2.2.10/32

Zone: V1-DMZ

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: servB

IP Address/Domain Name:

IP/Netmask: (select), 1.2.2.20/32

Zone: V1-DMZ

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: servC

IP Address/Domain Name:

IP/Netmask: (select), 1.2.2.30/32

Zone: V1-DMZ

2. IP Spoof Protection

Screening > Screen (Zone: V1-Trust): Select **IP Address Spoof Protection**, and then click **Apply**.

CLI

1. Addresses

```
set address v1-dmz servA 1.2.2.10/32
set address v1-dmz servB 1.2.2.20/32
set address v1-dmz servC 1.2.2.30/32
```

2. IP Spoof Protection

```
set zone v1-untrust screen ip-spoofing
save
```

IP Source Route Options

Source routing was designed to allow the user at the source of an IP packet transmission to specify the IP addresses of the routers (also referred to as “hops”) along the path that he or she wants an IP packet to take on its way to its destination. The original intent of the IP source route options was to provide routing control tools to aid diagnostic analysis. If, for example, the transmission of a packet to a particular destination meets with irregular success, you might first use either the record route or timestamp IP option to discover the addresses of routers along the path or paths that the packet takes. You can then use either the loose or strict source route option to direct traffic along a specific path, using the addresses you learned from the results that the record route or timestamp options produced. By changing router addresses to alter the path and sending several packets along different paths, you can note changes that either improve or lessen the success rate. Through analysis and the process of elimination, you might be able to deduce where the trouble lies.

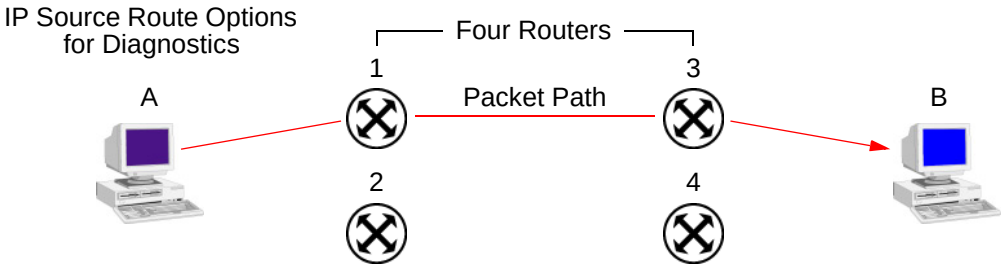

Transmission from A to B is successful 50% of the time using routers 1 and 3.

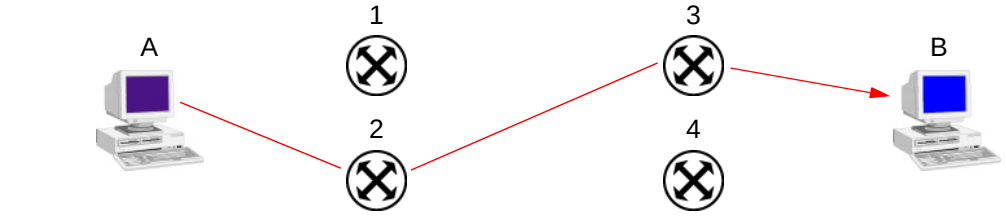

Using IP source routing, A sends traffic through routers 2 and 3. Transmission from A to B remains successful only 50% of the time.

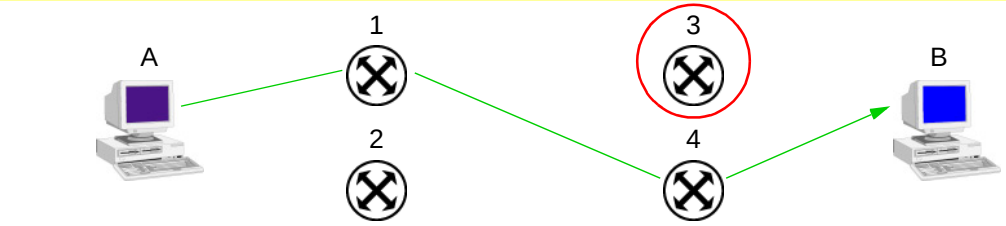

Using IP source routing, A sends traffic through routers 1 and 4. Transmission from A to B is successful 100% of the time. Therefore, we can conclude that the trouble lies in router #3.

Although the uses of IP source route options were originally benign, attackers have learned to put them to more devious uses. They can use IP source route options to hide their true address and access restricted areas of a network by specifying a different path. For an example showing how an attacker can put both deceptions to use, consider the following scenario.

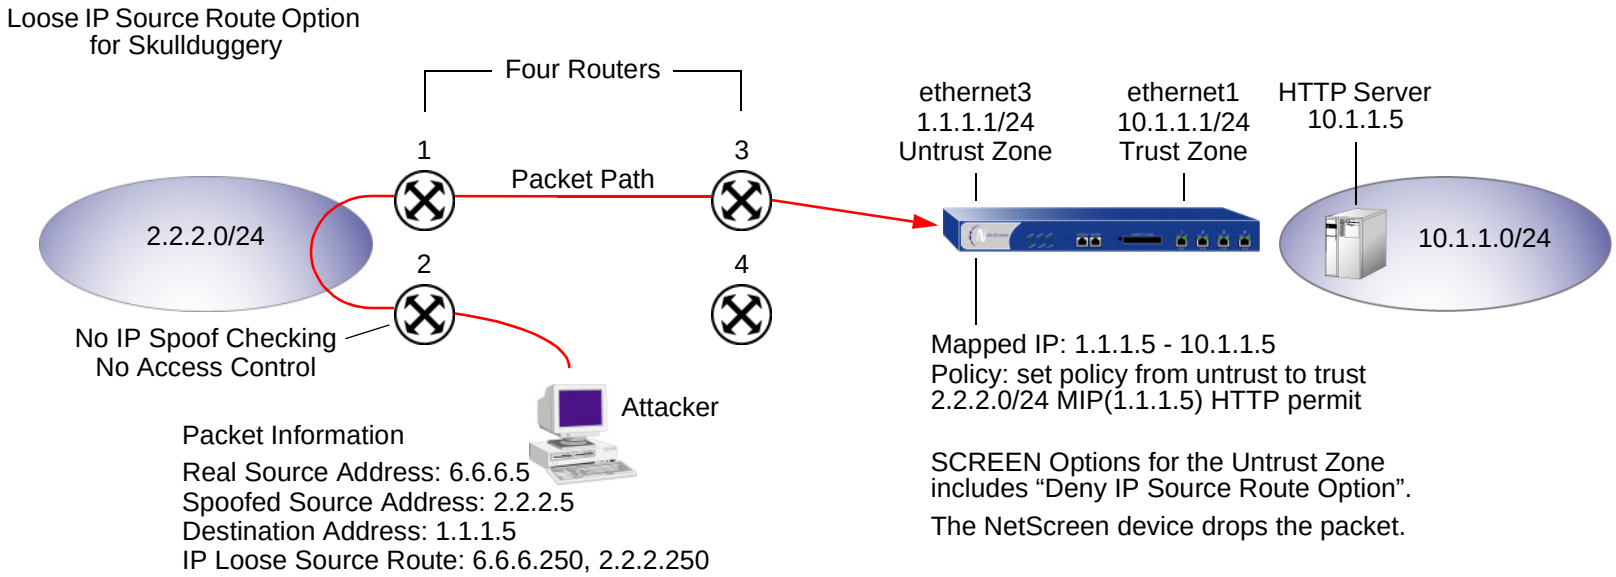

The NetScreen firewall only allows traffic 2.2.2.0/24 if it comes through ethernet1, an interface bound to the Untrust zone. Routers 3 and 4 enforce access controls but routers 1 and 2 do not. Furthermore, router 2 does not check for IP spoofing. The attacker spoofs the source address, and by using the loose source route option, directs the packet through router 2 to the 2.2.2.0/24 network and from there out router 1. Router 1 forwards it to router 3, which forwards it to the NetScreen device. Because the packet came from the 2.2.2.0/24 subnet and has a source address from that subnet, it seems to be valid. However, one remnant of the earlier chicanery remains: the loose source route option. In this example, you have enabled the “Deny IP Source Route Option” SCREEN option for the Untrust zone. When the packet arrives at ethernet3, the NetScreen device rejects it.

You can enable the NetScreen device to either block any packets with loose or strict source route options set or detect such packets and then record the event in the counters list for the ingress interface. The SCREEN options are as follows:

- **Deny IP Source Route Option:** Enable this option to block all IP traffic that employs the loose or strict source route option. Source route options can allow an attacker to enter a network with a false IP address.
- **Detect IP Loose Source Route Option:** The NetScreen device detects packets where the IP option is 3 (Loose Source Routing) and records the event in the SCREEN counters list for the ingress interface. This option specifies a partial route list for a packet to take on its journey from source to destination. The packet must proceed in the order of addresses specified, but it is allowed to pass through other routers in between those specified.
- **Detect IP Strict Source Route Option:** The NetScreen device detects packets where the IP option is 9 (Strict Source Routing) and records the event in the SCREEN counters list for the ingress interface. This option specifies the complete route list for a packet to take on its journey from source to destination. The last address in the list replaces the address in the destination field.

  (For more information about all the IP options, see “Network Reconnaissance Using IP Options” on page 12.)

To block packets with either a loose or strict source route option set, do either of the following, where the specified security zone is the one from which the packets originate:

WebUI

Screening > Screen (Zone: select a zone name): Select **IP Source Route Option Filter**, and then click **Apply**.

CLI

```
set zone zone screen ip-filter-src
```

To detect and record (but not block) packets with a loose or strict source route option set, do either of the following, where the specified security zone is the one from which the packets originate:

WebUI

Screening > Screen (Zone: select a zone name): Enter the following, and then click **Apply**:

IP Loose Source Route Option Detection: (select)

IP Strict Source Route Option Detection: (select)

CLI

```
set zone zone screen ip-loose-src-route
set zone zone screen ip-strict-src-route
```

3

Denial-of-Service Attack Defenses

The intent of a denial-of-service (DoS) attack is to overwhelm the targeted victim with a tremendous amount of bogus traffic so that the victim becomes so preoccupied processing the bogus traffic that it is unable to process legitimate traffic. The target can be the NetScreen firewall, the network resources to which the firewall controls access, or the specific hardware platform or operating system (OS) of an individual host.

If a DoS attack originates from multiple source addresses, it is known as a distributed denial-of-service (DDoS) attack. Typically, the source address of a DoS attack is spoofed. The source addresses in a DDoS attack might be spoofed or the actual addresses of hosts that the attacker has previously compromised and which he or she is now using as "zombie agents" from which to launch the attack.

The NetScreen device can defend itself and the resources it protects from DoS and DDoS attacks. The following sections describe the various defense options available:

- "Firewall DoS Attacks" on page 36
  - "Session Table Flood" on page 36
  - "SYN-ACK-ACK Proxy Flood" on page 43
- "Network DoS Attacks" on page 45
  - "SYN Flood" on page 45
  - "ICMP Flood" on page 59
  - "UDP Flood" on page 61
  - "Land Attack" on page 63
- "OS-Specific DoS Attacks" on page 65
  - "Ping of Death" on page 65
  - "Teardrop Attack" on page 67
  - "WinNuke" on page 69

FIREWALL DOS ATTACKS

If an attacker discovers the presence of the NetScreen firewall, he or she might launch a denial-of-service (DoS) attack against it instead of the network behind it. A successful DoS attack against a firewall amounts to a successful DoS attack against the protected network in that it thwarts attempts of legitimate traffic to traverse the firewall. This section explains two methods that an attacker might use to fill up the session table of a NetScreen device and thereby produce a DoS: "Session Table Flood" on page 36 and "SYN-ACK-ACK Proxy Flood" on page 43

Session Table Flood

A successful DoS attack overwhelms its victim with such a massive barrage of ersatz traffic that it becomes unable to process legitimate connection requests. DoS attacks can take many forms—SYN flood, SYN-ACK-ACK flood, UDP flood, ICMP flood, and so on—but they all seek the same objective: to fill up their victim's session table. When the session table is full, that host cannot create any new sessions and begins rejecting new connection requests. The following SCREEN options help mitigate such attacks:

- "Source- and Destination-Based Session Limits"
- "Aggressive Aging" on page 40

Source- and Destination-Based Session Limits

In addition to limiting the number of concurrent sessions from the same source IP address, you can also limit the number of concurrent sessions to the same destination IP address. One benefit of setting a source-based session limit is that it can stem an attack such as the Nimda virus (which is actually both a virus and a worm) that infects a server and then begins generating massive amounts of traffic from that server. Because all the virus-generated traffic originates from the same IP address, a source-based session limit ensures that the NetScreen firewall can curb such excessive amounts of traffic.

Source-Based Session Limiting: Nimda Virus/Worm Traffic Containment

A Web server is infected with the Nimda virus/worm hybrid, which causes the server to generate excessive amounts of traffic.

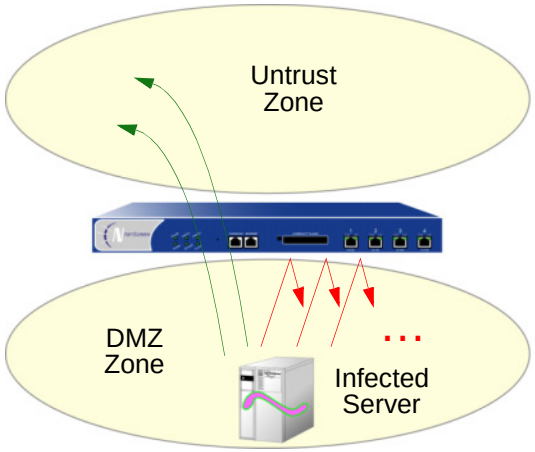

After the number of concurrent sessions from the infected server reaches the maximum limit, the NetScreen device begins blocking all further connection attempts from that server.

Another benefit of source-based session limiting is that it can mitigate attempts to fill up the NetScreen session table—if all the connection attempts originate from the same source IP address. However, a wily attacker can launch a distributed denial-of-service (DDoS) attack. In a DDoS attack, the malicious traffic can come from hundreds of hosts, known as “zombie agents”, that are surreptitiously under the control of an attacker. In addition to the SYN, UDP, and ICMP flood detection and prevention SCREEN options, setting a destination-based session limit can ensure that the NetScreen device allows only an acceptable number of concurrent connection requests—no matter what the source—to reach any one host.

Source-Based Session Limiting:
Denial-of-Service Attack

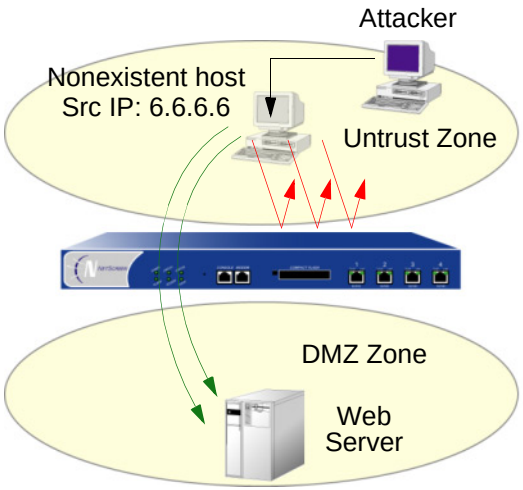

When the number of concurrent sessions from 6.6.6.6 surpasses the maximum limit, the NetScreen device begins blocking further connection attempts from that IP address.

Destination-Based Session Limiting:
Distributed Denial-of-Service Attack

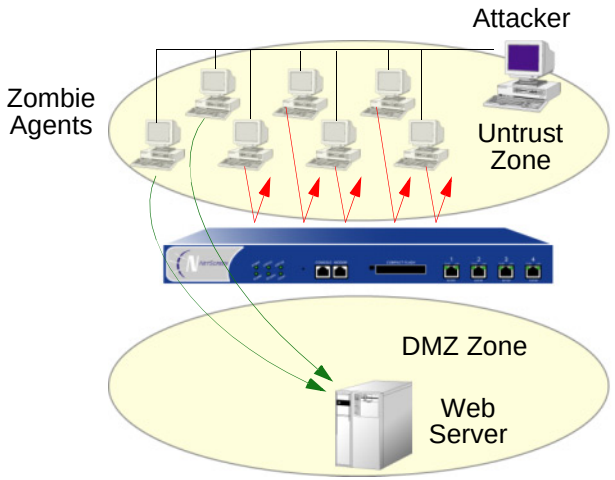

When the number of concurrent sessions to the Web server surpasses the maximum limit, the NetScreen device begins blocking further connection attempts to that IP address.

Determining what constitutes an acceptable number of connection requests requires a period of observation and analysis to establish a baseline for typical traffic flows. You also need to consider the maximum number of concurrent sessions required to fill up the session table of the particular NetScreen platform you are using. To see the maximum number of sessions that your session table supports, use the CLI command **get session**, and then look at the first line in the output, which lists the number of current (allocated) sessions, the maximum number of sessions, and the number of failed session allocations:

```
alloc 420/max 128000, alloc failed 0
```

The default maximum for both source- and destination-based session limits is 128 concurrent sessions, a value that might need adjustment to suit the needs of your network environment and the platform.

Example: Source-Based Session Limiting

In this example, you want to limit the amount of sessions that any one server in the DMZ and Trust zones can initiate. Because the DMZ zone only contains Web servers, none of which should initiate traffic, you set the source-session limit at the lowest possible value: 1 session. On the other hand, the Trust zone contains personal computers, servers, printers, and so on, many of which do initiate traffic. For the Trust zone, you set the source-session limit maximum to 80 concurrent sessions.

WebUI

Screening > Screen (Zone: DMZ): Enter the following, and then click **OK**:

Source IP Based Session Limit: (select)

Threshold: 1 Sessions

Screening > Screen (Zone: Trust): Enter the following, and then click **OK**:

Source IP Based Session Limit: (select)

Threshold: 80 Sessions

CLI

```
set zone dmz screen limit-session source-ip-based 1
set zone dmz screen limit-session source-ip-based
set zone trust screen limit-session source-ip-based 80
set zone trust screen limit-session source-ip-based
save
```

Example: Destination-Based Session Limiting

In this example, you want to limit the amount of traffic to a Web server at 1.2.2.5. The server is in the DMZ zone. After observing the traffic flow from the Untrust zone to this server for a month, you have determined that the average number of concurrent sessions it receives is 2000. Based on this information, you decide to set the new session limit at 4000 concurrent sessions. Although your observations show that traffic spikes sometimes exceed that limit, you opt for firewall security over occasional server inaccessibility.

WebUI

Screening > Screen (Zone: Untrust): Enter the following, and then click **OK**:

Destination IP Based Session Limit: (select)

Threshold: 4000 Sessions

CLI

```
set zone untrust screen limit-session destination-ip-based 4000
set zone untrust screen limit-session destination-ip-based
save
```

Aggressive Aging

By default, an initial TCP session 3-way handshake takes 20 seconds to time out (that is, to expire because of inactivity). After a TCP session has been established, the timeout value changes to 30 minutes. For HTTP and UDP sessions, the session timeouts are 5 minutes and 1 minute respectively. The session timeout counter begin when a session starts and is refreshed every 10 seconds if the session is active. If a session becomes idle for more than 10 seconds, the timeout counter begins to decrement.

NetScreen provides a mechanism to accelerate the timeout process when the number of sessions in the session table surpasses a specified high-watermark threshold. When the number of sessions dips below a specified low-watermark threshold, the timeout process returns to normal. During the period when the aggressive aging out process is in effect, a NetScreen device ages out the oldest sessions first, using the aging out rate that you specify. These aged-out sessions are tagged as invalid and are removed in the next “garbage sweep”, which occurs every 2 seconds.

The aggressive ageout option shortens default session timeouts by the amount you enter[1]. The aggressive ageout value can be between 2 and 10 units, where each unit represents a 10-second interval (that is, the aggressive ageout setting can be between 20 and 100 seconds). The default setting is 2 units, or 20 seconds. If you define the aggressive ageout setting at 100 seconds, for example, you shorten the TCP and HTTP session timeouts as follows:

- TCP: The session timeout value shortens from 1800 seconds (30 minutes) to 1700 seconds (28:20 minutes) during the time when the aggressive aging process is in effect. During that period, the NetScreen device automatically deletes all TCP sessions whose timeout value has passed 1700 seconds, beginning with the oldest sessions first.

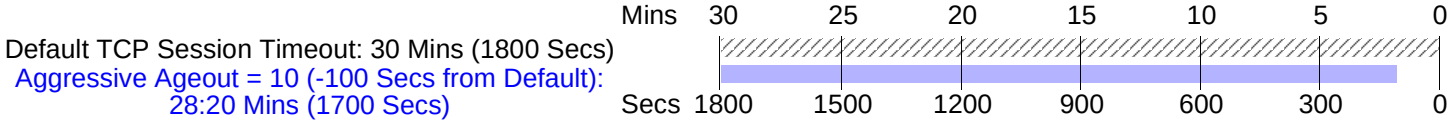

- HTTP: The session timeout value shortens from 300 seconds (5 minutes) to 200 seconds (3:20 minutes) during the time when the aggressive aging process is in effect. During that period, the NetScreen device automatically deletes all HTTP sessions whose timeout value has passed 200 seconds, beginning with the oldest sessions first.

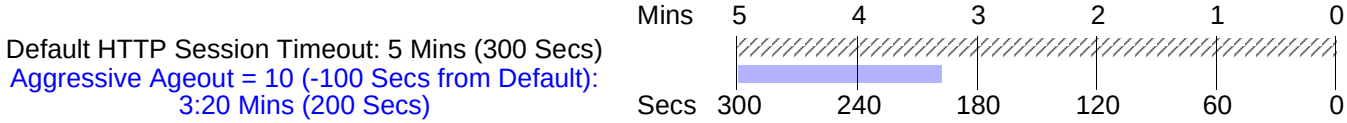

- UDP: Because the default UDP session timeout is 60 seconds, defining an early ageout setting at 100 seconds causes all UDP sessions to ageout and be marked for deletion in the next garbage sweep.

1. When you set and enable the aggressive ageout option, the normal session timeout value displayed in the configuration remains unchanged—1800 seconds for TCP, 300 seconds for HTTP, and 60 seconds for UDP sessions. However, when the aggressive ageout period is in effect, these sessions time out earlier—by the amount you specify for early ageout—instead of counting down all the way to zero.

Example: Aggressively Aging Out Sessions

In this example, you set the aggressive aging out process to commence when traffic exceeds a high-watermark of 80% and cease when it retreats below a low-watermark of 70%. You specify 40 seconds for the aggressive age-out interval. When the session table is more than 80% full (the high-mark threshold), the NetScreen device decreases the timeout for all sessions by 40 seconds and begins aggressively aging out the oldest sessions until the number of sessions in the table is under 70% (the low-mark threshold).

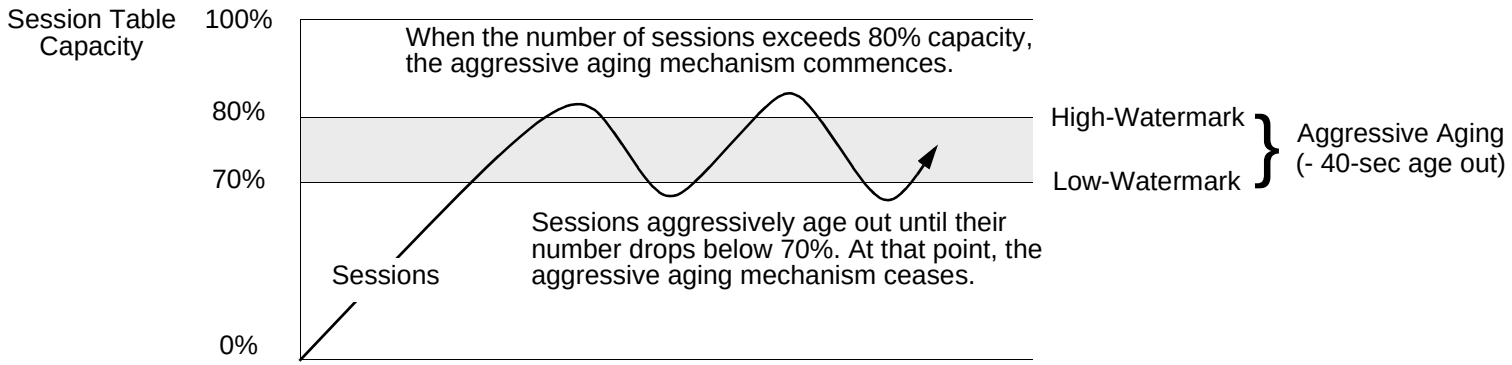

WebUI

> ***Note:*** *You must use the CLI to configure the aggressive age -out settings.*

CLI

```
set flow aging low-watermark 70
set flow aging high-watermark 80
set flow aging early-ageout 4
save
```

SYN-ACK-ACK Proxy Flood

When an authentication user initiates a Telnet or FTP connection, the user sends a SYN segment to the Telnet or FTP server. The NetScreen device intercepts the SYN segment, creates an entry in its session table, and proxies a SYN-ACK segment to the user. The user then replies with an ACK segment. At that point, the initial 3-way handshake is complete. The NetScreen device sends a login prompt to the user. If the user, with malicious intent, does not log in, but instead continues initiating SYN-ACK-ACK sessions, the NetScreen session table can fill up to the point where the device begins rejecting legitimate connection requests.

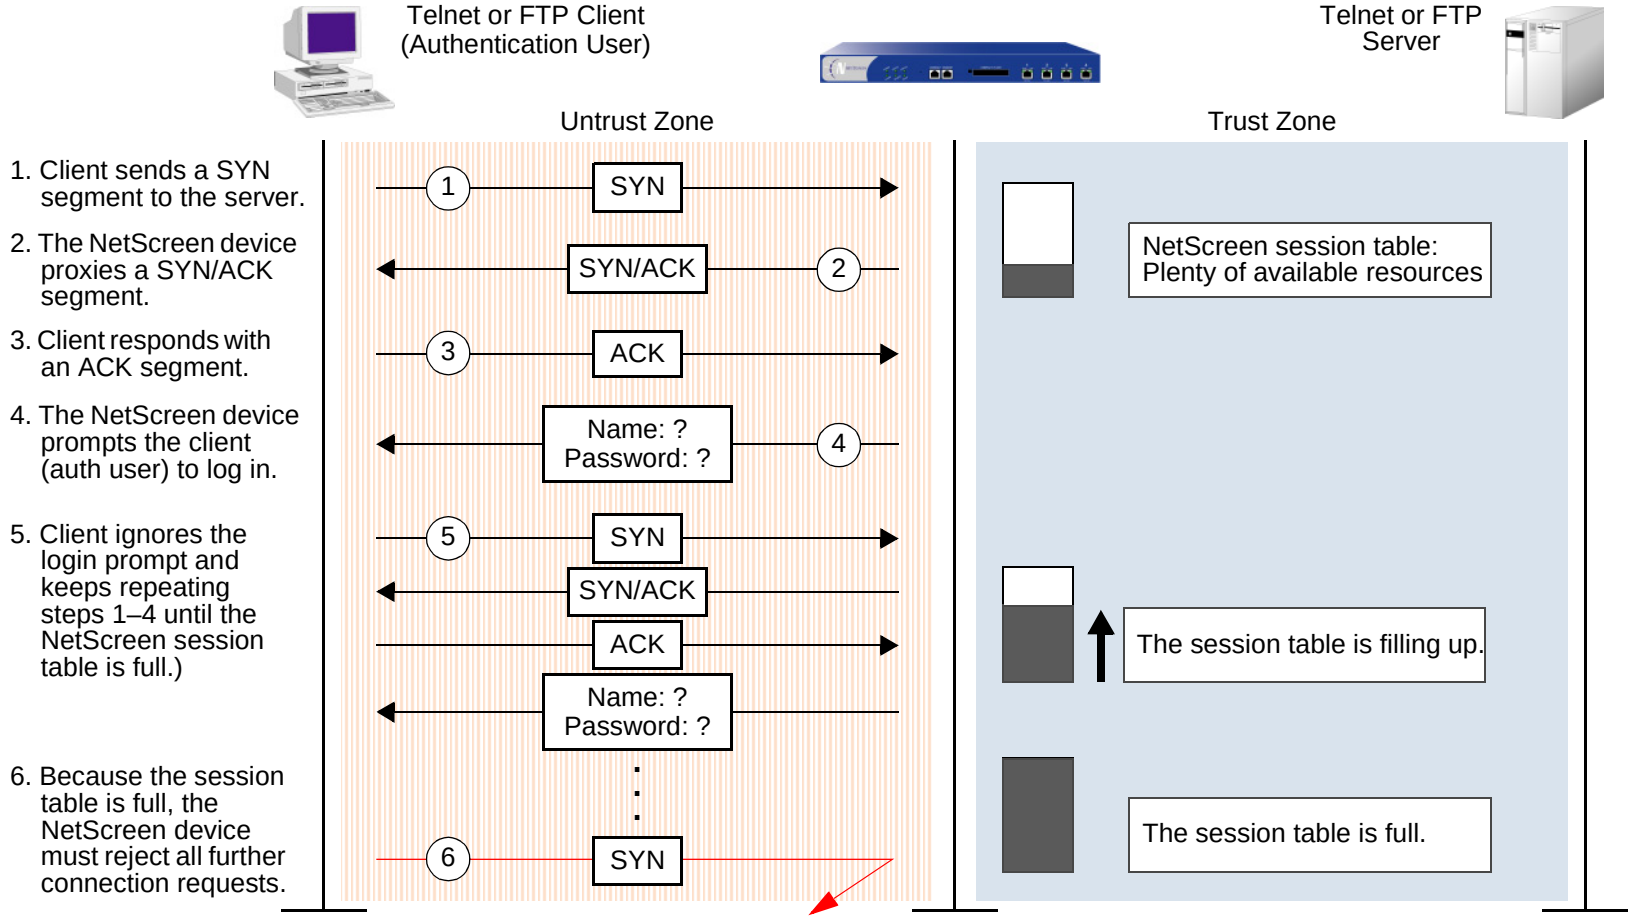

To thwart such an attack, you can enable the SYN-ACK-ACK proxy protection SCREEN option. After the number of connections from the same IP address reaches the SYN-ACK-ACK proxy threshold, the NetScreen device rejects further connection requests from that IP address. By default, the threshold is 512 connections from any single IP address. You can change this threshold (to any number between 1 and 2,500,000) to better suit the requirements of your network environment.

To enable protection against a SYN-ACK-ACK proxy flood, do either of the following, where the specified zone is that in which the attack originates:

WebUI

Screening > Screen (Zone: select a zone name): Enter the following, and then click **Apply**:

SYN-ACK-ACK Proxy Protection: (select)

Threshold: (enter a value to trigger SYN-ACK-ACK proxy flood protection[2])

CLI

```
set zone zone screen syn-ack-ack-proxy threshold number
set zone zone screen syn-ack-ack-proxy
```

2. The value unit is connections per source address. The default value is 512 connections from any single address.

Network DoS Attacks

A denial-of-service (DoS) attack directed against one or more network resources floods the target with an overwhelming number of SYN, ICMP, or UDP packets, or with an overwhelming number of SYN fragments. Depending on the attacker's purpose and the extent and success of previous intelligence gathering efforts, the attacker might single out a specific host, such as a router or server; or he or she might aim at random hosts across the targeted network. Either approach has the potential of upsetting service to a single host or to the entire network, depending on how critical the role of the victim is to the rest of the network.

SYN Flood

A SYN flood occurs when a host becomes so overwhelmed by SYN segments initiating uncompletable connection requests that it can no longer process legitimate connection requests.

Two hosts establish a TCP connection with a triple exchange of TCP segments known as a three-way handshake: A sends a SYN segment to B; B responds with a SYN/ACK segment; and A responds with an ACK segment. A SYN flood attack inundates a site with SYN segments containing forged ("spoofed") IP source addresses with nonexistent or unreachable addresses. B responds with SYN/ACK segments to these addresses and then waits for responding ACK segments. Because the SYN/ACK segments are sent to nonexistent or unreachable IP addresses, they never elicit responses and eventually time out.

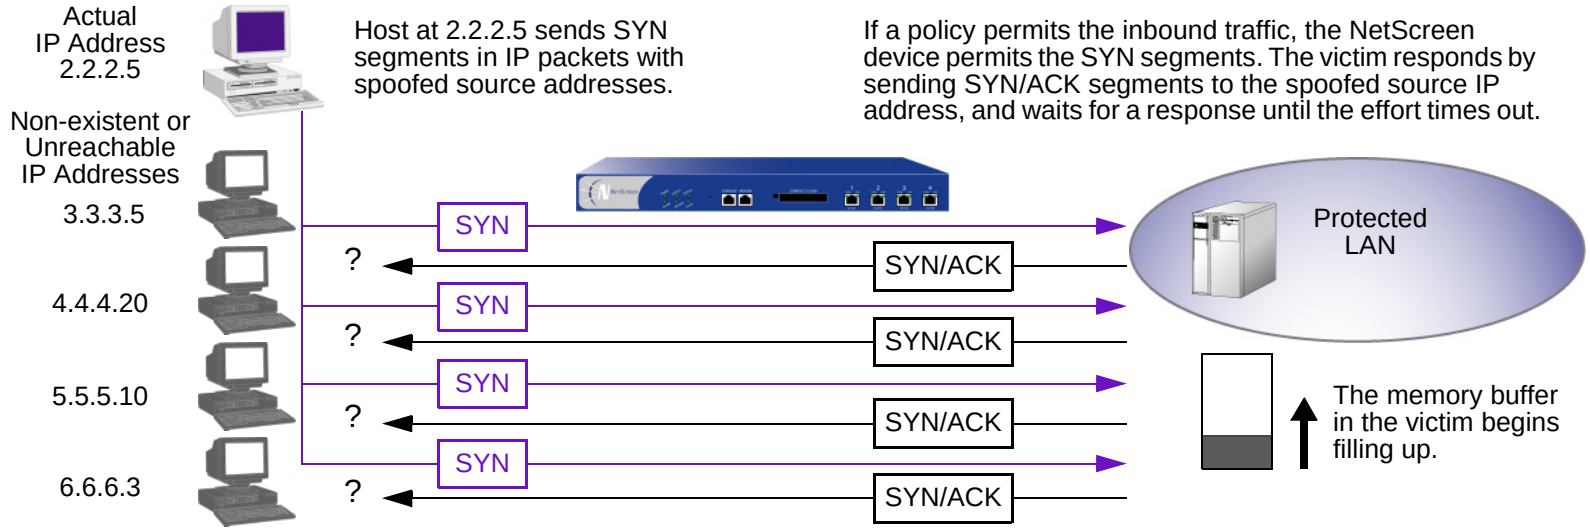

By flooding a host with uncompletable TCP connections, the attacker eventually fills the memory buffer of the victim. Once this buffer is full, the host can no longer process new TCP connection requests. The flood might even damage the victim's operating system. Either way, the attack disables the victim and its normal operations.

SYN Flood Protection

NetScreen devices can impose a limit on the number of SYN segments permitted to pass through the firewall per second. You can base the attack threshold on the destination address and port, the destination address only, or the source address only. When the number of SYN segments per second exceeds one of these thresholds, the NetScreen device starts proxying incoming SYN segments, replying with SYN/ACK segments and storing the incomplete connection requests in a connection queue. The incomplete connection requests remain in the queue until the connection is completed or the request times out. In the following illustration, the SYN attack threshold has been passed and the NetScreen device has started proxying SYN segments.

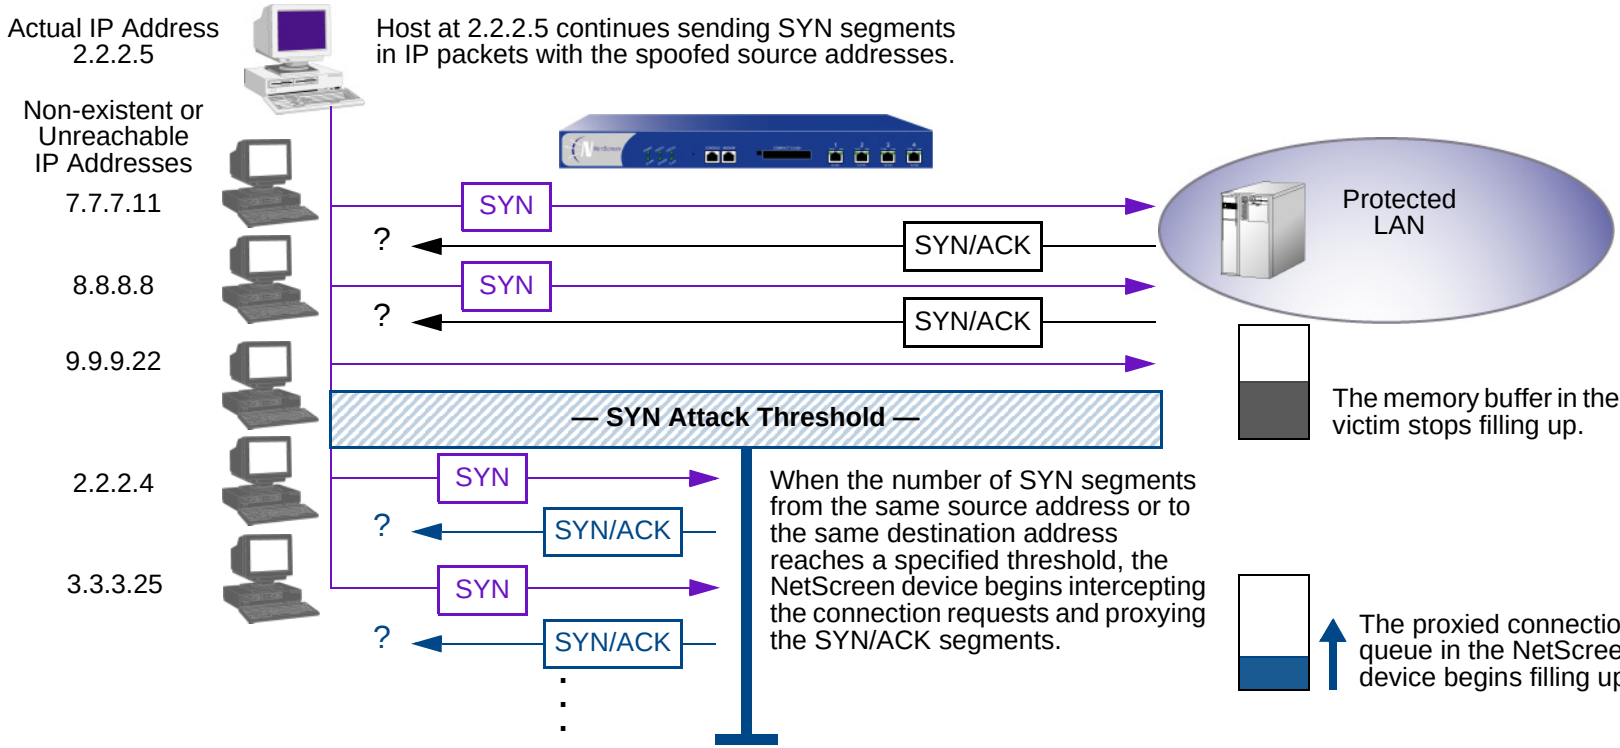

In the next illustration, the proxied connection queue has completely filled up, and the NetScreen device is rejecting new incoming SYN segments. This action shields hosts on the protected network from the bombardment of incomplete three-way handshakes.

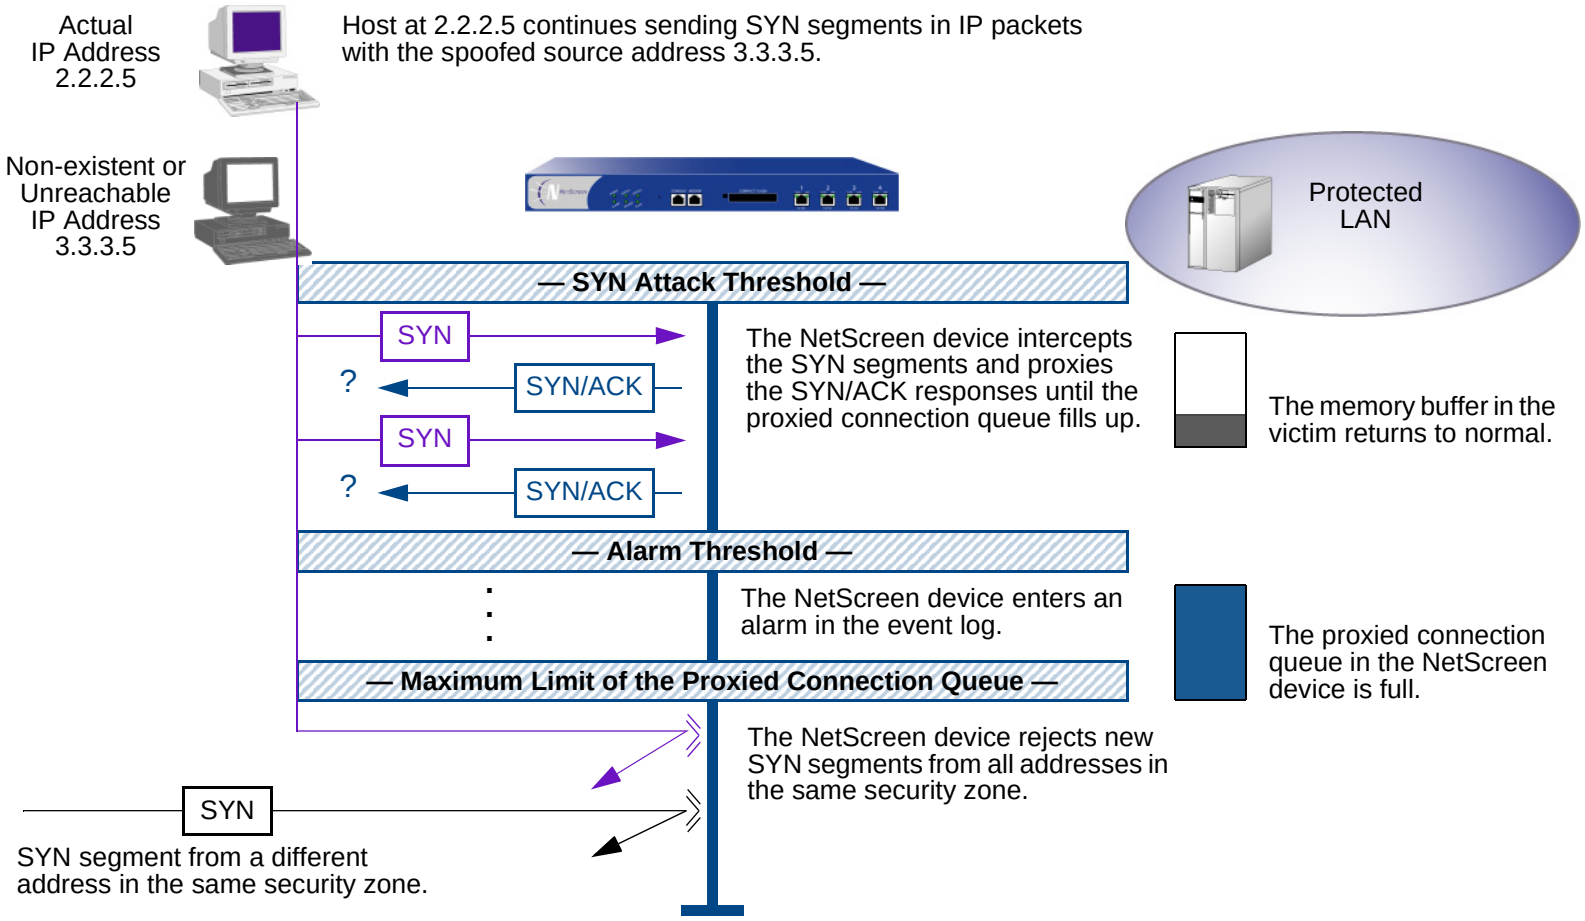

The NetScreen device starts receiving new SYN packets when the proxy queue drops below the maximum limit.

Note: *The procedure of proxying incomplete SYN connections above a set threshold pertains only to traffic permitted by existing policies. Any traffic for which a policy does not exist is automatically dropped.*

To enable the SYN flood protection SCREEN option and define its parameters, do either of the following, where the specified zone is that in which a flood might originate:

WebUI

Screening > Screen (Zone: select a zone name): Enter the following, and then click **Apply**:

SYN Flood Protection: (select to enable)

Threshold: (enter the number of SYN packets—that is, TCP segments with the SYN flag set—per second required to activate the SYN proxying mechanism[3])

Alarm Threshold: (enter the number of proxied TCP connection requests required to write an alarm in the event log)

Source Threshold: (enter the number SYN packets per second from a single IP address required to activate SYN proxying)

Destination Threshold: (enter the number SYN packets per second to a single IP address required to activate SYN proxying)

Timeout Value: (enter the length of time in seconds that the NetScreen device holds an incomplete TCP connection attempt in the proxied connection queue)

Queue Size: (enter the number of proxied TCP connection requests held in the proxied connection queue before the NetScreen device starts rejecting new connection requests)

3. For more details on each of these parameters, see the descriptions in the following CLI section.

CLI

To enable SYN flood protection.

```
set zone zone screen syn-flood
```

You can set the following parameters for proxying uncompleted TCP connection requests:

Attack Threshold: The number of SYN segments (that is, TCP segments with the SYN flag set) to the same destination address and port number per second required to activate the SYN proxying mechanism. Although you can set the threshold at any number, you need to know the normal traffic patterns at your site to set an appropriate threshold for it. For example, if it is an e-business site that normally gets 20,000 SYN segments per second, you might want to set the threshold at 30,000/second. If a smaller site normally gets 20 SYN segments/second, you might consider setting the threshold at 40.

```
set zone zone screen syn-flood attack-threshold number
```

Alarm Threshold: The number of proxied, half-complete TCP connection requests per second after which the NetScreen device enters an alarm in the event log. The value you set for an alarm threshold triggers an alarm when the number of proxied, half-completed connection requests to the same destination address and port number per second exceeds that value. For example, if you set the SYN attack threshold at 2000 SYN segments per second and the alarm at 1000, then a total of 3001 SYN segments to the same destination address and port number per second is required to trigger an alarm entry in the log. More precisely:

1. The firewall passes the first 2000 SYN segments per second that meet policy requirements.
2. The firewall proxies the next 1000 SYN segments in the same second.
3. The 1001st proxied connection request (or 3001st connection request in that second) triggers the alarm.

```
set zone zone screen syn-flood alarm-threshold number
```

For each SYN segment to the same destination address and port number in excess of the alarm threshold, the attack detection module generates a message. At the end of the second, the logging module compresses all similar messages into a single log entry that indicates how many SYN segments to the same destination address and port number arrived after exceeding the alarm threshold. If the attack persists beyond the first second, the event log enters an alarm every second until the attack stops.

Source Threshold: This option allows you to specify the number of SYN segments received per second from a single source IP address—regardless of the destination IP address and port number—before the NetScreen device executes the SYN proxying mechanism.

```
set zone zone screen syn-flood source-threshold number
```

Tracking a SYN flood by source address uses different detection parameters from tracking a SYN flood by destination address and destination port number. When you set a SYN attack threshold and a source threshold, you put both the basic SYN flood protection mechanism and the source-based SYN flood tracking mechanism in effect.

Destination Threshold: This option allows you to specify the number of SYN segments received per second for a single destination IP address before the NetScreen device executes the SYN proxying mechanism. If a protected host runs multiple services, you might want to set a threshold based on destination IP address only—regardless of the destination port number.

```
set zone zone screen syn-flood destination-threshold number
```

When you set a SYN attack threshold and a destination threshold, you put both the basic SYN flood protection mechanism and the destination-based SYN flood tracking mechanism in effect.

Tracking a SYN flood by destination address uses different detection parameters from tracking a SYN flood by destination address and destination port number. Consider the following case where the NetScreen device has policies permitting FTP requests (port 21) and HTTP requests (port 80) to the same server. If the SYN flood attack threshold is 1000 packets per second (pps) and an attacker sends 999 FTP packets and 999 HTTP packets per second, neither set of packets (where a set is defined as having the same destination address and port number) activates the SYN proxying mechanism. The basic SYN flood attack mechanism tracks destination address and port number, and neither set exceeds the attack threshold of 1000 pps. However, if the destination threshold is 1000 pps, the NetScreen device treats both FTP and HTTP packets with the same destination address as members of a single set. The 1001st packet—FTP or HTTP—to the same destination address triggers the SYN proxying mechanism.

Timeout: The maximum length of time before a half-completed connection is dropped from the queue. The default is 20 seconds, and you can set the timeout from 0–50 seconds. You might try decreasing the timeout value to a shorter length until you begin to see any dropped connections during normal traffic conditions. Twenty seconds is a very conservative timeout for a threeway-handshake ACK response.

```
set zone zone screen syn-flood timeout number
```

Queue size: The number of proxied connection requests held in the proxied connection queue before the NetScreen device starts rejecting new connection requests. The longer the queue size, the longer the NetScreen device needs to scan the queue to match a valid ACK response to a proxied connection request. This can slightly slow the initial connection establishment; however, because the time to begin data transfer is normally far greater than any minor delays in initial connection setup, users would not see a noticeable difference.

```
set zone zone screen syn-flood queue-size number
```

Drop Unknown MAC: When a NetScreen device detects a SYN attack, it proxies all TCP connection requests. However, a NetScreen device in Transparent mode cannot proxy a TCP connection request if the destination MAC address is not in its MAC learning table. By default, a NetScreen device in Transparent mode that has detected a SYN attack passes SYN packets containing unknown MAC addresses. You can use this option to instruct the device to drop SYN packets containing unknown destination MAC addresses instead of letting them pass.

```
set zone zone screen syn-flood drop-unknown-mac
```

Example: SYN Flood Protection

In this example, you protect four Web servers in the DMZ zone from SYN flood attacks originating in the Untrust zone by enabling the SYN flood protection SCREEN option for the Untrust zone.

Note: *NetScreen recommends that you augment the SYN flood protection that the NetScreen device provides with device-level SYN flood protection on each of the Web servers. In this example, the Web servers are running UNIX, which also provides some SYN flood defenses, such as adjusting the length of the connection request queue and changing the timeout period for incomplete connection requests.*

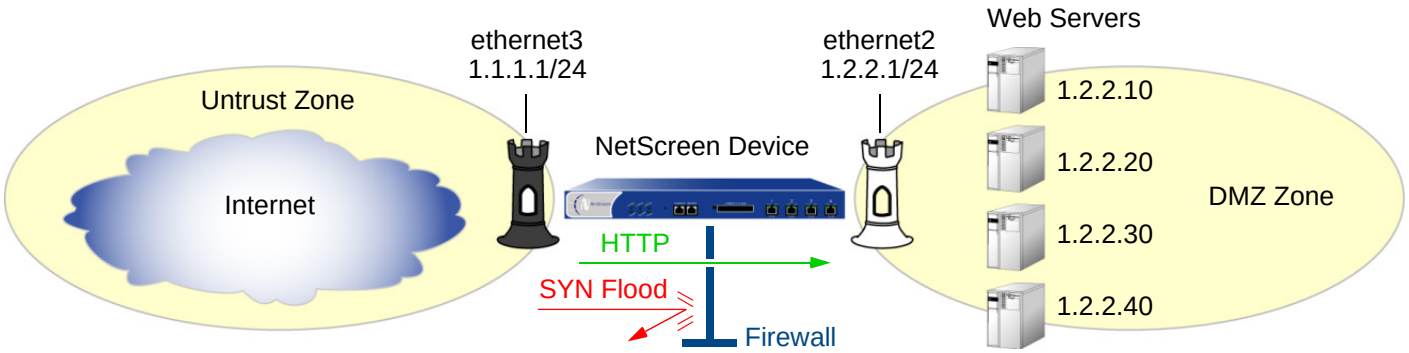

To configure the SYN flood protection parameters with appropriate values for your network, you must first establish a baseline of typical traffic flows. For one week, you run a sniffer[4] on ethernet3—the interface bound to the Untrust zone—to monitor the number of new TCP connection requests arriving every second for the four Web servers in the DMZ[5]. Your analysis of the data accumulated from one week of monitoring produces the following statistics:

- Average number of new connection requests per server: 250/second
- Average peak number of new connection requests per server: 500/second

4. A sniffer is a network analyzing device that captures packets on the network segment to which you attach it. Most sniffers allow you to define filters to collect only the type of traffic that interests you. Later, you can view and evaluate the accumulated information. In this example, you want the sniffer to collect all TCP packets with the SYN flag set arriving at ethernet3 and destined for one of the four Web servers in the DMZ.

5. You might want to continue running the sniffer at regular intervals to see if there are traffic patterns based on the time of day, days of the week, the time of month, or the season of the year. For example, traffic might increase dramatically during the Christmas season. Significant changes probably warrant adjusting the various thresholds.

Based on this information, you set the following SYN flood protection parameters for the Untrust zone:

Parameters	Values	Reason for Each Value
Attack Threshold	625 packets per second (pps)	This is 25% higher than the average peak number of new connection requests per second per server, which is unusual for this network environment. When the number of SYN packets per second for any one of the four Web servers exceeds this number, the NetScreen device begins proxying new connection requests to that server. (In other words, beginning with the 626th SYN packet to the same destination address and port number in one second, the NetScreen device begins proxying connection requests to that address and port number.)
Alarm Threshold	250 pps	250 pps is 1/4 of the queue size (1000 proxied, half-completed connection requests*). When the NetScreen device proxies 251 new connection requests in one second, it makes an alarm entry in the event log. By setting the alarm threshold somewhat higher than the attack threshold, you can avoid alarm entries for traffic spikes that only slightly exceed the attack threshold.
Source Threshold	25 pps	When you set a source threshold, the NetScreen device tracks the source IP address of SYN packets, regardless of the destination address and port number. (Note that this source-based tracking is separate from the tracking of SYN packets based on destination address and destination port number that constitutes the basic SYN flood protection mechanism.) In the one week of monitoring activity, you observed that no more than 1/25 of new connection requests for all servers came from any one source within a one-second interval. Therefore, connection requests exceeding this threshold are unusual and provide sufficient cause for the NetScreen device to execute its proxying mechanism. (25 pps is 1/25 of the attack threshold, which is 625 pps.)

Parameters	Values	Reason for Each Value
Destination Threshold	0 pps	When you set a destination threshold, the NetScreen device runs a separate tracking of only the destination IP address, regardless of the destination port number. Because the four Web servers only receive HTTP traffic (destination port 80)—no traffic to any other destination port number reaches them—setting a separate destination threshold offers no additional advantage.
Timeout	20 seconds	Because the queue size is relatively short (1000 proxied connection requests), the default value of 20 seconds is a reasonable length of time to hold incomplete connection requests in the queue for this configuration.
Queue Size	1000 proxied, half-completed connections	1000 proxied, half-completed connection requests is twice the average peak number of new connection requests (500 pps). The NetScreen device proxies up to 1000 requests per second before dropping new requests. Proxying twice the average peak number of new connection requests provides a conservative buffer for legitimate connection requests to get through.

* Half-completed connection requests are incomplete three-way handshakes. A three-way handshake is the initial phase of a TCP connection. It consists of a TCP segment with the SYN flag set, a response with the SYN and ACK flags set, and a response to that with the ACK flag set. For a complete description, see “Glossary” in Volume 1, “Overview”.

WebUI

1. Interfaces

Network > Interfaces > Edit (for ethernet2): Enter the following, and then click **OK**:

Zone Name: DMZ

Static IP: (select this option when present)

IP Address/Netmask: 1.2.2.1/24

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

2. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: ws1

IP Address/Domain Name:

IP/Netmask: (select), 1.2.2.10/32

Zone: DMZ

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: ws2

IP Address/Domain Name:

IP/Netmask: (select), 1.2.2.20/32

Zone: DMZ

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: ws3

IP Address/Domain Name:

IP/Netmask: (select), 1.2.2.30/32

Zone: DMZ

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: ws4

IP Address/Domain Name:

IP/Netmask: (select), 1.2.2.40/32

Zone: DMZ

Objects > Addresses > Group > (for Zone: DMZ) New: Enter the following group name, move the following addresses, and then click **OK**:

Group Name: web_servers

Select **ws1** and use the **<<** button to move the address from the Available Members column to the Group Members column.

Select **ws2** and use the **<<** button to move the address from the Available Members column to the Group Members column.

Select **ws3** and use the **<<** button to move the address from the Available Members column to the Group Members column.

Select **ws4** and use the **<<** button to move the address from the Available Members column to the Group Members column.

3. Policy

Policies > (From: Untrust, To: DMZ) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), web_servers

Service: HTTP

Action: Permit

4. SCREEN

Screening > Screen (Zone: Untrust): Enter the following, and then click **Apply**:

SYN Flood Protection: (select)

Threshold: 625

Alarm Threshold: 250

Source Threshold: 25

Destination Threshold: 0

Timeout Value: 20[6]

Queue Size: 1000

6. Because 20 seconds is the default setting, you do not have to set the timeout to 20 seconds unless you have previously set it to another value.

CLI

1. Interfaces

```
set interface ethernet2 zone dmz
set interface ethernet2 ip 1.2.2.1/24

set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
```

2. Addresses

```
set address dmz ws1 1.2.2.10/32
set address dmz ws2 1.2.2.20/32
set address dmz ws3 1.2.2.30/32
set address dmz ws4 1.2.2.40/32

set group address dmz web_servers add ws1
set group address dmz web_servers add ws2
set group address dmz web_servers add ws3
set group address dmz web_servers add ws4
```

3. Policy

```
set policy from untrust to dmz any web_servers HTTP permit
```

4. SCREEN

```
set zone untrust screen syn-flood attack-threshold 625
set zone untrust screen syn-flood alarm-threshold 250
set zone untrust screen syn-flood source-threshold 25
set zone untrust screen syn-flood timeout 20[7]
set zone untrust screen syn-flood queue-size 1000
set zone untrust screen syn-flood
save
```

7. Because 20 seconds is the default setting, you do not have to set the timeout to 20 seconds unless you have previously set it to another value.

ICMP Flood

An ICMP flood occurs when ICMP echo requests overload its victim with so many requests that it expends all its resources responding until it can no longer process valid network traffic. When enabling the ICMP flood protection feature, you can set a threshold that once exceeded invokes the ICMP flood attack protection feature. (The default threshold value is 1000 packets per second.) If the threshold is exceeded, the NetScreen device ignores further ICMP echo requests for the remainder of that second plus the next second as well.

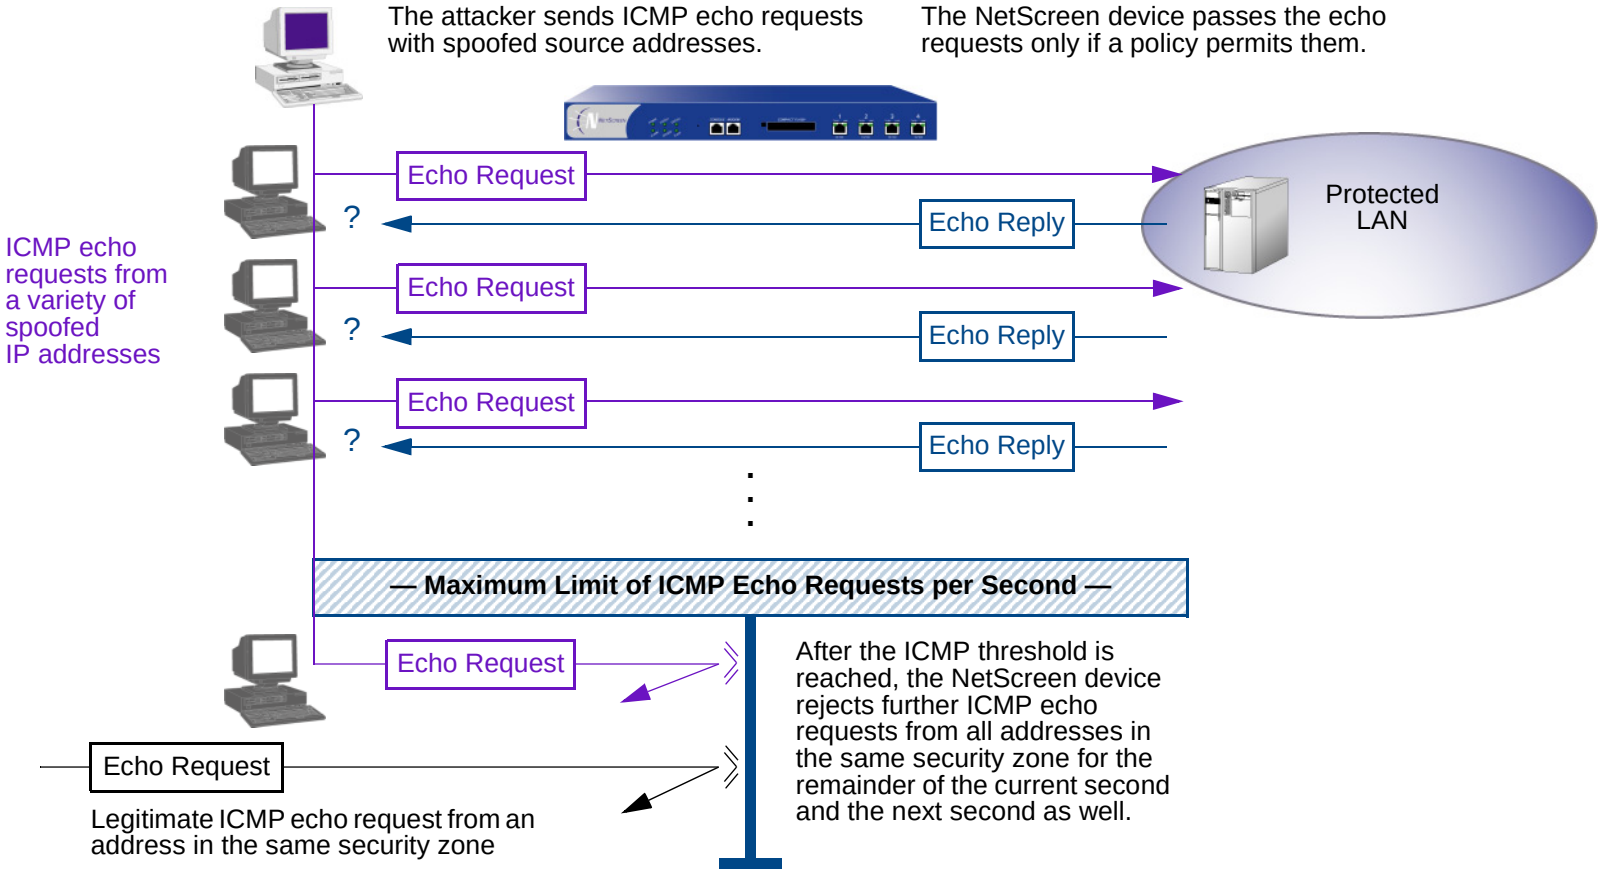

To enable ICMP flood protection, do either of the following, where the specified zone is that in which a flood might originate:

WebUI

Screening > Screen (Zone: select a zone name): Enter the following, and then click **Apply**:

ICMP Flood Protection: (select)

Threshold: (enter a value to trigger ICMP flood protection[8])

CLI

```
set zone zone screen icmp-flood threshold number
set zone zone screen icmp-flood
```

8. The value unit is ICMP packets per second. The default value is 1000 packets per second.

UDP Flood

Similar to the ICMP flood, UDP flooding occurs when an attacker sends IP packets containing UDP datagrams with the purpose of slowing down the victim to the point that it can no longer handle valid connections. After enabling the UDP flood protection feature, you can set a threshold that once exceeded invokes the UDP flood attack protection feature. (The default threshold value is 1000 packets per second.) If the number of UDP datagrams from one or more sources to a single destination exceeds this threshold, the NetScreen device ignores further UDP datagrams to that destination for the remainder of that second plus the next second as well.

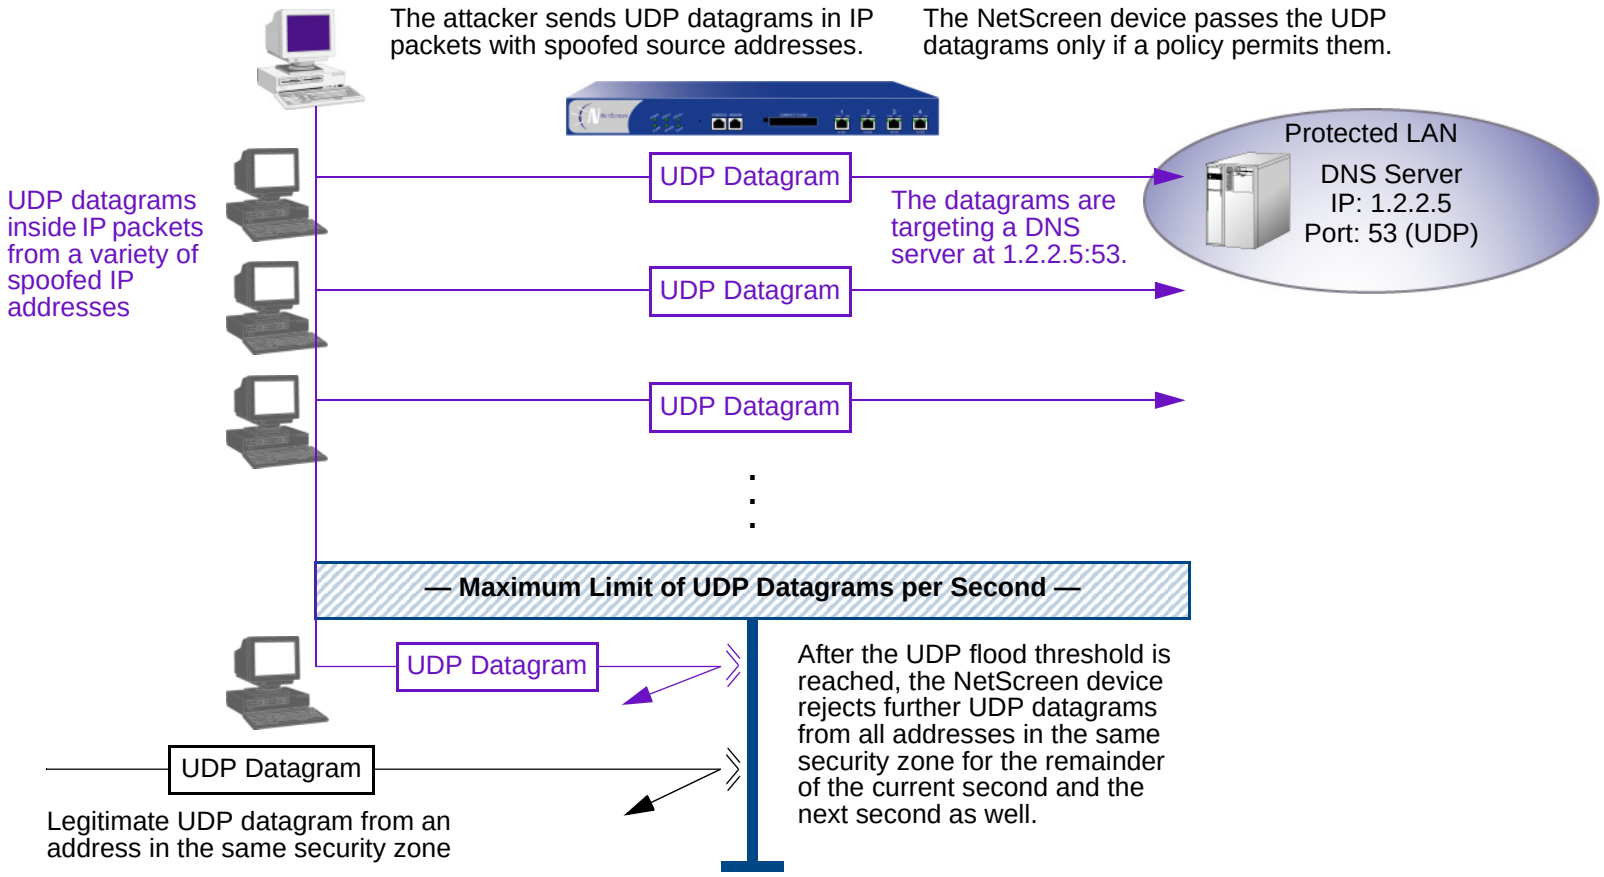

To enable UDP flood protection, do either of the following, where the specified zone is that in which a flood might originate:

WebUI

Screening > Screen (Zone: select a zone name): Enter the following, and then click **Apply**:

UDP Flood Protection: (select)

Threshold: (enter a value to trigger UDP flood protection[9])

CLI

```
set zone zone screen udp-flood threshold number
set zone zone screen udp-flood
```

9. The value unit is UDP packets per second. The default value is 1000 packets per second.

Land Attack

Combining a SYN attack with IP spoofing, a Land attack occurs when an attacker sends spoofed SYN packets containing the IP address of the victim as both the destination and source IP address. The receiving system responds by sending the SYN-ACK packet to itself, creating an empty connection that lasts until the idle timeout value is reached. Flooding a system with such empty connections can overwhelm the system, causing a DoS.

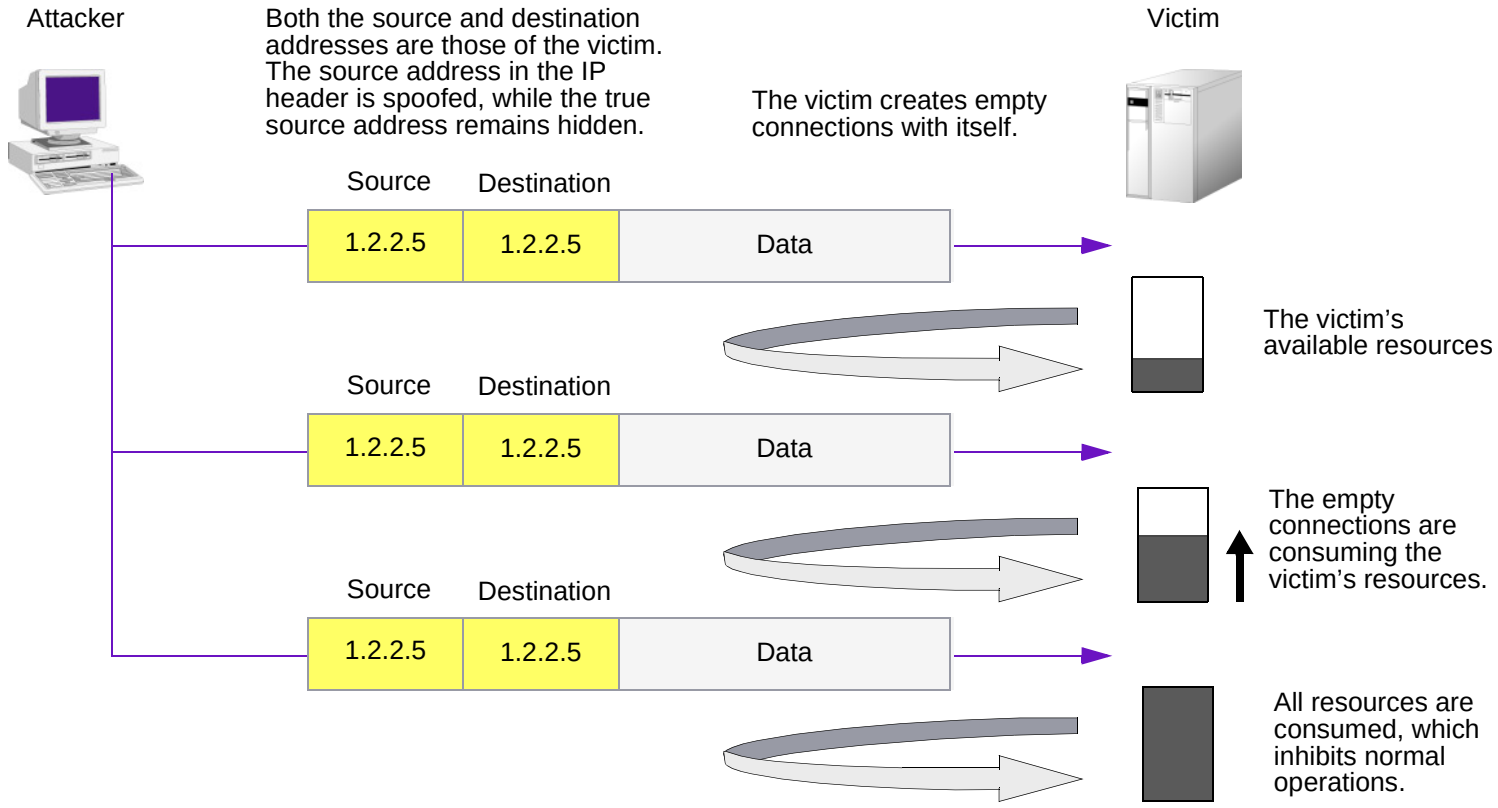

When you enable the SCREEN option to block Land attacks, the NetScreen device combines elements of the SYN flood defense and IP spoofing protection to detect and block any attempts of this nature.

To enable protection against a Land attack, do either of the following, where the specified zone is that in which the attack originates:

WebUI

Screening > Screen (Zone: select a zone name): Select **Land Attack Protection**, and then click **Apply**.

CLI

```
set zone zone screen land
```

OS-Specific DoS Attacks

If an attacker not only identifies the IP address and responsive port numbers of an active host but also its operating system (OS), instead of resorting to brute-force attacks, he or she can launch more elegant attacks that can produce one- or two-packet “kills”. The attacks presented in this section can cripple a system with minimum effort. If your NetScreen device is protecting hosts susceptible to these attacks, you can enable the NetScreen device to detect these attacks and block them before they reach their target.

Ping of Death

The maximum allowable IP packet size is 65,535 bytes, including the packet header, which is typically 20 bytes long[10]. An ICMP echo request is an IP packet with a pseudo header, which is 8 bytes long[11]. Therefore, the maximum allowable size of the data area of an ICMP echo request is 65,507 bytes (65,535 - 20 - 8 = 65,507).

However, many ping implementations allow the user to specify a packet size larger than 65,507 bytes. A grossly oversized ICMP packet can trigger a range of adverse system reactions such as denial of service (DoS), crashing, freezing, and rebooting.

When you enable the Ping of Death SCREEN option, the NetScreen device detects and rejects such oversized and irregular packet sizes even when the attacker hides the total packet size by purposefully fragmenting it.

Note: *For information about Ping of Death, see http://www.insecure.org/sploits/ping-o-death.html.*

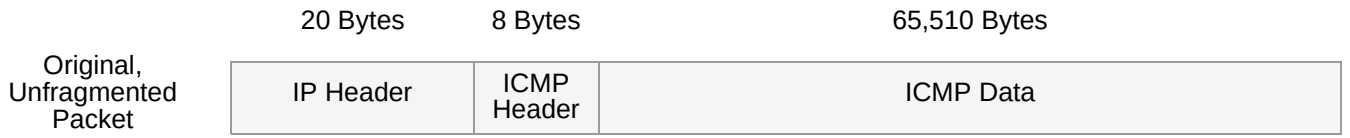

The size of this packet is 65,538 bytes. It exceeds the size limit prescribed by RFC 791, “Internet Protocol”, which is 65,535 bytes. As the packet is transmitted, it becomes broken into numerous fragments. The reassembly process might cause the receiving system to crash.

10. For information about IP specifications, see RFC 791, “Internet Protocol”.
11. For more information about ICMP specifications, see RFC 792, “Internet Control Message Protocol”.

To enable protection against a Ping of Death attack, do either of the following, where the specified zone is that in which the attack originates:

WebUI

Screening > Screen (Zone: select a zone name): Select **Ping of Death Attack Protection**, and then click **Apply**.

CLI

```
set zone zone screen ping-death
```

Teardrop Attack

Teardrop attacks exploit the reassembly of fragmented IP packets. In the IP header, one of the fields is the fragment offset field, which indicates the starting position, or “offset”, of the data contained in a fragmented packet relative to the data of the original unfragmented packet.

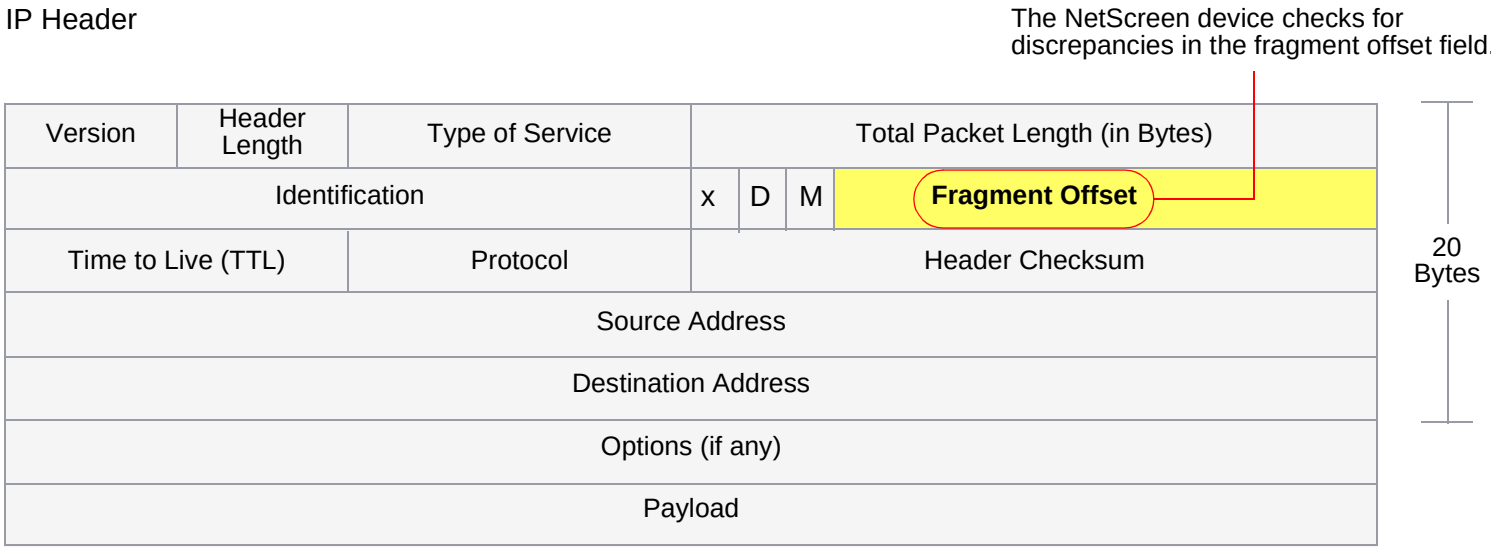

When the sum of the offset and size of one fragmented packet differ from that of the next fragmented packet, the packets overlap, and the server attempting to reassemble the packet can crash, especially if it is running an older operating system that has this vulnerability.

Fragment Discrepancy

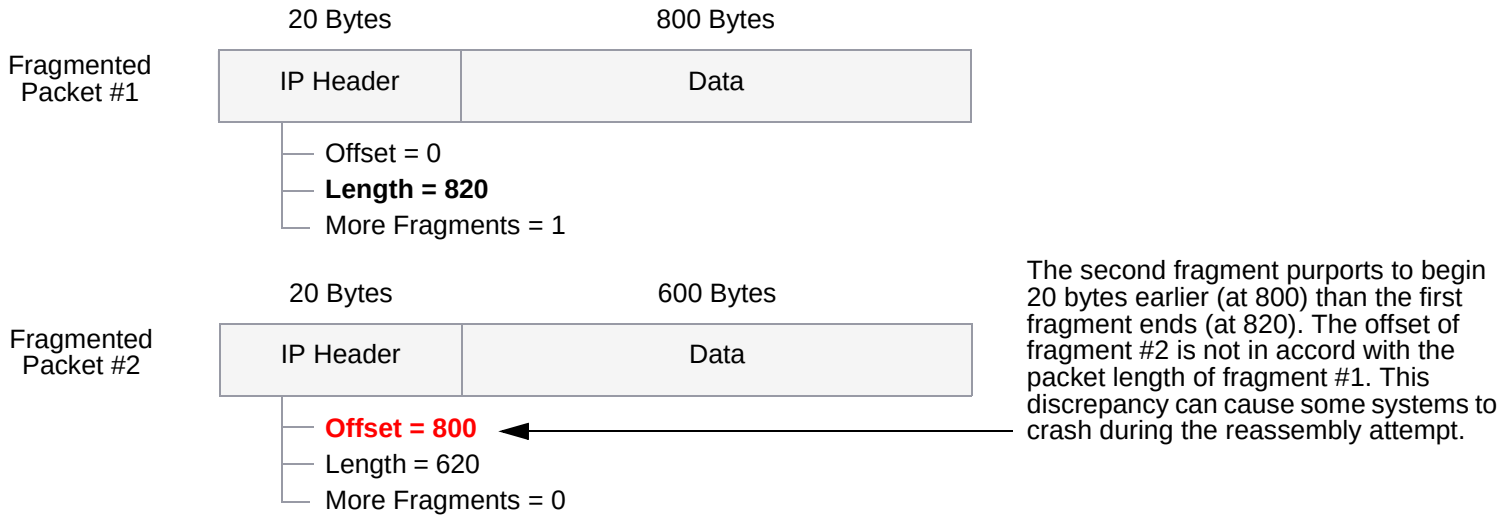

After you enable the Teardrop Attack SCREEN option, whenever the NetScreen detects this discrepancy in a fragmented packet, it drops it.

To enable protection against a Teardrop attack, do either of the following, where the specified zone is that in which the attack originates:

WebUI

Screening > Screen (Zone: select a zone name): Select **Teardrop Attack Protection**, and then click **Apply**.

CLI

```
set zone zone screen tear-drop
```

WinNuke

WinNuke is a DoS attack targeting any computer on the Internet running Windows. The attacker sends a TCP segment—usually to NetBIOS port 139 with the urgent (URG) flag set—to a host with an established connection. This introduces a NetBIOS fragment overlap, which causes many machines running Windows to crash. After rebooting the attacked machine, the following message appears, indicating that an attack has occurred:

> An exception OE has occurred at 0028:[address] in VxD MSTCP(01) + 000041AE. This was called from 0028:[address] in VxD NDIS(01) + 00008660. It may be possible to continue normally.
>
> Press any key to attempt to continue.
>
> Press CTRL+ALT+DEL to restart your computer. You will lose any unsaved information in all applications.
>
> Press any key to continue.

WinNuke Attack Indicators

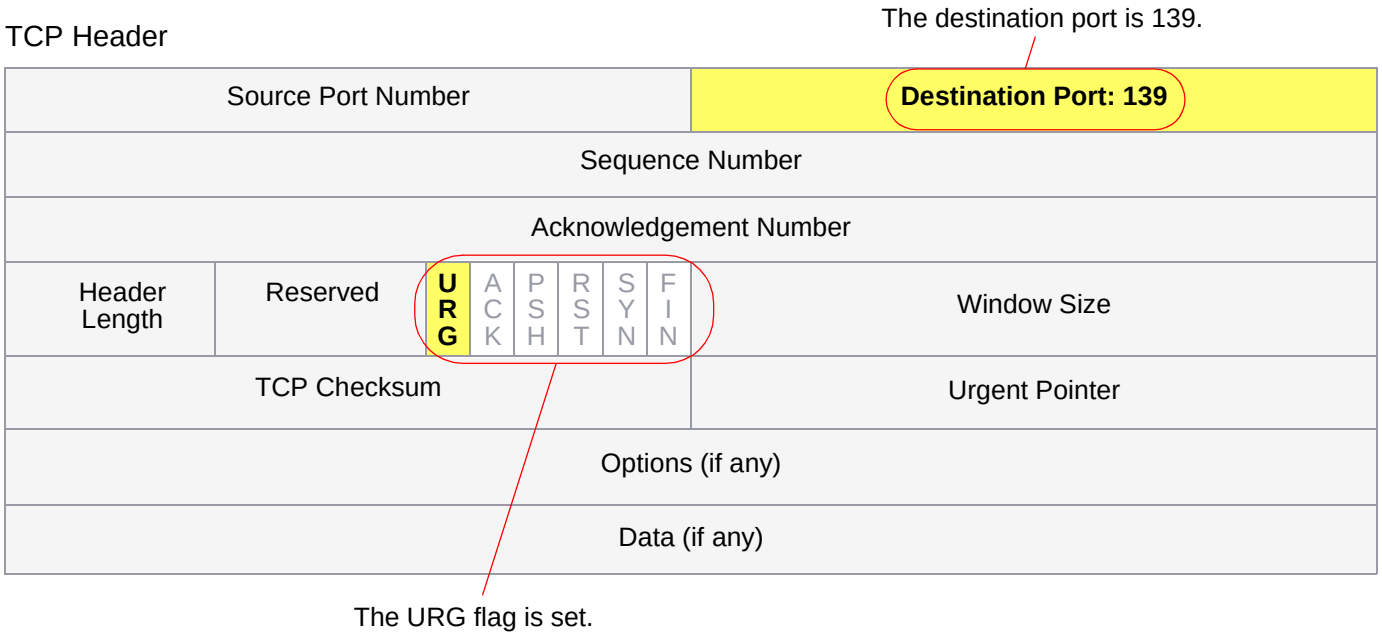

If you enable the WinNuke attack defense SCREEN option, the NetScreen device scans any incoming Microsoft NetBIOS session service (port 139) packets. If the NetScreen device observes that the URG flag is set in one of those packets, it unsets the URG flag, clears the URG pointer, forwards the modified packet, and makes an entry in the event log noting that it has blocked an attempted WinNuke attack.

To enable protection against a WinNuke attack, do either of the following, where the specified zone is that in which the attack originates:

WebUI

Screening > Screen (Zone: select a zone name): Select **WinNuke Attack Protection**, and then click **Apply**.

CLI

```
set zone zone screen winnuke
```

4

Content Monitoring and Filtering

NetScreen provides broad protection and control of network activity through ScreenOS features and the pairing of NetScreen with Websense and Trend Micro products.

NetScreen provides some content monitoring and filtering capabilities within ScreenOS in its malicious URL protection SCREEN option. Furthermore, through the fragment reassembly feature, the NetScreen device can detect URLs even among fragmented TCP segments and fragmented IP packets.

For antivirus (AV) protection, you have a choice on some NetScreen devices to obtain an advanced license key and an AV license key and use an internal AV scanning feature. You can also configure NetScreen devices to work with up to three external Trend Micro AV scanners (after you have first obtained and loaded the two license keys). For URL filtering, you can configure a NetScreen device to work with one or more Websense servers.

This chapter examines how to configure the NetScreen device to perform segment and packet reassembly, monitor HTTP and FTP traffic for malicious URLs, and communicate with other devices to perform AV scanning and URL filtering. The chapter is organized into the following sections:

- "Fragment Reassembly" on page 72
  - "Malicious URL Protection" on page 72
  - "Application Layer Gateway" on page 73
- "Antivirus Scanning" on page 76
  - "Internal AV Scanning" on page 77
  - "External AV Scanning" on page 90
- "URL Filtering" on page 113

FRAGMENT REASSEMBLY

Typically, a network forwarding device such as a router or switch does not reassemble fragmented packets that it receives. It is the responsibility of the destination host to reconstruct the fragmented packets when they all arrive. Because the purpose of forwarding devices is the efficient delivery of traffic, queuing fragmented packets, reassembling them, then refragmenting them, and forwarding them is unnecessary and inefficient. However, passing fragmented packets through a firewall is insecure. An attacker can intentionally break up packets to conceal traffic strings that the firewall otherwise would detect and block.

ScreenOS allows you to enable fragment reassembly on a per zone basis. Doing so allows the NetScreen device to expand its ability to detect and block malicious URL strings, and to improve its ability to provide an application layer gateway (ALG) to check the data portions of packets.

Malicious URL Protection

In addition to the URL filtering feature explained later in this chapter (see “URL Filtering” on page 113), you can define up to 16 malicious URL string patterns, each of which can be up to 24 characters long, for malicious URL protection at the zone level. With the Malicious URL blocking feature enabled, the NetScreen device examines the data payload of all HTTP and FTP packets. If it locates a URL and detects that the beginning of its string—up to a specified number of characters—matches the pattern you defined, the NetScreen device blocks that packet from passing the firewall.

A resourceful attacker, realizing that the string is known and might be guarded against, can deliberately fragment the IP packets or TCP segments and thereby make the pattern unrecognizable during a packet-by-packet inspection. For example, if the malicious URL string is **120.3.4.5/level/50/exec**, IP fragmentation might break up the string into the following sections:

- First packet: **120.**
- Second packet: **3.4.5/level/50**
- Third packet: **/exec**

Individually, the fragmented strings can pass undetected through the NetScreen device, even if you have the string defined as **120.3.4.5/level/50/exec** with a length of 20 characters. The string in the first packet—”120.”— matches the first part of the defined pattern, but it is shorter than the required length of 20 matching characters. The strings in the second and third packets do not match the beginning of the defined pattern, and so too pass without impedance.

However, if the packets are reassembled, the fragments combine to form a recognizable string that the NetScreen device can block. Using the Fragment Reassembly feature, the NetScreen device can buffer fragments in a queue, reassemble them into a complete packet, and then inspect that packet for a malicious URL. Depending on the results of this reassembly process and subsequent inspection, the NetScreen device performs one of the following steps:

- If the NetScreen device discovers a malicious URL, it drops the packet and enters the event in the log.
- If the NetScreen device cannot complete the reassembly process, a time limit is imposed to age out and discard fragments.
- If the NetScreen device determines that the URL is not malicious but the reassembled packet is too big to forward, the NetScreen device fragments that packet into multiple packets and forwards them.
- If the NetScreen device determines that the URL is not malicious and does not need to fragment it, it then forwards the packet.

Application Layer Gateway

NetScreen provides an application layer gateway (ALG) for a number of protocols, such as DNS, FTP, H.323, and HTTP. Of these, fragment reassembly can be an important component in the enforcement of policies involving FTP and HTTP services. The ability of the NetScreen firewall to screen packets for protocols such as FTP-Get and FTP-Put requires it to examine not only the packet header but also the data in the payload. For example, there might be two policies, one denying FTP-put from the Untrust to DMZ zones, and another permitting FTP-get from the Untrust to the DMZ zones:

```
set policy from untrust to dmz any any ftp-put deny
set policy from untrust to dmz any any ftp-get permit
```

To distinguish the two types of traffic, the NetScreen firewall examines the payload. If it reads **RETR** *filename*, the FTP client has sent a request to get (or “retrieve”) the specified file from the FTP server, and the NetScreen device allows the packet to pass. If the NetScreen device finds **STOR** *filename*, the client has sent a request to put (or “store”) the specified file on the server, and the NetScreen device blocks the packet.

To get around this defense, an attacker can deliberately fragment a single FTP-put packet into two packets that contain the following text in their respective payloads: packet 1: **ST**; packet 2: **OR** *filename*. When the NetScreen device inspects each packet individually, it does not find the string **STOR** *filename*, and consequently allows them both to pass.

However, if the packets are reassembled, the fragments combine to form a recognizable string upon which the NetScreen device can act. Using the Fragment Reassembly feature, the NetScreen device buffers the FTP fragments in a queue, reassembles them into a complete packet, and then inspects that packet for the complete FTP request. Depending on the results of this reassembly process and subsequent inspection, the NetScreen device performs one of the following steps:

- If the NetScreen device discovers an FTP-put request, it drops the packet and enters the event in the log.
- If the NetScreen device cannot complete the reassembly process, a time limit is imposed to age out and discard fragments.
- If the NetScreen device discovers an FTP-get request but the reassembled packet is too big to forward, the NetScreen device fragments that packet into multiple packets and forwards them.
- If the NetScreen device discovers an FTP-get request and does not need to fragment it, it then forwards the packet.

Example: Blocking Malicious URLs in Packet Fragments

In this example, you define the following three malicious URL strings and enable the malicious URL blocking option:

- Malicious URL #1
  - ID: Perl
  - Pattern: scripts/perl.exe
  - Length: 14
- Malicious URL #2
  - ID: CMF
  - Pattern: cgi-bin/phf
  - Length: 11
- Malicious URL #3
  - ID: DLL
  - Pattern: 210.1.1.5/msadcs.dll
  - Length: 18

The values for “length” indicate the number of characters in the pattern that must be present in a URL—starting from the first character—for a positive match. Note that for #1 and #3, not every character is required.

You then enable fragment reassembly for the detection of the URLs in fragmented HTTP and FTP traffic arriving at an Untrust zone interface.

WebUI

Screening > Mal-URL (Zone: Untrust): Enter the following, and then click **OK**:

ID: perl

Pattern: /scripts/perl.exe

Length: 14

Screening > Mal-URL (Zone: Untrust): Enter the following, and then click **OK**:

ID: cmf

Pattern: cgi-bin/phf

Length: 11

Screening > Mal-URL (Zone: Untrust): Enter the following, and then click **OK**:

ID: dll

Pattern: 210.1.1.5/msadcs.dll

Length: 18

Screening > Mal-URL (Zone: Untrust): Select the **IP/TCP Reassembly for ALG** check box, and then click **OK**.

CLI

```
set zone untrust screen mal-url perl “scripts/perl.exe“ 14
set zone untrust screen mal-url cmf “cgi-bin/phf” 11
set zone untrust screen mal-url dll “210.1.1.5/msadcs.dll” 18
set zone untrust screen reassembly-for-alg
save
```

ANTIVIRUS SCANNING

A virus is an executable code that infects or attaches itself to other executable code so that it can reproduce itself. Some viruses are malicious, erasing files or locking up systems. Others present a problem merely in the act of infecting other files, as their propagation may overwhelm the infected host or network with excessive amounts of bogus data.

In conjunction with Trend Micro antivirus (AV) technology, NetScreen provides two AV solutions:

- Internal AV scanning
- External AV scanning

With internal AV scanning, the AV scanner is inside the NetScreen device as part of ScreenOS. Using a NetScreen device that supports internal AV simplifies deployment and management. It is a cost-effective choice for remote sites, small offices, retail outlets, and telecommuters. For information on configuring the internal AV scanning feature, see “Internal AV Scanning” on page 77.

With external AV scanning, the AV scanner is a separate device to which the NetScreen device forwards traffic that requires scanning. Using a NetScreen device that supports one or more external AV scanners provides a flexible and scalable approach. You can begin with one AV scanner, but if the protected network grows, you can add more scanners (up to three total) to process increased traffic loads. For information on configuring the external AV scanning feature, see “External AV Scanning” on page 90.

Internal AV Scanning

Some NetScreen devices provide antivirus (AV) scanning for specific application-layer transactions using an internal AV scanner developed by Trend Micro. To use the internal AV scanner to scan network traffic for viruses, you reference the internal AV scanner in your security policy.

You can configure the internal AV scanner to examine network traffic from several protocols including Simple Mail Transfer Protocol (SMTP), Hypertext Transfer Protocol (HTTP), and Post Office Protocol - version 3 (POP3). After verifying that it has received the entire content of the SMTP, HTTP or POP3 packet, the internal AV scanner examines the data for viruses. It does this by referencing a virus pattern file to identify virus signatures. When the internal AV scanner detects a virus, the NetScreen device drops the content and sends a message to the client indicating that the content is infected. If the scanner does not detect a virus, the NetScreen device forwards the content to its intended destination.

For SMTP traffic scanning, the NetScreen device redirects traffic from a local SMTP client to the internal AV scanner before sending it to the local mail server.

SMTP Antivirus Scanning

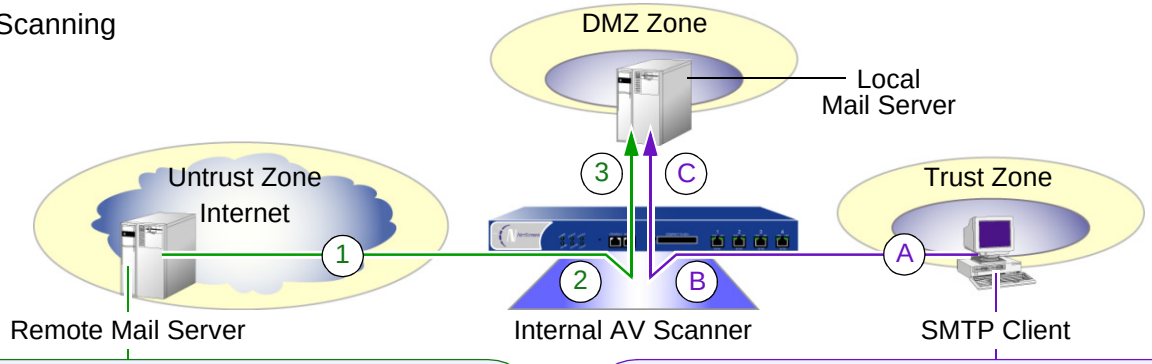

1. A remote mail server forwards an e-mail message via SMTP to the local mail server.
2. The NetScreen device intercepts the e-mail message and passes the data to the internal AV scanner, which scans it for viruses.
3. After completing the scan, the NetScreen device follows one of two courses:
   - If there is no virus, it forwards the message to the local server.
   - If there is a virus, it sends a message reporting the infection to the remote server.

A. An SMTP client sends an e-mail message to a local mail server.

B. The NetScreen device intercepts the e-mail message and passes the data to the internal AV scanner, which scans it for viruses.

C. After completing the scan, the NetScreen device follows one of two courses:
   - If there is no virus, it forwards the message to the local server.
   - If there is a virus, it sends a message reporting the infection to the client.

For POP3 traffic scanning, the NetScreen device redirects traffic from a local mail server to the internal AV scanner before sending it to the local POP3 client.

POP3 Antivirus Scanning

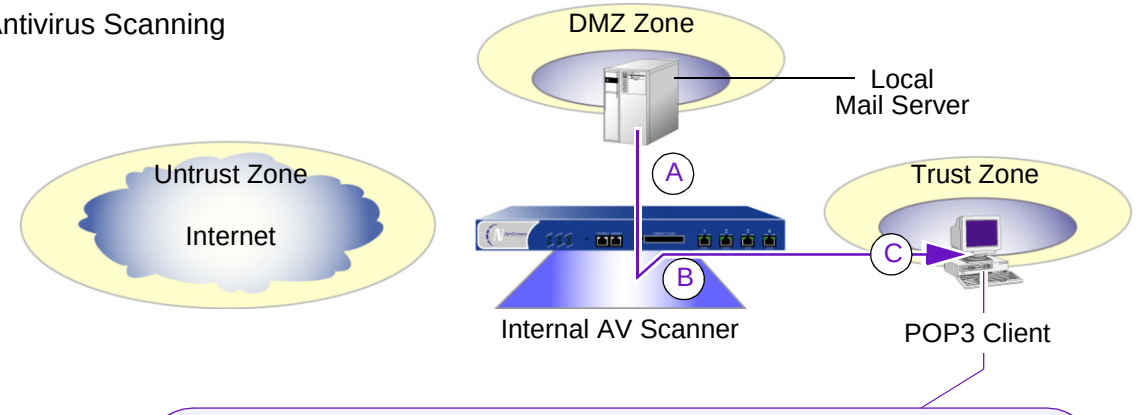

A. The POP3 client downloads an e-mail message from the local mail server.

B. The NetScreen device intercepts the e-mail message and passes the data to the internal AV scanner, which scans it for viruses.

C. After completing the scan, the NetScreen device follows one of two courses:

- If there is no virus, it forwards the message to the client.
- If there is a virus, it sends a message reporting the infection to the client.

For HTTP traffic scanning, the NetScreen device redirects replies from a Web server responding to the client that made HTTP requests to the internal AV scanner before forwarding the traffic to the client.

HTTP Antivirus Scanning

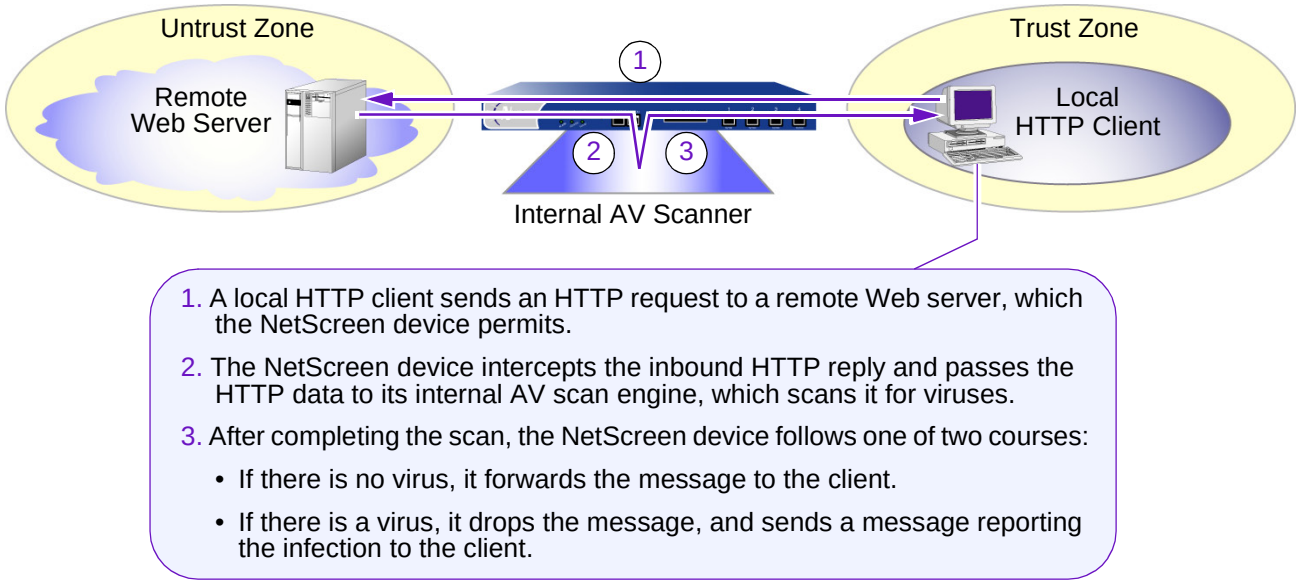

Note: *The internal AV scanner examines HTTP downloads; that is, HTTP data in replies from a Web server to HTTP requests from a client. The internal AV scanner does not scan uploads, such as when an HTTP client completes a questionnaire on a Web server or when a client writes a message in an e-mail originating on a Web server.*

For HTTP webmail traffic scanning, the NetScreen device redirects replies from a Web server responding to the client that made HTTP web-mail requests to the internal AV scanner before forwarding the traffic to the client.

HTTP Web-Mail Antivirus Scanning

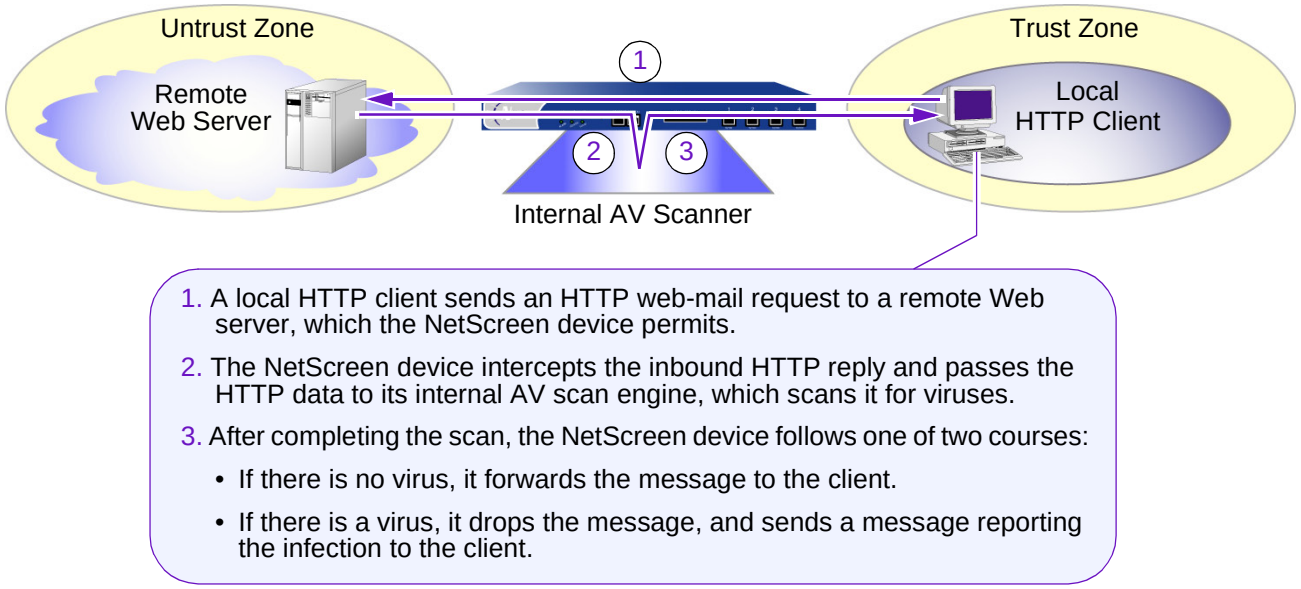

Enabling Internal AV Scanning

Internal AV scanning requires that you load a database of AV patterns onto the NetScreen device and update the pattern file periodically. To do so, you must register the device and purchase a subscription for the AV signature service. The subscription allows you to load the current version of the database and update it as newer versions become available for the life of the subscription. The procedure for initiating the AV signature service varies:

- If you purchased a NetScreen device with AV functionality, you can load an AV pattern file for a short period of time after the initial purchase. You must, however, register the device and purchase a subscription for AV signatures to continue receiving pattern updates.
- If you are upgrading a current NetScreen device to use internal AV scanning, you must register the device and purchase a subscription for AV signatures before you can begin loading the AV pattern file. After completing the registration process, you must wait for a period of up to 4 hours before initiating the AV pattern file download.

> ***Note:*** *For more information about the AV signature service, see "Registration and Activation of Signature Services" on page* ***2*** *-538.*

The process of updating the AV pattern file is as follows:

1. From the NetScreen device, specify the URL address of the external pattern file server to retrieve a server initialization file called server.ini.

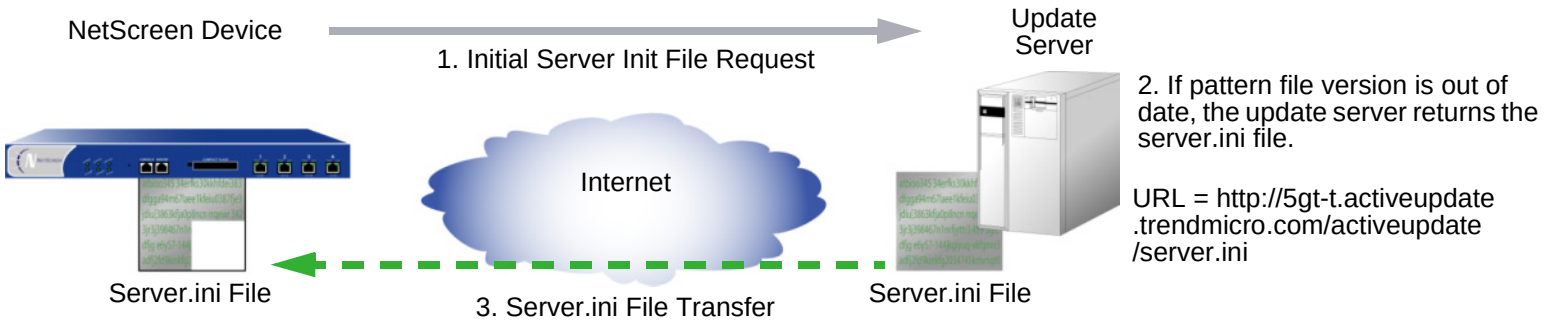

2. After the NetScreen device downloads the server initialization file, it parses it to obtain information about the updated pattern file including the pattern file version and size, and the location of the external pattern file server.

   Note: *ScreenOS contains a CA certificate for authenticating communication with the pattern file server.*

3. If the current pattern file is out of date, the NetScreen device retrieves the updated pattern file from the external pattern file server automatically.

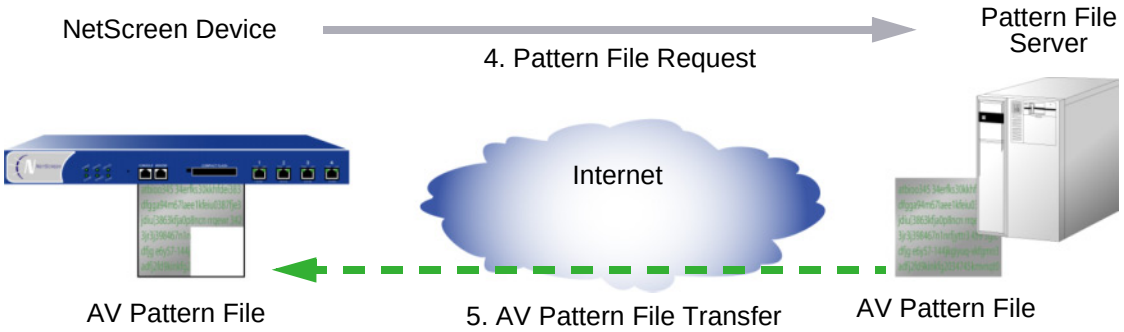

4. After the NetScreen device downloads the pattern file, it verifies that the AV subscription is still valid. If the AV signature service subscription is valid, the pattern file is updated. If the subscription is expired, the pattern file update fails and an error message appears indicating that the AV subscription is expired.

Note: *The total estimated time to complete a pattern update is approximately 3 minutes. This time may vary depending upon the pattern file size and existing network traffic. After completing the pattern file update, the NetScreen device re-initializes the internal AV scanner in order to use the new pattern.*

Updating the Pattern File Automatically or Semi-Automatically

Updates to the pattern file are added as new viruses propagate. You can configure the NetScreen device to update the pattern file either automatically on a regular basis or semi-automatically.

Note: *Once your subscription expires, the update server no longer permits you to update the AV pattern file.*

Example: Automatic Pattern Update

In this example, you configure the NetScreen device to update the pattern file automatically every 15 minutes. (The default AV pattern update interval is 60 minutes.) The pattern update server is located at the following URL address: http://5gt-t.activeupdate.trendmicro.com/activeupdate/server.ini

WebUI

Screening > Antivirus > Scan Manager: Enter the following, and then click **OK**:

Pattern Update Server:
http://5gt-t.activeupdate.trendmicro.com/activeupdate/server.ini

Auto Pattern Update: (select), Interval: 15 minutes (10~10080)

CLI

```
set av scan-mgr pattern-update-url http://5gt-t.activeupdate.trendmicro.com/
    activeupdate/server.ini interval 15
save
```

Example: Semi-Automatic Pattern Update

In this example, you configure the NetScreen device to update the pattern file semi-automatically. The pattern update server is located at the following URL address:
http://5gt-t.activeupdate.trendmicro.com/activeupdate/server.ini

WebUI

Screening > Antivirus > Scan Manager: Enter the following, and then click **OK**:

Pattern Update Server:
http://5gt-t.activeupdate.trendmicro.com/activeupdate/server.ini

Update Now: (select)

CLI

```
set av scan-mgr pattern-update-url
    http://5gt-t.activeupdate.trendmicro.com/activeupdate/server.ini
exec av scan-mgr pattern-update
```

Configuring Content Processing

The internal AV scanner examines SMTP, HTTP (webmail-only) and POP3 traffic by default.

Note: *The internal AV scanner examines specific HTTP webmail patterns only. The patterns for Yahoo!, Hotmail, and AOL mail services are pre-defined.*

You can change the default behavior so that the internal AV scanner examines specific network traffic only.

Example: Internal AV Scanning for SMTP

In this example, you configure the internal AV scanner to examine SMTP traffic only.

WebUI

Screening > Antivirus > Scan Manager: Enter the following, and then click **OK**:

Protocols to be scanned:

SMTP: (select)

CLI

```
set av scan-mgr content smtp timeout 20
save
```

Example: Internal AV Scanning for SMTP and HTTP

In this example, you configure the internal AV scanner to examine all SMTP and HTTP traffic.

WebUI

Screening > Antivirus > Scan Manager: Enter the following, and then click **OK**:

Protocols to be scanned:

SMTP: (select)

HTTP: (select); ALL HTTP: (select)

CLI

```
set av scan-mgr content smtp timeout 20
set av scan-mgr content http timeout 20
unset av http webmail enable
save
```

Configuring Decompression and Maximum Content Size

When it receives content, the internal AV scanner decompresses any compressed files. It decompresses up to 2 layers of compressed files by default. For example, if the scanner receives a file with an attachment, and the attachment is a compressed file layered within another compressed file, the scanner may decompress both layers in order to detect any viruses. You can configure the internal AV scanner to decompress up to 20 compressed files layered within another.

The internal AV scanner examines a maximum of 8 messages and 16 MB of “decompressed” file content at any specific time. If the total number of messages or size of the content received concurrently exceeds these limits, the scanner passes the content without checking for viruses by default. For example, the scanner can receive and examine four 4-MB messages concurrently. If the scanner receives nine 2-MB messages concurrently, it passes the content without scanning it. You can change this default behavior so that the internal AV scanner drops traffic instead of passing it.

Example: Dropping Large Files

In this example, you configure the internal AV scanner to decompress up to 10 files layered within another. You also configure the scanner to drop content if either the total number of messages received concurrently exceeds 4 messages or the total “decompressed” size of the content exceeds 12 MB.

WebUI

Screening > Antivirus > Scan Manager: Enter the following, and then click **OK**:

File decompression: 10 layers (1~4)

Drop: (select) file if it exceeds 3000 KB (4000~20000)

Drop: (select) file if the number of files exceeds 4 files (1~8)

CLI

```
set av scan-mgr decompress-layer 10
set av scan-mgr max-msgs 4
set av scan-mgr max-content-size 3000
set av scan-mgr max-content-size drop
save
```

Applying Internal AV Scanning

To apply internal AV scanning to SMTP, HTTP, or POP3 network traffic, you must reference the pre-defined internal AV scanner in policies.

Example: Internal AV Scanning (POP3)

In this example, you reference the internal AV scanner in a firewall policy permitting POP3 traffic from addresses in the Trust zone to the mail server (“mailsrv1”, 1.2.2.5) in the DMZ zone. All zones are in the trust-vr routing domain.

WebUI

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Enter the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet2): Enter the following, and then click **OK**:

Zone Name: DMZ

Static IP: (select this option when present)

IP Address/Netmask: 1.2.2.1/24

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

2. Address

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: mailsrv1

IP Address/Domain Name:

IP/Netmask: (select), 1.2.2.5/32

Zone: DMZ

3. POP3 Internal AV Scanning

Screening > Antivirus > Scan Manager: Enter the following, and then click **OK**:

Protocols to be scanned:

POP3: (select)

4. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 1.1.1.250

5. Policy

Policies > (From: Trust, To: DMZ) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), mailsrv1

Service: POP3

Action: Permit

> Advanced**:** Move the following AV objects, and then click **Return** to set the advanced options and return to the basic configuration page:

Select **scan-mgr** and use the **<<** button to move the AV object from the Available AV Object Names column to the Attached AV Object Names column.

CLI

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet2 zone dmz
set interface ethernet2 ip 1.2.2.1/24
set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
```

2. Address

```
set address dmz mailsvr1 1.2.2.5/32
```

3. POP3 Internal AV Scanning

```
set av-scan-mgr content pop3 timeout 20
```

4. Route

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
```

5. Policy

```
set policy from trust to dmz any mailsvr1 pop3 permit av scan-mgr
save
```

External AV Scanning

Most NetScreen devices can interoperate with an external antivirus (AV) scanner produced by Trend Micro called InterScan VirusWall Edition 3.6. You can configure the NetScreen device to forward Simple Mail Transfer Protocol (SMTP) and Hypertext Transfer Protocol (HTTP) traffic to the VirusWall AV scanner. The protocol for communication between the NetScreen device and the VirusWall scanner is called Content Scanning Protocol (CSP), version 1.5.

Note: *NetScreen does not support AV for virtual systems. On systems that support both virtual systems and AV, AV is only available at the root level.*

When the VirusWall scanner receives the entire content of an SMTP or HTTP packet, it examines the data for viruses. It has a database of virus patterns that it uses to identify virus signatures. If it finds anything amiss, the VirusWall quarantines the infected data for further study and returns the SMTP or HTTP file—without the infected data—to the NetScreen device. The NetScreen device then forwards the file to the intended recipient.

Whenever the VirusWall detects a virus, both the NetScreen device and the VirusWall make an event log entry identifying the detected virus.

Note: *To learn how to configure the Trend Micro InterScan VirusWall to communicate with the NetScreen device, as well as how to configure other settings, refer to the Trend Micro product documentation.*

For SMTP traffic scanning, the NetScreen device can redirect traffic from a remote mail server or from a local SMTP client to the VirusWall antivirus scanner before sending it to the local mail server.

SMTP Antivirus Scanning

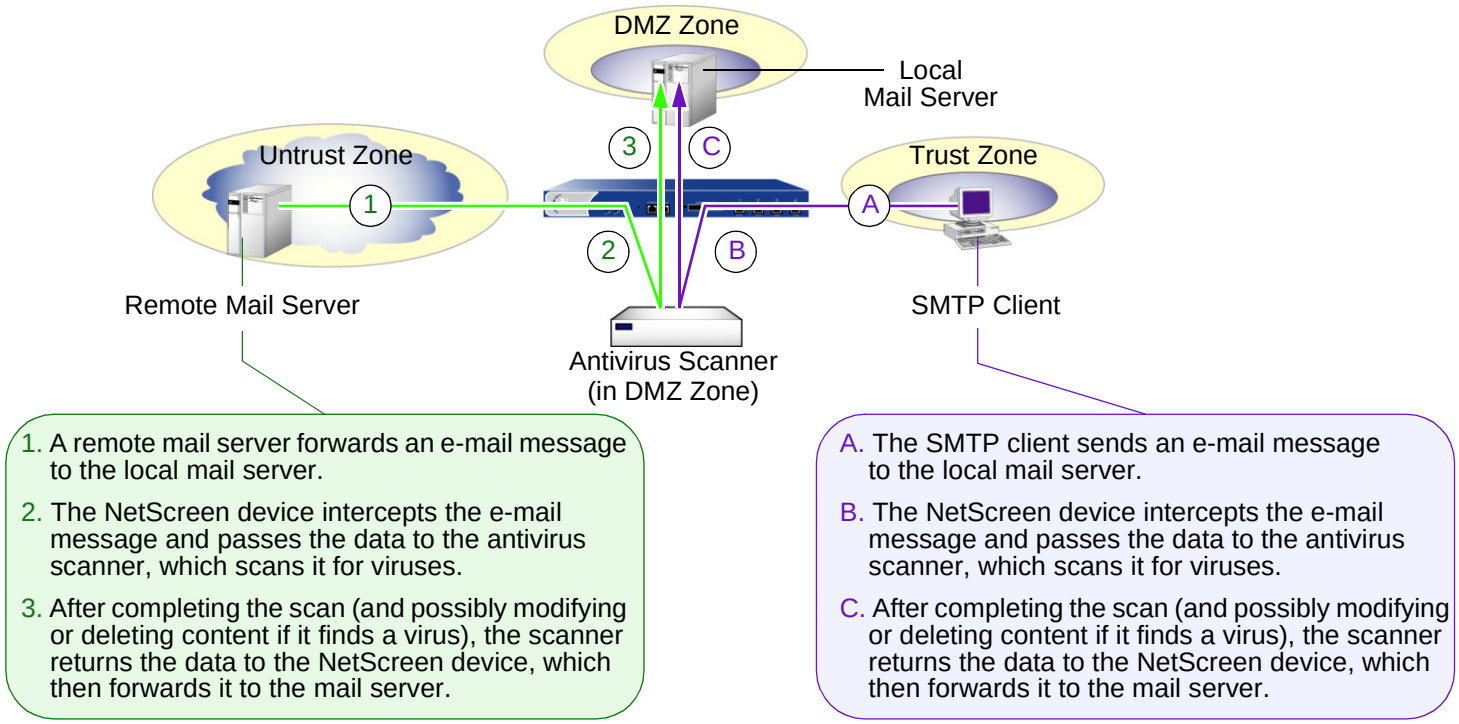

For HTTP traffic scanning, the NetScreen device can redirect replies from a Web server responding to the client that made HTTP requests to an antivirus scanner before forwarding the traffic to the client.

HTTP Antivirus Scanning

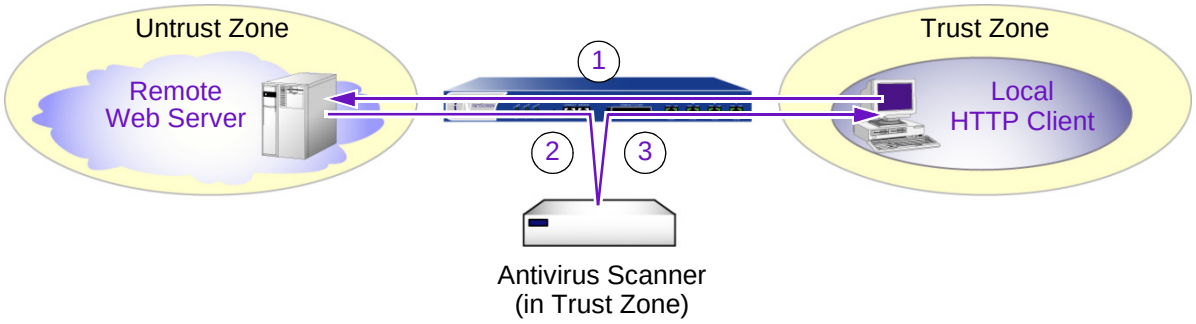

1. A local HTTP client sends an HTTP request to a remote Web server, which the NetScreen device permits.
2. The NetScreen device intercepts the inbound HTTP reply and passes the HTTP data to the antivirus server, which scans it for viruses.
3. After completing the scan (and possibly modifying or deleting content if it finds a virus), the antivirus server returns the data to the NetScreen device, which then forwards it to the client.

Note: *The antivirus scanner scans HTTP downloads; that is, HTTP data in replies from a Web server to HTTP requests from a client. The antivirus scanner does not scan uploads, such as when an HTTP client completes a questionnaire on a Web server or when a client writes a message in an e-mail originating on a Web server.*

Defining AV Objects

An antivirus object (AV object) is the term NetScreen uses to refer to an external antivirus scanner. You can define up to three AV objects to increase bandwidth capacity. When you create an AV object, you must define the following three components:

- AV object name
- IP address or domain name (resolved to an IP address by DNS) of the antivirus scanner
- Content type: HTTP, or SMTP, or both

When you define only one or two of the above components, the state of the AV object is considered incomplete. When you define all three, it is considered complete. For example:

```
set av scanner1 server-name 1.2.2.25

ns208A_5.0.0_beta3-> get av scanner1
<AV object scanner1>
    scanner name:    1.2.2.25
    scanner ip:      1.2.2.25
    scanner port:    3300
    status:          incomplete
    applications:    0
    scanned bytes:   0
    policy ref cnt:  0
```

The AV object is **incomplete** because it has a name (“scanner1”) and an address (1.2.2.25) but not a content type.

```
set av scanner1 server-name 1.2.2.25
set av scanner1 content http

get av scanner1
<AV object scanner1>
    scanner name:    1.2.2.25
    scanner ip:      1.2.2.25
    scanner port:    3300
    HTTP:            timeout 180 seconds
    status:          complete
    applications:    0
    scanned bytes:   0
    policy ref cnt:  0
```

The AV object is **complete** because it has a name, an address, and a content type (HTTP).

There are a few optional parameters that you can set for an AV object:

- **Port number:** By default, the port number that Content Scanning Protocol (CSP) uses for communication between a NetScreen device and a Trend Micro InterScan VirusWall is 3300. You can change this number on a per-AV object basis.

  ```
  set av name_str server-name { ip_addr | domain_name } port number
  unset av name_str server-name { ip_addr | domain_name } port
  ```

  The above **unset av** command returns the port number to the default (3300).

- **Timeout value (in seconds):** By default, a CSP connection times out after 180 seconds of inactivity. You can change this value on a per-AV object basis. The range is 1 to 1800 seconds.

  ```
  set av name_str content { http | smtp } timeout number
  unset av name_str content { http | smtp } timeout number
  ```

  The above **unset av** command returns the timeout value to the default (180 seconds).

In addition to the above options that you can set per AV object, you can also set the following parameters, which apply to the antivirus feature at large:

- **Maximum simultaneous TCP connections:** This specifies the maximum number of simultaneous TCP connections between the NetScreen device and all AV objects as a group, not between the NetScreen device and each individual AV object. The default value varies from platform to platform. (Refer to the NetScreen marketing literature for information relevant to your platform.)

WebUI

Screening > Antivirus: Enter a number in the Maximum Number of TCP Connections field, and then click **Apply**.

CLI

```
set av all max-connections number
unset av all max-connections
```

- **CSP resources per source:** A malicious user might simultaneously send a large number of SMTP or HTTP traffic to consume all available Content Scanning Protocol (CSP) resources and thereby hinder the ability of the NetScreen device to forward any other traffic to the AV scanner. To prevent such activity from succeeding, the NetScreen device can impose a maximum percentage of CSP resources that traffic from a single source can consume at any one time. The default maximum percentage is 70%. You can change this setting to any value between 1% and 100%, where 100% does not impose any restriction on the amount of CSP resources that traffic from a single source can consume.

  ### *WebUI*

  Screening > Antivirus: Enter a number in the Maximum AV Resources Allowed per AV Client field, and then click **Apply**.

  ### *CLI*

  ```
  set av all resources number
  unset av all resources
  ```

  The above **unset av** command returns the maximum percentage of CSP resources per source to the default (70%).

- **Fail mode behavior:** Fail mode is the behavior that the NetScreen device applies when it loses connectivity with the VirusWall scanner—either permit the unexamined traffic or block it. By default, if a NetScreen device cannot reach a VirusWall scanner, it blocks HTTP and SMTP traffic that a policy with antivirus checking enabled permits. You can change the default behavior from block to permit.

  ### *WebUI*

  Screening > Antivirus: Select the Fail Mode Traffic Permit check box to permit unexamined traffic, or clear the check box to block it, and then click **Apply**.

  ### *CLI*

  ```
  set av all fail-mode traffic permit
  unset av all fail-mode traffic
  ```

  The above **unset av** command returns the fail mode behavior to the default (block unexamined traffic).

- **Fail mode threshold:** Fail mode is the state when a number of consecutive failed connection attempts to an AV object exceeds a threshold. By default, that threshold is 150, and it applies to all AV objects. If the number of consecutive failed attempts exceeds this threshold, the NetScreen device waits for a defined interval of time (five minutes) before renewing its connection efforts. You can change the threshold if the default setting seems too high or too low for your needs.

WebUI

Screening > Antivirus: Enter a number in the Fail Mode Scanner Threshold field, and then click **Apply**.

CLI

```
set av all fail-mode scanner threshold number
unset av all fail-mode scanner
```

If the you want the NetScreen device to resume its efforts to connect to a particular AV object before the wait interval has elapsed, you can enter the following command:

```
clear av name_str fail-mode
```

This command clears the failure status so that when the next SMTP or HTTP traffic arrives, the NetScreen device immediately attempts to connect to the AV scanner. If it is successful, the NetScreen device resumes forwarding files for virus scanning to the AV scanner. If its connection attempts are unsuccessful, the status returns to fail mode.

- **HTTP keep-alive:** By default, the NetScreen device uses the HTTP “close” connection option for indicating the end of data transmission. (If necessary, the NetScreen device changes the token in the connection header field from “keep-alive” to “close”.) In this method, when the HTTP server completes its data transmission, it sends a TCP FIN to close the TCP connection and thereby indicate that it has finished sending data. When the NetScreen device receives a TCP FIN, it has all the HTTP data from the server and can instruct the AV scanner to begin scanning.

You can change the default behavior of the NetScreen device to use the HTTP “keep-alive” connection option, which does not send a TCP FIN to indicate the termination of data transmission. The HTTP server must indicate that it has sent all the data in another way, such as by sending the content length in the HTTP header or by some form of encoding. (The method that a server uses varies by server type.) This method keeps the TCP connection open while the antivirus examination occurs, which decreases latency and improves CPU performance. However, it is not as secure as the “close” connection method. You can change this behavior if you find that HTTP connections are timing out during the antivirus examination.

WebUI

Screening > Antivirus: Select the Keep Alive check box to use the “keep-alive” connection option, or clear the check box to use the “close” connection option, and then click **Apply**.

CLI

```
set av http keep-alive
unset av http keep-alive
```

- **HTTP trickling:** HTTP trickling is the forwarding of specified amounts of unscanned HTTP traffic to the requesting HTTP client to prevent the browser window from timing out while the VirusWall examines downloaded HTTP files. (The NetScreen device forwards small amounts of data in advance of transferring an entire scanned file.) By default, HTTP trickling is disabled. To enable it and use the default HTTP trickling parameters, do either of the following:

WebUI

Screening > Antivirus: Select the Trickling Default check box, and then click **Apply**.

CLI

```
set av http trickling default
```

With the default parameters, the NetScreen device employs trickling if the size of an HTTP file is 3 megabytes or larger. Then it forwards 500 bytes of content for every 1 megabyte sent for scanning.

To change the parameters for HTTP trickling, do either of the following:

WebUI

Screening > Antivirus: Enter the following, and then click **Apply**:

Trickling:

Custom: (select)

Minimum Length to Start Trickling: Enter *number1*.

Trickle Size: Enter *number2*.

Trickle for Every MB Sent for Scanning: Enter *number3*.

CLI

```
set av http trickling number1 number3 number2
```

The three number variables have the following meanings:

- *number1:* The minimum size (in megabytes) of an HTTP file to trigger trickling
- *number2:* The size (in bytes) of unscanned traffic that the NetScreen device forwards
- *number3*: The size (in megabytes) of a block of traffic to which the NetScreen device applies trickling

Note: *Data trickled to the client's hard drive appears as a small, unusable file. Because trickling works by forwarding a small amount of data to a client without scanning it, virus code might be among the data that the NetScreen device has trickled to the client. NetScreen advises users to delete such files.*

You can disable HTTP trickling in the WebUI (Screening > Antivirus: Click **Disable** in the Trickling section.) or with the CLI command **set av http trickling 0 0 0**. However, if a file being downloaded is larger than eight megabytes and HTTP trickling is disabled, the browser window will most likely time out.

Example: Defining Three AV Objects

In this example, you define the following AV objects:

- AV Object 1
  - Name: scanner1
  - IP address: 1.2.2.20
  - Port number for Content Scanning Protocol (CSP): 3300 (default)
  - Content: HTTP
  - Timeout: 180 seconds (default)
- AV Object 2
  - Name: scanner2
  - IP address: 1.2.2.30
  - Port number for CSP: 6830
  - Content: SMTP
  - Timeout: 200 seconds
- AV Object 3
  - Name: scanner3
  - IP address: 1.2.2.40
  - Port number for CSP: 6840
  - Content: HTTP and SMTP
  - HTTP Timeout: 120 seconds
  - SMTP Timeout: 200 seconds

The NetScreen device accesses the above addresses through ethernet2, which has IP address 1.2.2.1/24 and is bound to the DMZ zone.

WebUI

1. AV Object 1

   Objects > Antivirus > New: Enter the following, and then click **OK**:

   AV Object Name: scanner1

   Scan Server Name/IP: 1.2.2.20

   Scan Server Port: 3300

   Contents:

   HTTP: (select), Timeout: 180 Seconds

2. AV Object 2

   Objects > Antivirus > New: Enter the following, and then click **OK**:

   AV Object Name: scanner2

   Scan Server Name/IP: 1.2.2.30

   Scan Server Port: 6830

   Contents:

   SMTP: (select), Timeout: 200 Seconds

3. AV Object 3

   Objects > Antivirus > New: Enter the following, and then click **OK**:

   AV Object Name: scanner3

   Scan Server Name/IP: 1.2.2.40

   Scan Server Port: 6840

   Contents:

   HTTP: (select), Timeout: 120 Seconds

   SMTP: (select), Timeout: 200 Seconds

CLI

1. AV Object 1

```
set av scanner1 server-name 1.2.2.20
set av scanner1 content http
```

2. AV Object 2

```
set av scanner2 server-name 1.2.2.30 port 6830
set av scanner2 content smtp timeout 200
```

3. AV Object 3

```
set av scanner3 server-name 1.2.2.40 port 6840
set av scanner3 content http timeout 120
set av scanner3 content smtp timeout 200
save
```

Applying External AV Scanning

After you create one or more AV objects, you can reference them in policies to apply antivirus scanning to HTTP and SMTP traffic. A single AV object can scan HTTP traffic or SMTP traffic or both kinds of traffic. If you reference two or three AV objects in the same policy, then the NetScreen device sends traffic appropriate for scanning to those objects in a sequence that provides load balancing.

The order in which you reference the three AV objects in the policy configuration defines the order in which the NetScreen device sends HTTP and SMTP traffic to them. The AV object that you reference first is the one to which the NetScreen device sends the first file, such as an e-mail message or an HTTP reply, for scanning. In other words, the first AV object has the highest priority. The AV object that you reference second is the one to which the NetScreen device sends a second file if the first AV object is currently scanning another file. It has the second highest priority. The AV object that you reference in the policy configuration third gets a third file if the first two AV objects are both scanning other files. It has the lowest priority.

For example, if you create three AV objects “scanner1”, “scanner2”, and “scanner3” and then reference them in a policy in the following order,

```
set policy id 1 from trust to untrust any any http permit av scanner1
set policy id 1
ns(policy:1)-> set av scanner2
ns(policy:1)-> set av scanner3
```

then the order for sending files to each scanner proceeds as follows:

1. The NetScreen device sends the first file for scanning to scanner1.
2. When a second file to be scanned arrives, the NetScreen device sends it to scanner1 or scanner2 under the following conditions:
   - scanner1 if it has completed its scan of the first file
   - scanner2 if scanner1 is still scanning the first file
3. When a third file arrives, the NetScreen sends it to one of the three AV objects under the following conditions:
   - scanner1 if it is not scanning a file
   - scanner2 if scanner1 is scanning a file but scanner2 is not
   - scanner3 if scanner1 and scanner2 are both scanning files

The above sequence continues when all the scanners are busy scanning multiple files. If all scanners are scanning the same number of files, the NetScreen device sends the next file to scanner1. If scanner1 is scanning fewer files than scanner2 and scanner3, the NetScreen device sends the next file to scanner1. If scanner2 is scanning fewer files than scanner1 and scanner3, the NetScreen device sends the next file to scanner2. If scanner3 is scanning fewer files than scanner1 and scanner2, then the NetScreen device sends the next file to scanner3.

Example: Antivirus with One AV Object

In this example, you create a single AV object named “scanner1” to perform virus scanning on HTTP replies from Web servers in the Untrust zone to clients in the Trust zone. The antivirus scanner is also in the Trust zone. Although you enable antivirus checking for HTTP traffic between the Trust and Untrust zones, no additional policy is necessary to permit CSP traffic between the NetScreen device and scanner1. All zones are in the trust-vr routing domain.

WebUI

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Enter the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

2. AV Object

Objects > Antivirus > New: Enter the following, and then click **OK**:

AV Object Name: scanner1

Scan Server Name/IP: 1.2.2.20

Scan Server Port: 3300

Contents:

HTTP: (select), Timeout: 180 Seconds

3. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 1.1.1.250

4. Policy

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Any

Service: HTTP

Action: Permit

> Advanced**:** Move the following AV object, and then click **Return** to set the advanced options and return to the basic configuration page:

Select **scanner1** and use the **<<** button to move the AV object from the Available AV Object Names column to the Attached AV Object Names column.

CLI

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
```

2. AV Object

```
set av scanner1 server-name 1.2.2.20
set av scanner1 content http
```

3. Route

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
```

4. Policy ID 1

```
set policy id 1 from trust to untrust any any http permit av scanner1
save
```

Example: Antivirus with Two AV Objects

In this example, you define two AV objects named “scanner1” and “scanner2” for scanning HTTP and SMTP traffic. You then reference the AV objects in policies permitting HTTP between the Trust and Untrust zones and SMTP traffic from addresses in the Untrust and Trust zones to the mail server in the DMZ zone. To balance the traffic load sent to the two AV objects, you set up the distribution of antivirus scanning requests to them as follows:

- The NetScreen device redirects all HTTP antivirus scanning replies to the two AV objects. The two policies permitting HTTP traffic each reference both AV objects.

1. First HTTP reply -> scanner1

2. Second HTTP reply -> scanner2 (if scanner1 has not finished scanning the first HTTP reply; if scanner1 is free, then the NetScreen device redirects the second HTTP reply to scanner1)

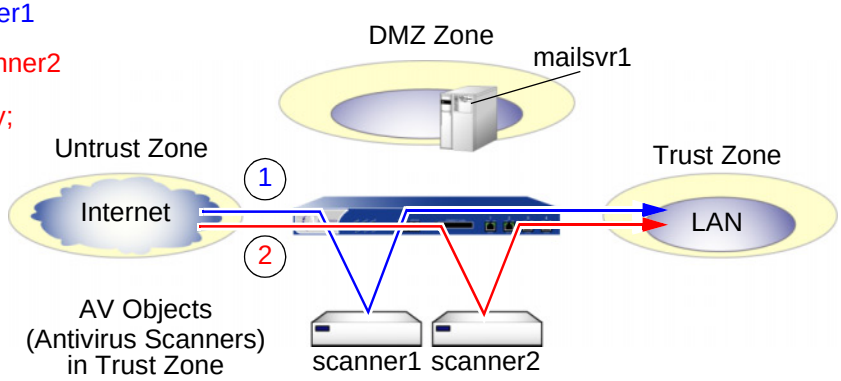

- The NetScreen device sends antivirus scanning requests to scanner1 for all SMTP traffic from the remote mail server (named “r-mail”) in the Untrust zone to the local mail server (named “mailsvr1”) in the DMZ. The NetScreen device sends antivirus scanning requests to scanner2 for all SMTP traffic from the Trust zone.

From the Untrust zone
All SMTP traffic -> scanner1

From the Trust zone
All SMTP traffic -> scanner2

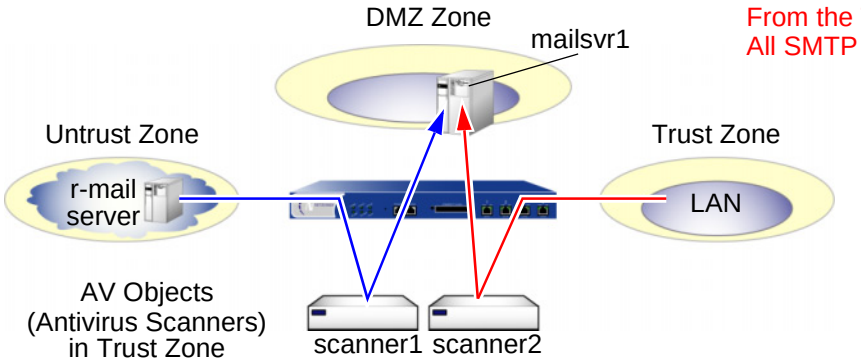

Both AV objects are in the Trust zone. Although you enable antivirus checking on traffic at the policy level, no additional policy is necessary to permit CSP traffic between the NetScreen device and the antivirus scanners. All zones are in the trust-vr routing domain.

WebUI

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Enter the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet2): Enter the following, and then click **OK**:

Zone Name: DMZ

Static IP: (select this option when present)

IP Address/Netmask: 1.2.2.1/24

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

2. AV Object 1

Objects > Antivirus > New: Enter the following, and then click **OK**:

AV Object Name: scanner1

Scan Server Name/IP: 10.1.1.20

Scan Server Port: 3300

Contents:

HTTP: (select), Timeout: 180 Seconds

SMTP: (select), Timeout: 180 Seconds

3. AV Object 2

Objects > Antivirus > New: Enter the following, and then click **OK**:

AV Object Name: scanner2

Scan Server Name/IP: 10.1.1.30

Scan Server Port: 3300

Contents:

HTTP: (select), Timeout: 180 Seconds

SMTP: (select), Timeout: 180 Seconds

4. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: mailsrv1

IP Address/Domain Name:

IP/Netmask: (select), 1.2.2.6/32

Zone: DMZ

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: r-mail

IP Address/Domain Name:

IP/Netmask: (select), 2.2.2.5/32

Zone: Untrust

5. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 1.1.1.250

6. Policy ID 1

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Any

Service: HTTP

Action: Permit

> Advanced**:** Move the following AV objects, and then click **Return** to set the advanced options and return to the basic configuration page:

Select **scanner1** and use the **<<** button to move the AV object from the Available AV Object Names column to the Attached AV Object Names column.

Select **scanner2** and use the **<<** button to move the AV object from the Available AV Object Names column to the Attached AV Object Names column.

7. Policy ID 2

Policies > (From: Untrust, To: DMZ) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), r-mail

Destination Address:

Address Book Entry: (select), mailsrv1

Service: MAIL

Action: Permit

> Advanced**:** Move the following AV objects, and then click **Return** to set the advanced options and return to the basic configuration page:

Select **scanner1** and use the **<<** button to move the AV object from the Available AV Object Names column to the Attached AV Object Names column.

8. Policy ID 3

Policies > (From: Trust, To: DMZ) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), mailsrv1

Service: MAIL

Action: Permit

> Advanced**:** Move the following AV objects, and then click **Return** to set the advanced options and return to the basic configuration page:

Select **scanner2** and use the **<<** button to move the AV object from the Available AV Object Names column to the Attached AV Object Names column.

CLI

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet2 zone dmz
set interface ethernet2 ip 1.2.2.1/24
set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
```

2. AV Object 1

```
set av scanner1 server-name 10.1.1.20
set av scanner1 content http
set av scanner1 content smtp
```

3. AV Object 2

```
set av scanner1 server-name 10.1.1.30
set av scanner1 content http
set av scanner1 content smtp
```

4. Addresses

```
set address dmz mailsvr1 1.2.2.6/32
set address untrust r-mail 2.2.2.5/32
```

5. Route

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
```

6. Policy ID 1

```
ns-> set policy id 1 from trust to untrust any any http permit av scanner1
ns-> set policy id 1
ns(policy:1)-> set av scanner2
ns(policy:1)-> exit
ns->
```

7. Policy ID 2

```
set policy id 2 from untrust to dmz r-mail mailsvr1 mail permit av scanner1
```

8. Policy ID 3

```
set policy id 3 from trust to dmz any mailsvr1 mail permit av scanner2
save
```

URL FILTERING

NetScreen supports URL filtering using the Websense Enterprise Engine, which enables you to block or permit access to different sites based on their URLs, domain names, and IP addresses. With the Websense API built directly into the NetScreen firewall, the NetScreen device can link directly to a Websense URL-filtering server.

The following illustration shows the basic sequence of events when a host in the Trust zone attempts an HTTP connection to a server in the Untrust zone. However, URL filtering determines that the requested URL is prohibited.

A **Blocked** URL

set policy from trust to untrust any any http permit url-filter

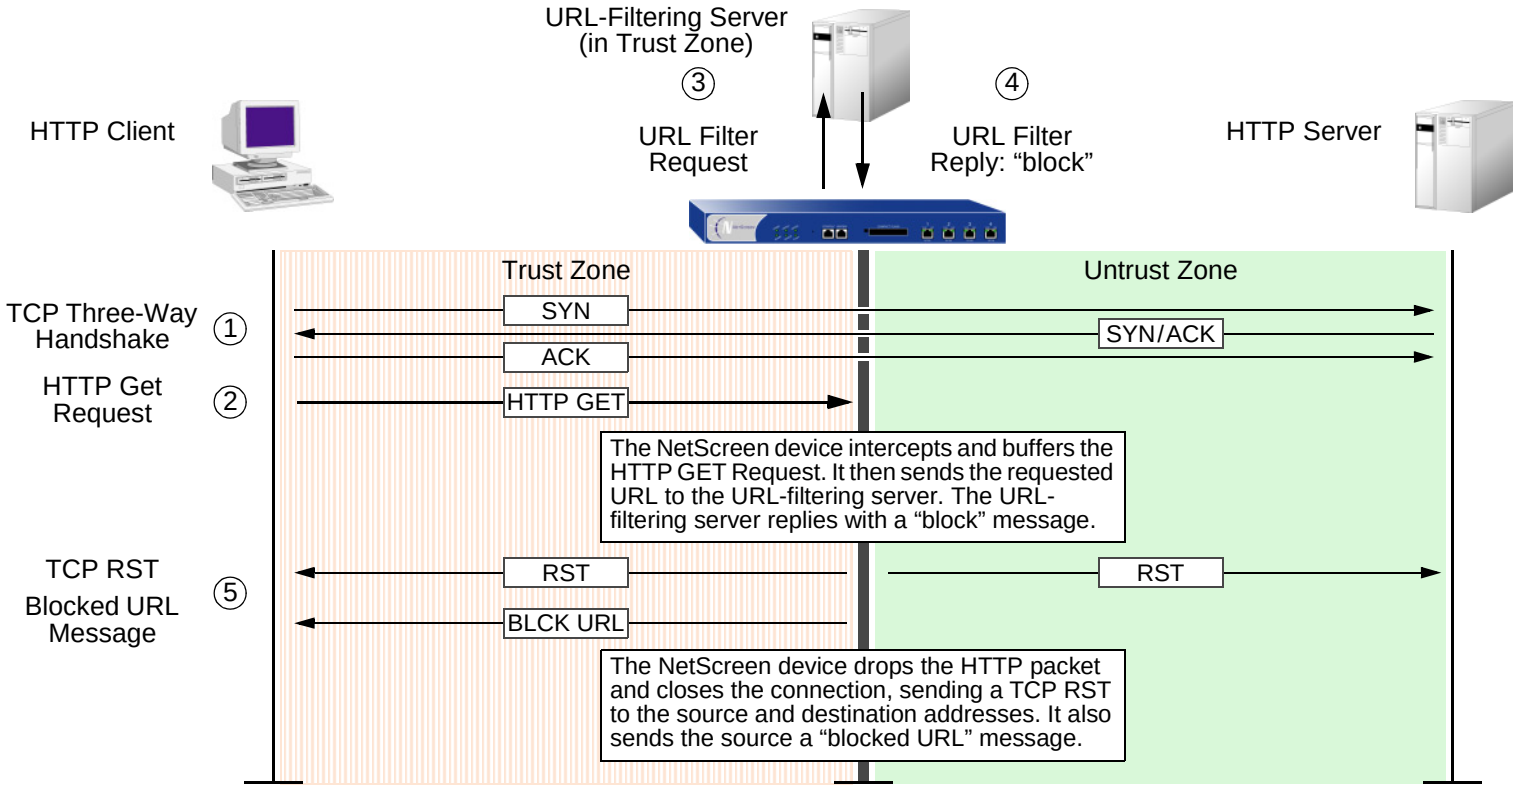

If the URL-filtering server permits access to the URL, the sequence of events in the HTTP connection attempt proceeds as follows:

A **Permitted** URL

set policy from trust to untrust any any http permit url-filter

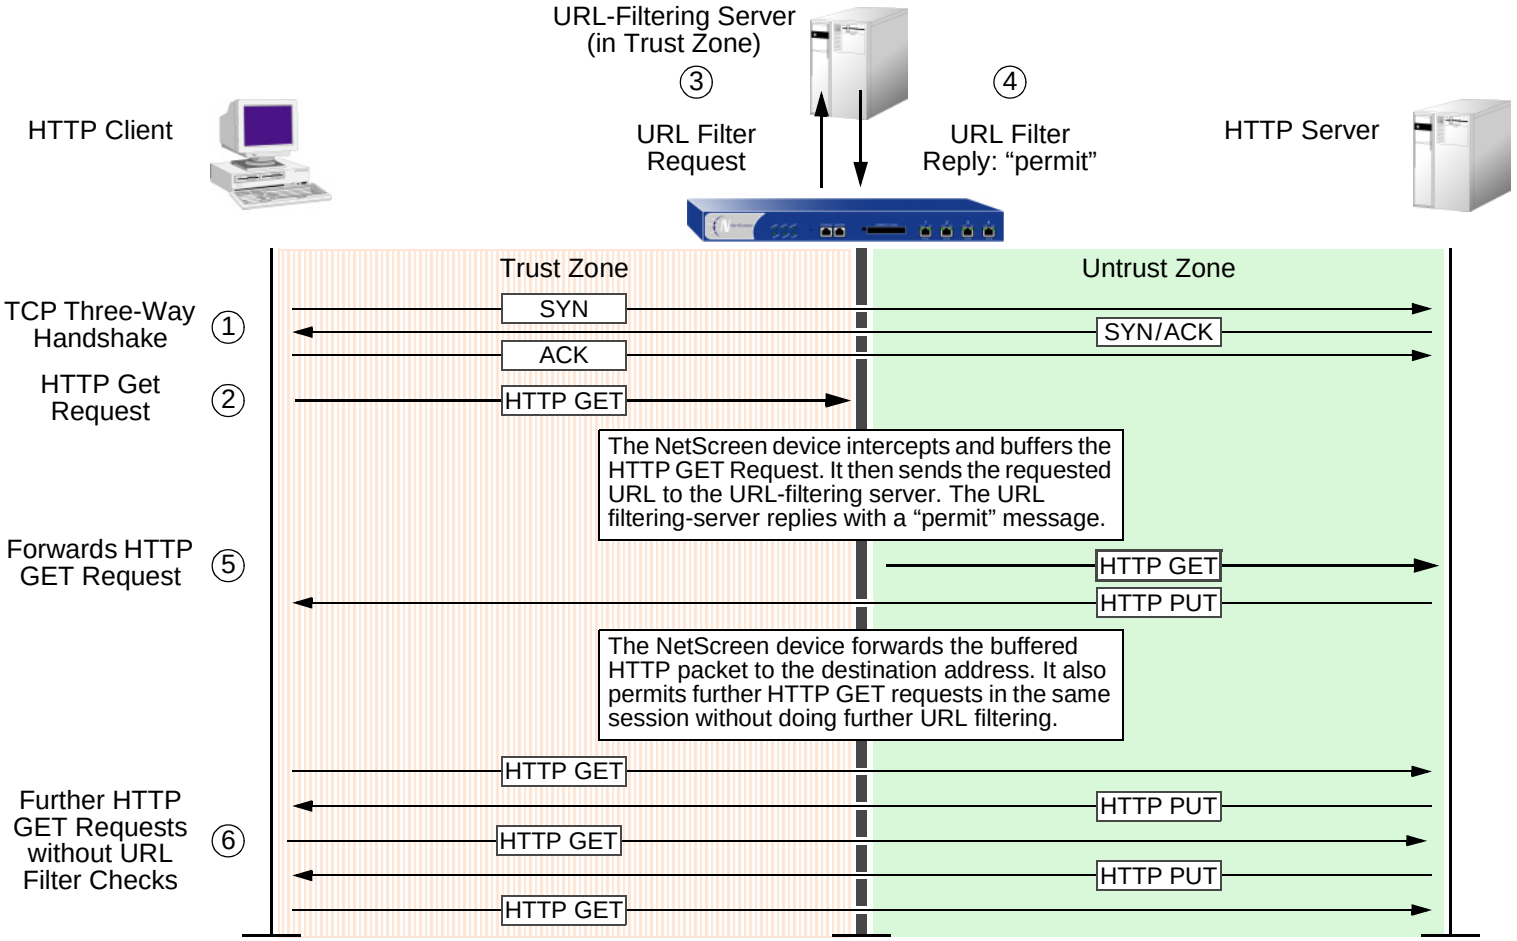

Using Websense, the administrator can do the following:

- Alter the URL-filtering database to block or allow access to selected sites
- Schedule different URL filtering profiles for different times of the day
- Download Websense Reporter logs of blocked or viewed URLs

Note: *For additional information about Websense, visit www.websense.com.*

NetScreen devices with virtual systems support up to eight different URL-filtering servers—one server reserved for the root system, and which can be shared with an unrestricted number of virtual systems; and seven URL-filtering servers for private use by the virtual systems. A root-level admin can configure the URL-filtering module at the root and virtual system (vsys) levels. A vsys-level admin can configure the URL module for his or her own vsys if that vsys has its own dedicated URL-filtering server. If the vsys-level admin uses the root URL-filtering server settings, that admin can see—but not modify—the root-level URL-filtering settings.

To configure a NetScreen device for URL filtering, you must perform the following tasks:

1. Set up communications with up to eight URL-filtering servers.
2. Define some system-level behavioral parameters. One set of parameters can apply to the root system and any vsys that shares the URL filtering configuration with the root system. Other sets can apply to virtual systems that have their own dedicated URL filtering server.
3. Activate URL filtering at the root and vsys levels.
4. Enable URL filtering in individual policies.

Details of these tasks are provided below.

1. Device-to-Device Communications

You first define settings for the Websense server and parameters for the behavior that you want the NetScreen device to take when applying URL filtering. If you configure these settings in the root system, they also apply to any virtual system that shares the URL-filtering configuration with the root system. For a vsys that has its own dedicated URL-filtering server, the root admin or vsys admin must configure the settings separately for that vsys.

The URL-filtering settings that you must define at the system level for device-to-device communications are as follows:

- **Websense Server Name:** The IP address or fully qualified domain name (FQDN) of the computer running the Websense server.
- **Websense Server Port:** The default port for Websense is 15868. If you have changed the default port on the Websense server you must also change it on the NetScreen device. Please see your Websense documentation for full details.
- **Source Interface:** The source from which the NetScreen device initiates URL filter requests to a Websense server when sending them through a VPN tunnel. (Note that the source interface is different from the outgoing interface, which is the egress interface for VPN traffic.) Typically, the URL-filtering server belongs in the Trust zone. However, if you want several NetScreen devices to access a single URL-filtering server, you might configure VPN tunnels from each remote device to the local NetScreen device protecting the server. From the remote peers' perspective, the server is in their Untrust zones, and they send URL filter requests to it through the tunnels.
- **Communication Timeout:** The time interval, in seconds, that the NetScreen device waits for a response from the Websense filter. If Websense does not respond within the time interval, the NetScreen device either blocks the request or allows it, as you choose (see below).

You can use the following CLI command to configure these settings:

```
set url server { ip_addr | dom_name } port_num timout_num
```

In the WebUI, enter these settings in their respective fields on the Screening > URL Filtering page.

2. System-Level Behavioral Parameters

Second, you define the behavior parameters that you want the system—root or vsys—to take when applying URL filtering. The behavior options are as follows:

- **Fail/Pass Mode:** If the NetScreen device loses contact with the Websense server, you can specify whether to **Block** or **Permit** all HTTP requests.

– **Blocked URL Message Type:** The source of the message the user receives when Websense blocks a site. If you select **NetPartners Websense**, the NetScreen device forwards the message it receives in the "block" response from the Websense server. When you select **NetScreen**, the NetScreen device sends the message that you have previously entered in the NetScreen Blocked URL Message field.

> ***Note:*** *If you select* ***NetScreen****, some of the functionality that Websense provides, such as redirection, are suppressed.*

– **NetScreen Blocked URL Message:** This is the message the NetScreen device returns to the user after blocking a site. You can use the message sent from the Websense server, or create a message (up to 500 characters) to be sent from the NetScreen device.

You can use the following CLI commands to configure these settings:

```
set url fail-mode { block | permit }
set url type { NetScreen | Websense }
set url message string
```

In the WebUI, enter these settings in their respective fields on the Screening > URL Filtering page.

3. System-Level Activation

When you complete the configuration, you must enable URL filtering at the system level. For a NetScreen device hosting virtual systems, you must enable URL filtering for each system that in which you want to apply it. For example, if you want the root system and a vsys to apply URL filtering, you must enable URL filtering in both the root system and that vsys.

You can use the following CLI command to activate and deactivate URL filtering at the system level:

```
set url config { disable | enable }
```

In the WebUI, select or clear the **Enable URL Filtering via Websense Server** check box on the Screening > URL Filtering page.

When you enable URL filtering at the system level, the NetScreen device checks all HTTP traffic to which policies (defined in that system) that require URL filtering apply by redirecting the HTTP requests to a Websense server. If you disable URL filtering at the system level, the NetScreen device ignores the URL filtering component in policies and treats them as simple “permit” policies.

4. Policy-Level Application

Finally, you configure the NetScreen device to contact the URL-filtering server on a per-policy basis.

You can use the following CLI command to enable URL filtering in a policy:

```
set policy from zone to zone src_addr dst_addr service permit url-filter
```

In the WebUI, select the **URL Filter** check box on the Advanced policy configuration page for the policy to which you want to apply URL filtering.

Note: *The NetScreen device reports the status of the Websense server. To update the status report, click the Server Status icon on the Screening > URL Filtering page in the WebUI.*

Example: URL Filtering Configuration

In this example, you configure the NetScreen device to work with a URL-filtering server at IP address 10.1.2.5, with port number 15868 (default). The URL-filtering server is in the Trust zone. You want to do URL filtering on all outbound HTTP traffic from hosts in the Trust zone to hosts in the Untrust zone. If the NetScreen device loses connectivity with the URL-filtering server, you want the NetScreen device to permit outbound HTTP traffic. When an HTTP client requests access to a prohibited URL, you want the NetScreen device to send the following message: "We're sorry, but the requested URL is prohibited. If this prohibition appears to be in error, contact ntwksec@mycompany.com."

The interface for the Untrust zone is ethernet3 and has IP address 1.1.1.1/24. The interface for the Trust zone is ethernet1 and has IP address 10.1.1.1/24. Both zones are in the trust-vr routing domain. Because the URL-filtering server is not in the immediate subnet of one of the NetScreen device interfaces, you add a route to it through ethernet1 and the internal router at 10.1.1.250.

WebUI

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Enter the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

2. URL-Filtering Server

Screening > URL Filtering: Enter the following, and then click **Apply**:

Enable URL Filtering via Websense Server: (select)

Websense Server Name: 10.1.2.5

Websense Server Port: 15868

Communication Timeout: 10 (seconds)

If connectivity to the Websense server is lost … all HTTP requests: Permit

Blocked URL Message Type: NetScreen

NetScreen Blocked URL Message: We're sorry, but the requested URL is prohibited. If this prohibition appears to be in error, contact ntwksec@mycompany.com.

3. Routes

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 1.1.1.250

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 10.1.2.0/24

Gateway: (select)

Interface: ethernet1

Gateway IP Address: 10.1.1.250

4. Policy

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Any

Service: HTTP

Action: Permit

> Advanced**:** Select the **URL Filter** check box, and then click **Return** to set the advanced options and return to the basic configuration page.

CLI

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
```

2. URL-Filtering Server

```
set url server 10.1.2.5 15868 10
set url fail-mode permit
set url type NetScreen
set url message “We’re sorry, but the requested URL is prohibited. If this
    prohibition appears to be in error, contact ntwksec@mycompany.com.”
set url config enable
```

3. Routes

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
set vrouter trust-vr route 10.1.2.0/24 interface ethernet1 gateway 10.1.1.250
```

4. Policy

```
set policy from trust to untrust any any http permit url-filter
save
```

5

Deep Inspection

You can enable Deep Inspection (DI) in policies to examine permitted traffic and take action if the DI module in ScreenOS finds attack signatures or protocol anomalies. The following sections in this chapter present the Deep Inspection elements that appear in policies and explains how to configure them:

- "Deep Inspection Overview" on page 124
- "Attack Object Database Server" on page 128
- "Attack Objects and Groups" on page 136
  - "Stateful Signatures" on page 138
  - "TCP Stream Signatures" on page 139
  - "Protocol Anomalies" on page 139
  - "Attack Object Groups" on page 140
- "Attack Actions" on page 142
- "Mapping Custom Services to Applications" on page 152
- "Customized Attack Objects and Groups" on page 156
  - "User-Defined Stateful SIgnature Attack Objects" on page 156
  - "TCP Stream Signature Attack Objects" on page 164

You can also enable Deep Inspection at the security zone level for HTTP components. These SCREEN options are explained in the final section of this chapter:

- "Granular Blocking of HTTP Components" on page 167
  - "ActiveX Controls" on page 167
  - "Java Applets" on page 168
  - "EXE Files" on page 168
  - "ZIP Files" on page 168

DEEP INSPECTION OVERVIEW

Deep Inspection (DI) is a mechanism for filtering the traffic permitted by the NetScreen firewall. Deep Inspection examines Layer 3 and 4 packet headers and Layer 7 application content and protocol characteristics in an effort to detect and prevent any attacks or anomalous behavior that might be present[1].

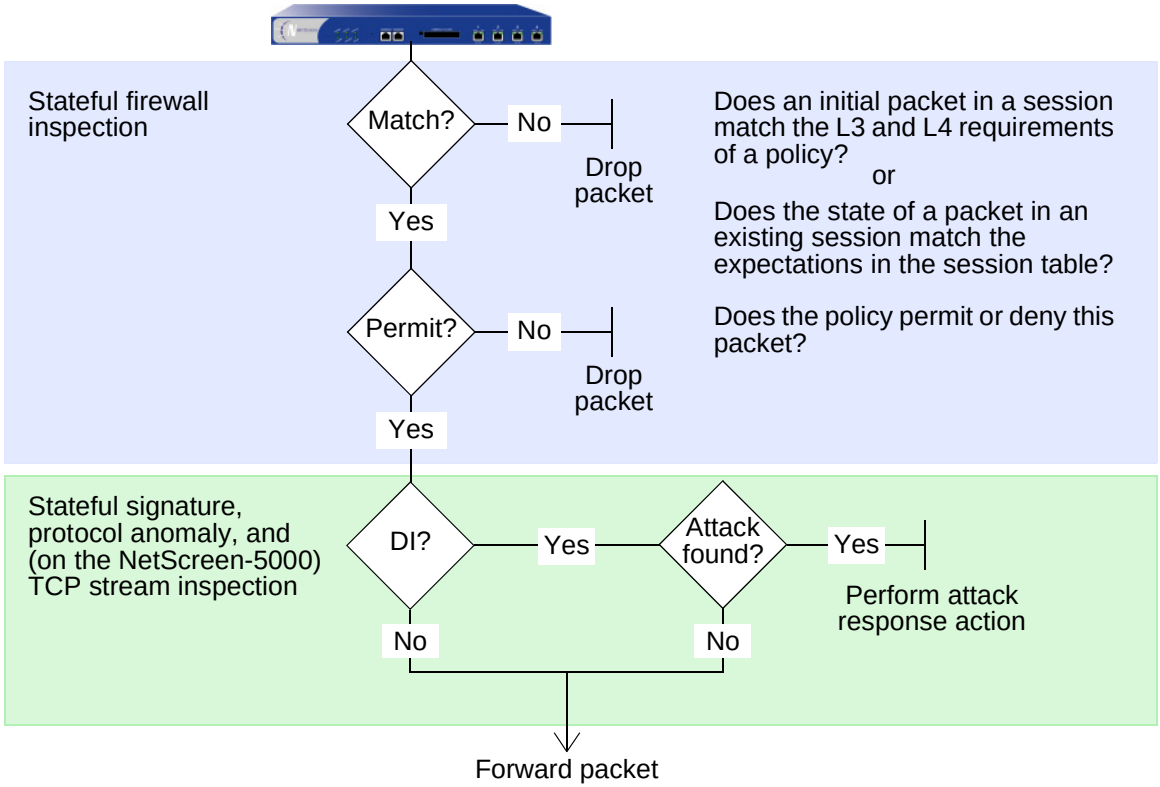

When the NetScreen device receives the first packet of a session, it inspects the source and destination IP addresses in the IP packet header (Layer 3 inspection) and the source and destination port numbers and protocol in

1. NetScreen detects anomalous traffic patterns at Layers 3 and 4 (IP and TCP) via SCREEN options set at the zone level, not the policy level. Examples of IP and TCP traffic-anomaly detection are "IP Address Sweep" on page 8, "Port Scanning" on page 10, and the various flood attacks described in "Network DoS Attacks" on page 45.

the TCP segment or UDP datagram header (Layer 4 inspection). If the Layer 3 and 4 components match the criteria specified in a policy, the NetScreen device then performs the specified action on the packet—permit, deny, or tunnel[2]. When the NetScreen device receives a packet for an established session, it compares it with the state information maintained in the session table to determine if it indeed belongs to the session.

If you have enabled Deep Inspection in the policy that applies to this packet and the policy action is “permit” or “tunnel”, then the NetScreen device further inspects it and its associated data stream for attack objects. Attack objects can be attack signatures or protocol anomalies, which you can either define yourself or download to the NetScreen device from an attack object database server[3]. (For more information, see “Attack Objects and Groups” on page 136 and “Customized Attack Objects and Groups” on page 156.) Based on the attack objects specified in the policy, the NetScreen device might perform the following inspections:

- Examine header values and payload data for stateful attack signatures
- Compare the format of the transmitted protocol with the standards specified in the RFCs and RFC extensions for that protocol to determine if someone has altered it, possibly for malicious purposes

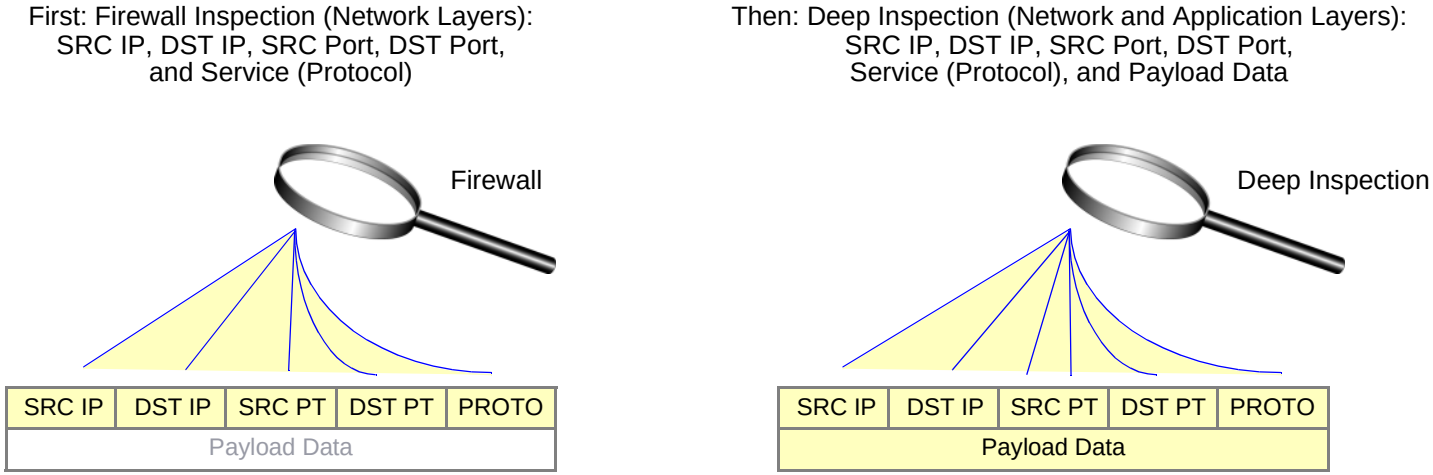

2. If the specified action is tunnel, the notion of permission is implied. Note that if you enable Deep Inspection (DI) in a policy whose action is tunnel, the NetScreen device performs the specified DI operations before encrypting an outbound packet and after decrypting an inbound packet.
3. The ability to download attack objects from the attack object database server requires that you first subscribe for the service. For more information, see “Registration and Activation of Signature Services” on page **2**-538.

If the NetScreen device detects an attack object, it performs the action specified in the DI component of the policy: close, close-client, close-server, drop, drop-packet, ignore, or none. If it does not find one of the specified attack objects, it forwards the packet. (For more information about attack actions, see "Attack Actions" on page 142.)

The following **set policy** command includes a DI component:

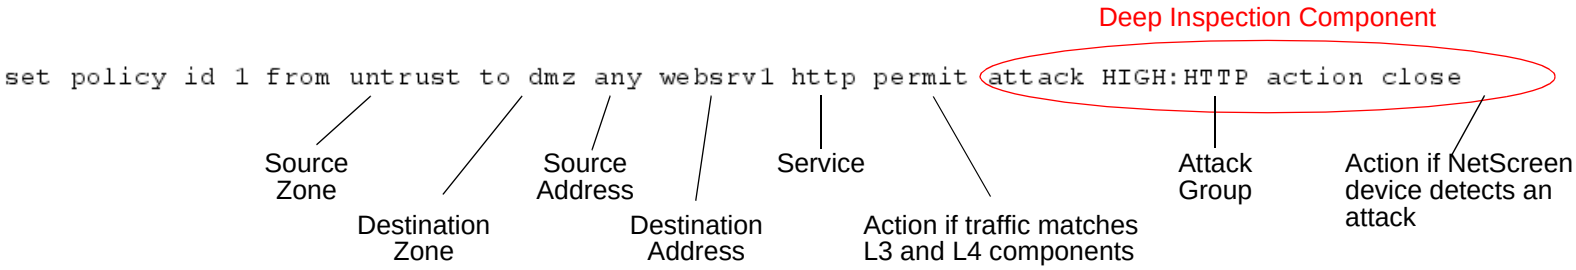

The above command directs the NetScreen device to permit HTTP traffic from any address in the Untrust zone to the destination address "websrv1" in the DMZ zone. It also instructs the NetScreen device to inspect all HTTP traffic permitted by this policy. If it finds any attack objects defined in the attack object group "HIGH:HTTP:ANOM", the NetScreen device closes the connection by dropping the packet and sending TCP RST notifications to the source and destination.

You can conceptually separate a **set policy** command into two parts—the core section and the DI component:

- The core section contains the source and destination zones, source and destination addresses, one or more services, and an action[4].
- The DI component instructs the NetScreen device to inspect traffic permitted by the core section of the policy for attack objects contained in one or more specified attack object groups. If the NetScreen device detects an attack object the NetScreen device then performs the action stated in the DI component.

4. You can optionally add other extensions to the core component of a **set policy** command: VPN and L2TP tunnel references, a schedule reference, address translation specifications, user authentication specifications, antivirus checking, logging, counting, and traffic management settings. Whereas these extensions are optional, the elements that constitute the core of a policy—source and destination zones, source and destination addresses, service (or services), and action—are required. (An exception to this is a global policy, in which no source and destination zones are specified: **set policy global *src_addr dst_addr service action***. For more information about global policies, see "Global Policies" on page **2**-201.)

It is possible to enter the context of an existing policy by using its ID number. For example:

```
ns-> set policy id 1
ns(policy:1)->
```

Note: *The command prompt changes to signal that a subsequent command is within a particular context.*

Entering a policy context is convenient if you want to enter several commands related to a single policy. For example, the following set of commands creates a policy that permits HTTP and HTTPS traffic from the any address in the Untrust to websrv1 and websrv2 in the DMZ zone and looks for medium, high, and critical HTTP stateful signature and protocol anomaly attacks:

```
ns-> set policy id 1 from untrust to dmz any websrv1 http permit attack
    CRITICAL:HTTP:ANOM action close
ns-> set policy id 1
ns(policy:1)-> set dst-address websrv2
ns(policy:1)-> set service https
ns(policy:1)-> set attack CRITICAL:HTTP:SIGS
ns(policy:1)-> set attack HIGH:HTTP:ANOM
ns(policy:1)-> set attack HIGH:HTTP:SIGS
ns(policy:1)-> exit
ns-> save
```

The above configuration permits both HTTP and HTTPS traffic, but only looks for attacks in HTTP traffic. To be able to add attack object groups with a policy context, you must first specify a DI attack and action in the top-level command. In the above example, you can add CRITICAL:HTTP:SIGS, HIGH:HTTP:ANOM, and HIGH:HTTP:SIGS attack object groups because you first configured the policy for Deep Inspection with the CRITICAL:HTTP:ANOM group.

Note: *You can specify only one attack action per policy. For information about the seven attack actions, see* *"Attack Actions" on page 142**.*

ATTACK OBJECT DATABASE SERVER

The attack object database contains all the predefined attack objects, organized into attack object groups by protocol and severity level. NetScreen stores the attack object database on a server at https://services.netscreen.com/restricted/sigupdates. To use the predefined attack objects, you must download the database from this server, load it on your NetScreen device, and then reference specific attack object groups in policies. To gain access to the attack object database server, you must first subscribe to the DI signature service for your NetScreen device. (For information on how to do that, see “Registration and Activation of Signature Services” on page **2**-538.)

Note: *ScreenOS contains a CA certificate for authenticating communication with the attack object database server.*

There are four ways to update the database:

- **Immediate Update:** With this option, you update the attack object database on the NetScreen device immediately with the database stored on the attack object database server. For this operation to work, you must first configure the attack object database server settings. (For an example, see “Example: Immediate Update” on page 129.)

  Note: *Before performing an immediate database update, you can use the* ***exec attack-db check*** *command to check if the attack object database on the server is more recent than the one on the NetScreen device.*

- **Automatic Update:** With this option, the NetScreen device downloads the attack object database directly to the NetScreen device at user-scheduled times if the database on the server is a newer version than that previously loaded on the NetScreen device. NetScreen updates the database on a regular basis with newly discovered attack patterns. Therefore, because of its changing nature, it behooves you to update your NetScreen device regularly too. For this operation to work, you must first configure the attack object database server settings. (For an example, see “Example: Automatic Updates” on page 130.)
- **Automatic Notification and Immediate Update:** With this option, the NetScreen device checks at user-scheduled times if the data on the attack object database server is more recent than that on the NetScreen device. If the data on the server is more recent, a notice appears on the Home page in the WebUI, and in the CLI after you log in to the NetScreen device. You can then enter the **exec attack-db**

update command or click the **Update Now** button on the Configuration > Update > Attack Signature page in the WebUI to save the database from the server to the NetScreen device. For the server-checking operation semi-automatic procedure to work, you must first configure the attack object database server settings. (For an example, see "Example: Automatic Notification and Immediate Update" on page 132.)

- **Manual Update:** With this option, you first use a Web browser to download the attack object database to a local directory or TFTP server directory. You can then load the database on the NetScreen device using either the WebUI (from the local directory) or CLI (from the TFTP server directory). (For an example, see "Example: Manual Update" on page 134.)

Example: Immediate Update

In this example, you save the attack object database (the attacks.bin file) from the attack object database server to the NetScreen device immediately. You use the default URL: https://services.netscreen.com/restricted/sigupdates. You do not have to set this URL for the database server. The NetScreen device uses it by default.

You do not set a schedule for updating the database on the NetScreen device. Instead you save the database from the server to the NetScreen device immediately.

Note: *This example assumes that you have already obtained and activated a subscription for the DI signature service for the NetScreen device. (For information about subscriptions, see "License Keys" on page* ***2*** *-536.)*

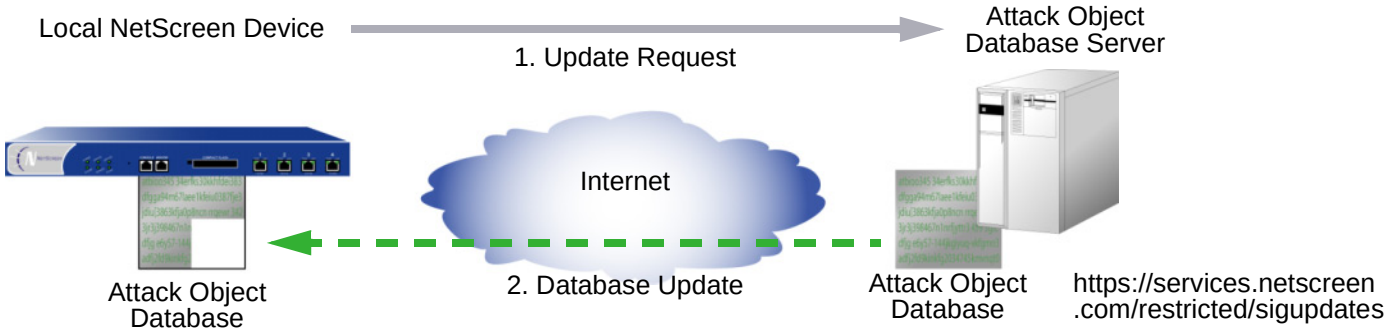

WebUI

Configuration > Update > Attack Signature: Click the **Update Now** button.

CLI

```
ns-> exec attack-db update
Loading attack database.............
Done.
Done.
Switching attack database...Done
Saving attack database to flash...Done.
ns->
```

Example: Automatic Updates

In this example, you set a schedule to update the database on the NetScreen device every Monday at 4:00 AM. At that scheduled time, the NetScreen device compares the version of the database on the server with that on the NetScreen device. If the version on the server is more recent, the NetScreen device automatically replaces its database with the newer version.

Note: *This example assumes that you have already obtained and activated a subscription for the DI signature service for the NetScreen device. (For information about subscriptions, see "License Keys" on page* ***2*** *-536.)*

You use the default URL: https://services.netscreen.com/restricted/sigupdates. You do not have to set this URL for the database server. The NetScreen device uses it by default.

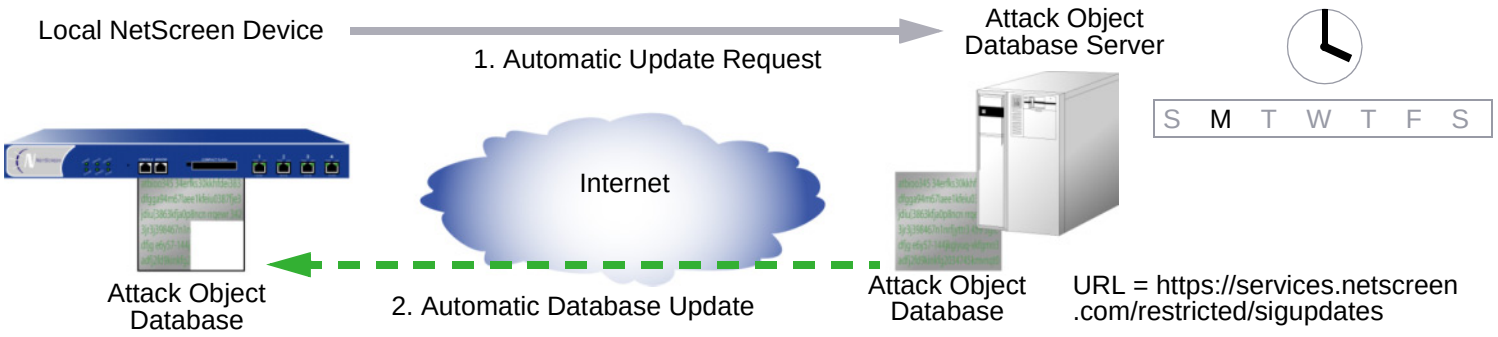

WebUI

Configuration > Update > Attack Signature: Enter the following, and then click **OK**:

Database Server: (leave empty)

Update Mode: Automatic Update

Schedule:

Weekly on: Monday[5]

Time (hh:mm): 04:00

CLI

```
set attack db mode update
set attack db schedule weekly monday 04:00
save
```

5. If you schedule updates on a monthly basis and the date you choose does not occur in a month (for example, 31 does not occur in several months), the NetScreen device uses the last possible date of the month in its place.

Example: Automatic Notification and Immediate Update

In this example, you set a schedule to check the database on the NetScreen device on a daily basis at 7:00 AM.

When you receive a notice that the database on the server has been updated, you click the **Update Now** button on the Configuration > Update > Attack Signature page in the WebUI or enter the **exec attack-db update** command to save the database from the server to the NetScreen device.

Note: *This example assumes that you have already obtained and activated a subscription for the DI signature service for the NetScreen device. (For information about subscriptions, see "License Keys" on page* ***2****-536.)*

You use the default URL: https://services.netscreen.com/restricted/sigupdates. You do not have to set this URL for the database server. The NetScreen device uses it by default.

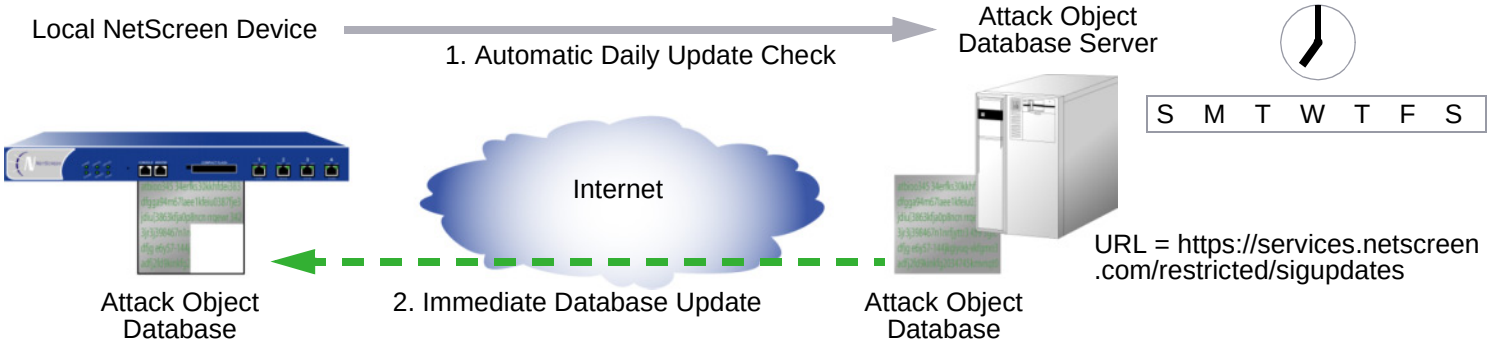

WebUI

1. Scheduled Database Checking

Configuration > Update > Attack Signature: Enter the following, and then click **OK**:

Database Server: (leave empty)
Update Mode: Automatic Notification
Schedule:
  Daily
Time (hh:mm): 07:00

2. Immediate Database Update

When you receive a notice that the attack database on the server is more current than the one on the NetScreen device, do the following:

Configuration > Update > Attack Signature: Click the **Update Now** button.

CLI

1. Scheduled Database Checking

```
set attack db mode notification
set attack db schedule daily 07:00
```

2. Immediate Database Update

When you receive a notice that the attack database on the server is more current than the one on the NetScreen device, do the following:

```
exec attack-db update
```

Example: Manual Update

In this example, you manually save the latest attack object database to the local directory “C:\netscreen\attacks-db” (if you want to use the WebUI to load the database) or C:\Program Files\TFTP Server (if you want to use the CLI to load it). You then load the database on the NetScreen device from your local directory[6].

For an automatic update, the NetScreen device automatically adds the following elements to the URL:

- Serial number of the NetScreen device
- Number of the major ScreenOS version running on the device
- Platform type

When you manually update the database, you must add these elements yourself. In this example, the serial number is 0043012001000213, the ScreenOS version is 5.0, and the platform is NetScreen-208 (ns200). Consequently, the resulting URL is:

https://services.netscreen.com//restricted/sigupdates/5.0/ns200/attacks.bin?sn=0043012001000213

Note: *This example assumes that you have already obtained and activated a subscription for the DI signature service for the NetScreen device. (For information about subscriptions, see “License Keys” on page* ***2****-536.)*

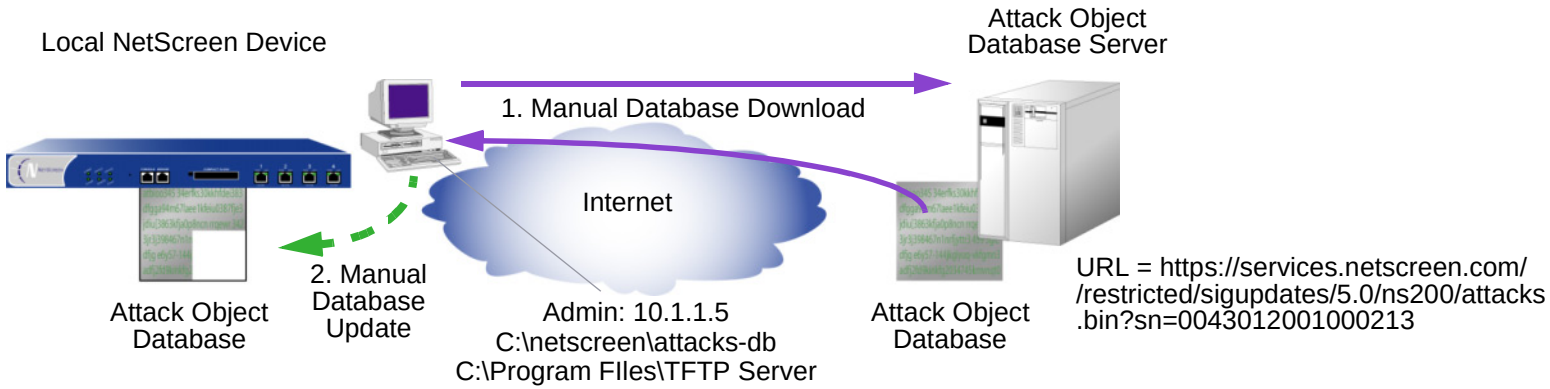

6. After downloading the attack object database, you can also post it on a local server and set it up for other NetScreen devices to access. The admins for the other devices must then change the database server URL to that of the new location. They can either enter the new URL in the Database Server field on the Configuration > Update > Attack Signature page or use the following CLI command: **set attack db server** *url_string*.

1. Database Download

Enter the following URL in the address field of your Web browser:

https://services.netscreen.com//restricted/sigupdates/5.0/ns200/attacks.bin?sn=0043012001000213

Save *attacks.bin* to the local directory “C:\netscreen\attacks-db” (for loading via the WebUI) or to your TFTP server directory C:\Program Files\TFTP Server (when you want to use the CLI to load it).

WebUI

2. Database Update

Configuration > Update > Attack Signature: Enter the following, and then click **OK**:

Deep Inspection Signature Update:

Load File: Enter **C:\netscreen\attacks-db\attacks.bin**, or click **Browse** and navigate to that directory, select **attacks.bin**, and then click **Open**.

CLI

2. Database Update

```
save attack-db from tftp 10.1.1.5 attacks.bin to flash
```

ATTACK OBJECTS AND GROUPS

Attack objects are stateful signatures and protocol anomalies that a NetScreen device uses to detect attacks aimed at compromising one or more hosts on a network. Attack objects are in groups organized by protocol type and then by severity. When you add Deep Inspection (DI) to a policy, the NetScreen device inspects the traffic that the policy permits for any patterns matching those in the referenced attack object group (or groups).

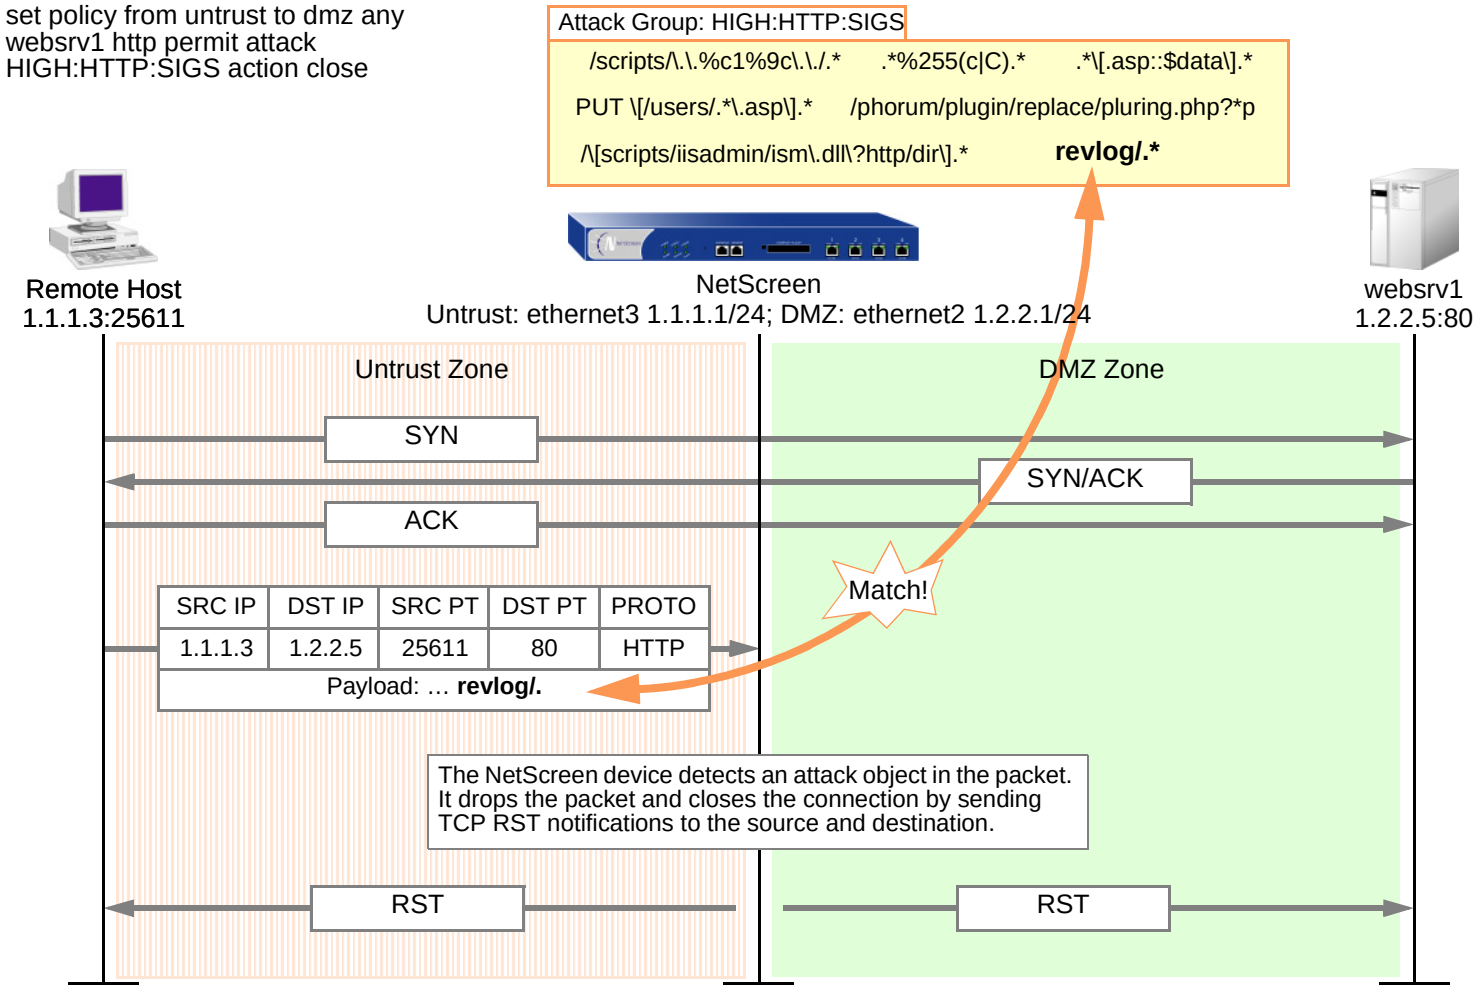

The attack object groups that you reference in the DI component must target the same service type that the policy permits. For example, if the policy permits SMTP traffic, the attack object group must aim at attacks on SMTP traffic. The following policy exemplifies a valid configuration:

✓
```
set policy id 2 from trust to untrust any any smtp permit attack CRIT:SMTP:SIGS
    action close
```

The next policy is erroneous because the policy permits SMTP traffic, but the attack object group is for POP3 traffic:

✗
```
set policy id 2 from trust to untrust any any smtp permit attack CRIT:POP3:SIGS
    action close
```

The second policy is misconfigured and, if implemented, would cause the NetScreen device to expend unnecessary resources inspecting SMTP traffic for POP3 attack objects that it could never find. If policy 2 permits both SMTP and POP3 traffic, you can configure the DI component to check for SMTP attack objects, POP3 attack objects, or for both.

```
set group service grp1
set group service grp1 add smtp
set group service grp1 add pop3
```
✓
```
set policy id 2 from trust to untrust any any grp1 permit attack
    CRIT:SMTP:SIGS action close
```
✓
```
set policy id 2 attack CRIT:POP3:SIGS
```

If the NetScreen device has access to http://help.netscreen.com/sigupdates/english, you can see the contents of all the predefined attack object groups and descriptions of the predefined attack objects. Open your Web browser, and enter one of the following URLs in the Address field:

http://help.netscreen.com/sigupdates/english/DNS.html

http://help.netscreen.com/sigupdates/english/FTP.html

http://help.netscreen.com/sigupdates/english/HTTP.html

http://help.netscreen.com/sigupdates/english/IMAP.html

http://help.netscreen.com/sigupdates/english/POP3.html

http://help.netscreen.com/sigupdates/english/SMTP.html

Each of the above URLs links to an HTML page containing a list of all the predefined attack objects—organized in groups by severity—for a particular protocol. To see a description of an attack object, click its name.

Stateful Signatures

An attack signature is a pattern that exists when a particular exploit is in progress[7]. The signature can be a Layer 3 or 4 traffic pattern, such as when one address sends lots of packets to different port numbers at another address (port scan), or a textual pattern, such as when a malicious URL string appears in the data payload of a single HTTP or FTP packet. The string can also be a specific segment of code or a specific value in the packet header. However, when searching for a textual pattern, the Deep Inspection (DI) module in a NetScreen device looks for more than just a signature in a packet; it looks for the signature in a particular portion of the packet (even if fragmented or segmented), in packets sent at a particular time in the life of the session, and sent by either the connection initiator or the responder.

When the DI module checks for a textual pattern, it considers the roles of the participants as client or server and monitors the state of the session to narrow its search to just those elements relevant to the exploit for which attackers use the pattern. Using contextual information to refine packet examination greatly reduces false alarms—or “false positives”—and avoids unnecessary processing. The term “stateful signatures” conveys this concept of looking for signatures within the context of the participants’ roles and session state.

To see the advantage of considering the context in which a signature occurs, note the way the NetScreen DI module examines packets when enabled to detect the EXPN Root attack. Attackers use the EXPN Root attack to expand and expose mailing lists on a mail server. To detect the EXPN Root attack, the NetScreen device searches for the signature “expn root” in the control portion of a Simple Mail Transfer Protocol (SMTP) session. The NetScreen device examines only the control portion because that is only where the attack can occur. If “expn root” occurs in any other portion of the session, it is not an attack.

Using a simple textual packet signature detection technique, the signature “expn root” triggers an alarm even if it appears in the data portion of the SMTP connection; that is, in the body of an e-mail message. If, for example, you were writing to a colleague about EXPN Root attacks, a simple packet signature detector would regard this as an attack. Using stateful signatures, the NetScreen DI module can distinguish between text strings that signal attacks and those that are harmless occurrences.

7. Because the NetScreen DI module supports regular expressions, it can use wildcards when searching for patterns. Thus, a single attack signature definition can apply to multiple attack pattern variations.

TCP Stream Signatures

Like a stateful signature, a TCP stream signature is a pattern that exists when an exploit is in progress. However, when the DI module examines traffic for stateful signatures, it searches only within specific contexts. When the DI module examines traffic for TCP stream signatures, it does so without regard for contexts. Another distinction between the two types of signatures is that although stateful signatures can be predefined or user-defined, TCP stream signatures must be user-defined. After you add a stream signature attack object to an attack object group, you can then reference that group in a policy that applies Deep Inspection. (For more about TCP stream signatures, see “TCP Stream Signature Attack Objects” on page 164.)

Note: *You can define TCP stream signatures on NetScreen-5000 series systems only.*

Protocol Anomalies

Attack objects that search for protocol anomalies detect traffic that deviates from the standards defined in RFCs and common RFC extensions. With signature attack objects, you must use a predefined pattern or create a new one; therefore, they can only detect known attacks. Protocol anomaly detection is particularly useful for catching new attacks or those attacks that cannot be defined by a textual pattern. ScreenOS supports protocol anomaly attack objects for the following protocols:

- DNS
- FTP
- HTTP
- IMAP
- POP3
- SMTP

Attack Object Groups

Predefined attack object groups contain attack objects for a specific protocol. For each protocol, the groups are separated into protocol anomalies and stateful signatures, and then roughly organized by severity. The three attack object group severity levels are critical, high, and medium:

Critical: Contains attack objects matching exploits that attempt to evade detection, cause a network device to crash, or gain system-level privileges.

High: Contains attack objects matching exploits that attempt to disrupt a service, gain user-level access to a network device, or activate a Trojan horse previously loaded on a device.

Medium: Contains attack objects matching exploits that detect reconnaissance efforts attempting to access vital information through directory traversal or information leaks.

Changing Severity Levels

Although attack object groups are classified by protocol and severity level (critical, high, medium), each attack object has its own specific severity level:

- Critical
- Very High
- High
- Medium
- Low
- Info

These severity levels are meaningful for NetScreen-Security Manager (NSM).

It is possible to override the default severity level of all attack objects in one or more attack object groups referenced in a policy. You do this at the policy level by entering the context of an existing policy and then assigning a new severity level.

The following shows how to change the severity level of the attack object groups referenced in a policy through the WebUI and CLI:

WebUI

Policies > Edit (for an existing policy): Do the following, and then click **OK**:

> Deep Inspection: Select a severity option in the Severity drop-down list, and then click **OK**.

CLI

```
ns-> set policy id number
ns(policy:number)-> set severity string
```

To return the severity level for each attack object to its original setting, you again enter the context of a policy and issue the following **unset policy** command:

WebUI

Policies > Edit (for an existing policy): Do the following, and then click **OK**:

> Deep Inspection: Select **Default** in the Severity drop-down list, and then click **OK**.

CLI

```
ns-> set policy id number
ns(policy:number)-> unset policy id number severity
```

ATTACK ACTIONS

When the NetScreen Deep Inspection (DA) module detects an attack, it then performs the action that you specify. The seven possibilities are as follows:

- **Close** (severs connection and sends RST to client and server[8])

  Use this option for TCP connections. The NetScreen device drops the connection and sends a TCP RST to both the client (source) and server (destination). Because the delivery of RST notifications is unreliable, by sending a RST to both client and server, there is a greater chance that at least one gets the RST and closes the session.

- **Close Client** (severs connection and sends RST to client)

  Use this option for outbound TCP connections from a protected client to an untrusted server. If, for example, the server sends a malicious URL string, the NetScreen device drops the connection and sends a RST only to the client for it to clear its resources while the server is left hanging.

- **Close Server** (severs connection and sends RST to server)

  Use this option for inbound TCP connections from an untrusted client to a protected server. If the client tries to launch an attack, the NetScreen device drops the connection and sends a TCP RST only to the server for it to clear its resources while the client is left hanging.

- **Drop** (severs connection without sending anyone a RST)

  Use this option for UDP or other non-TCP connections, such as DNS. The NetScreen device drops all packets in a session, but does not send a TCP RST.

- **Drop Packet** (drops a particular packet, but does not sever connection)

  This option drops the packet in which an attack signature or protocol anomaly occurs but does not terminate the session itself. Use this option to drop malformed packets without disrupting the entire session. For example, if the NetScreen device detects an attack signature or protocol anomaly from an AOL proxy, dropping everything would disrupt all AOL service. Instead, dropping just the packet stops the problem packet without stopping the flow of all the other packets.

8. The client is always the initiator of a session; that is, the source address in a policy. The server is always the responder, or the destination address.

- **Ignore** (after detecting an attack signature or anomaly, the NetScreen device makes a log entry and stops checking—or ignores—the remainder of the connection)

  If the NetScreen device detects an attack signature or protocol anomaly, it makes an event log entry but does not sever the session itself. Use this option to tweak false positives during the initial setup phase of your Deep Inspection (DI) implementation. Also, use this option when a service uses a standard port number for nonstandard protocol activities; for example, Yahoo Messenger uses port 25 (SMTP port) for non-SMTP traffic. The NetScreen device logs the occurrence once per session (so that it does not fill the log with false positives), but takes no action.

- **None** (no action)

  It is useful when first identifying attack types during the initial setup phase of your DI implementation. When the NetScreen device detects an attack signature or protocol anomaly, it makes an entry in the event log but takes no action on the traffic itself. The NetScreen device continues to check subsequent traffic in that session and make log entries if it detects other attack signatures and anomalies.

Example: Attack Actions – Close Server, Close, Close Client

In this example, there are three zones: Trust, Untrust, and DMZ. You have finished analyzing attacks and have concluded you need the following three policies:

- **Policy ID 1:** Permit HTTP, HTTPS, PING, and FTP-GET traffic from any address in the Untrust zone to the Web servers (websrv1 and websrv2) in the DMZ.

  Attack Settings for Policy ID 1:

  – CRITICAL:HTTP:ANOM, CRITICAL:HTTP:SIGS
  – HIGH:HTTP:ANOM, HIGH:HTTP:SIGS
  – MEDIUM:HTTP:ANOM, MEDIUM:HTTP:SIGS
  – CRITICAL:FTP:SIGS
  – Action: Close Server

  You choose to drop the connection and send a TCP RST notification only to the protected Web servers so they can terminate sessions and clear resources. You anticipate attacks coming from the Untrust zone.

- **Policy ID 2:** Permit HTTP, HTTPS, PING, and FTP traffic from any address in the Trust zone to the Web servers (websrv1 and websrv2) in the DMZ

  Attack Settings for Policy ID 2:

  – CRITICAL:HTTP:ANOM, CRITICAL:HTTP:SIGS
  – HIGH:HTTP:ANOM, HIGH:HTTP:SIGS
  – MEDIUM:HTTP:ANOM, MEDIUM:HTTP:SIGS
  – CRITICAL:FTP:SIGS
  – Action: Close

  You choose to drop the connection and send a TCP RST notification to both the protected clients and servers so they both can terminate their sessions and clear their resources regardless of the severity level of the attack.

- **Policy ID 3:** Permit FTP-GET, HTTP, HTTPS, PING traffic from any address in the Trust zone to any address in the Untrust zone

  Attack Settings for Policy ID 3:

  – CRITICAL:HTTP:ANOM, CRITICAL:HTTP:SIGS
  – HIGH:HTTP:ANOM, HIGH:HTTP:SIGS
  – MEDIUM:HTTP:ANOM, MEDIUM:HTTP:SIGS
  – CRITICAL:FTP:SIGS
  – Action: Close Client

  You choose to drop the connection and send a TCP RST notification to the protected clients so they both can terminate their sessions and clear their resources. In this case, you anticipate an attack coming from an untrusted HTTP or FTP server.

Although the policies permit HTTP, HTTPS, Ping, and FTP-Get or FTP, the NetScreen device activates Deep Inspection only for HTTP and FTP traffic. All zones are in the trust-vr routing domain.

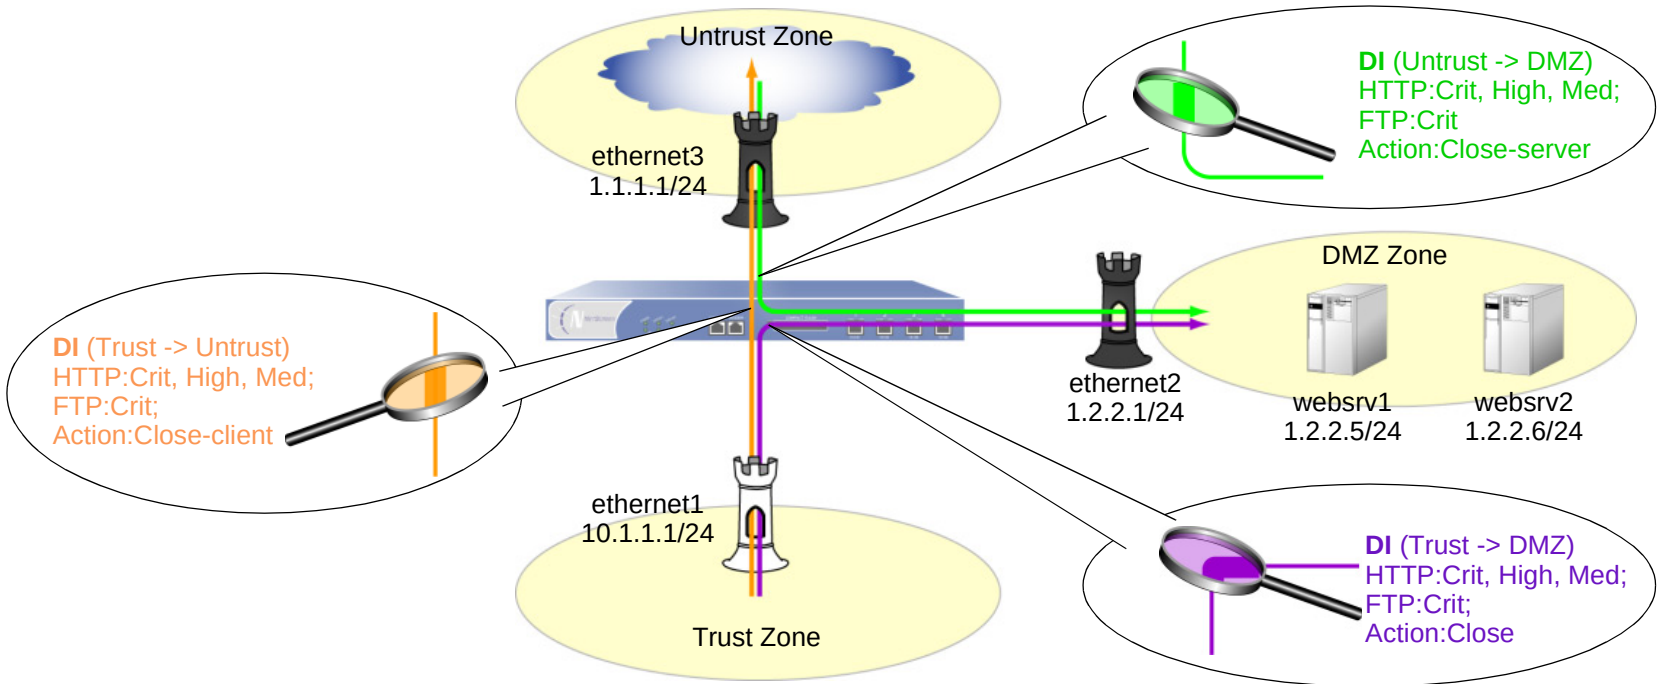

WebUI

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Enter the following, and then click **OK**:

Interface Mode: NAT

Service Options:

Management Services: (select all)

Other services: Ping

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust
Static IP: (select this option when present)
IP Address/Netmask: 1.1.1.1/24

Network > Interfaces > Edit (for ethernet2): Enter the following, and then click **OK**:

Zone Name: DMZ
Static IP: (select this option when present)
IP Address/Netmask: 1.2.2.1/24

2. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: websrv1
IP Address/Domain Name:
IP/Netmask: (select), 1.2.2.5/32
Zone: DMZ

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: websrv2
IP Address/Domain Name:
IP/Netmask: (select), 1.2.2.6/32
Zone: DMZ

3. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0
Gateway: (select)
Interface: ethernet3
Gateway IP Address: 1.1.1.250

4. Policy ID 1

Policies > (From: Untrust, To: DMZ) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), websrv1

> Click **Multiple**, select **websrv2**, and then click **OK** to return to the basic policy configuration page.

Service: HTTP

> Click **Multiple**, select **FTP-GET**, **HTTPS**, **PING**, and then click **OK** to return to the basic policy configuration page.

Action: Permit

> Click **Deep Inspection**, enter the following, and then click **OK** to return to the basic policy configuration page:

Action: Close Server

Use the **<<** button to move the following attack groups from the Available Members column to the Selected Members column:

CRITICAL:HTTP:ANOM

CRITICAL:HTTP:SIGS

HIGH:HTTP:ANOM

HIGH:HTTP:SIGS

MEDIUM:HTTP:ANOM

MEDIUM:HTTP:SIGS

CRITICAL:FTP:SIGS

5. Policy ID 2

Policies > (From: Trust, To: DMZ) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), websrv1

> Click **Multiple**, select **websrv2**, and then click **OK** to return to the basic policy configuration page.

Service: HTTP

> Click **Multiple**, select **FTP-GET**, **HTTPS**, **PING**, and then click **OK** to return to the basic policy configuration page.

Action: Permit

> Click **Deep Inspection**, enter the following, and then click **OK** to return to the basic policy configuration page:

Action: Close

Use the **<<** button to move the following attack groups from the Available Members column to the Selected Members column:

CRITICAL:HTTP:ANOM

CRITICAL:HTTP:SIGS

HIGH:HTTP:ANOM

HIGH:HTTP:SIGS

MEDIUM:HTTP:ANOM

MEDIUM:HTTP:SIGS

CRITICAL:FTP:SIGS

6. Policy ID 3

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Any

Service: HTTP

> Click **Multiple**, select **FTP-GET**, **HTTPS**, **PING**, and then click **OK** to return to the basic policy configuration page.

Action: Permit

> Click **Deep Inspection**, enter the following, and then click **OK** to return to the basic policy configuration page:

Action: Close Client

Use the **<<** button to move the following attack groups from the Available Members column to the Selected Members column:

CRITICAL:HTTP:ANOM

CRITICAL:HTTP:SIGS

HIGH:HTTP:ANOM

HIGH:HTTP:SIGS

MEDIUM:HTTP:ANOM

MEDIUM:HTTP:SIGS

CRITICAL:FTP:SIGS

CLI

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 manage
set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
set interface ethernet2 zone dmz
set interface ethernet2 ip 2.1.1.1/24
```

2. Addresses

```
set address dmz websrv1 1.2.2.5/32
set address dmz websrv2 1.2.2.6/32
```

3. Route

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
```

4. Policy ID 1

```
set policy id 1 from untrust to dmz any websrv1 http permit attack
    CRITICAL:HTTP:ANOM action close-server
set policy id 1
ns(policy:1)-> set dst-address websrv2
ns(policy:1)-> set service ftp-get
ns(policy:1)-> set service https
ns(policy:1)-> set service ping
ns(policy:1)-> set attack CRITICAL:HTTP:SIGS
ns(policy:1)-> set attack HIGH:HTTP:ANOM
ns(policy:1)-> set attack HIGH:HTTP:SIGS
ns(policy:1)-> set attack MEDIUM:HTTP:ANOM
ns(policy:1)-> set attack MEDIUM:HTTP:SIGS
ns(policy:1)-> set attack CRITICAL:FTP:SIGS
ns(policy:1)-> exit
```

5. Policy ID 2

```
set policy id 2 from trust to dmz any websrv1 http permit attack
    CRITICAL:HTTP:ANOM action close
set policy id 2
ns(policy:2)-> set dst-address websrv2
ns(policy:2)-> set service ftp
ns(policy:2)-> set service https
ns(policy:2)-> set service ping
ns(policy:2)-> set attack CRITICAL:HTTP:SIGS
ns(policy:2)-> set attack HIGH:HTTP:ANOM
ns(policy:2)-> set attack HIGH:HTTP:SIGS
ns(policy:2)-> set attack MEDIUM:HTTP:ANOM
ns(policy:2)-> set attack MEDIUM:HTTP:SIGS
ns(policy:2)-> set attack CRITICAL:FTP:SIGS
ns(policy:2)-> exit
```

6. Policy ID 3

```
set policy id 3 from trust to untrust any any http permit attack
    CRITICAL:HTTP:ANOM action close-client
set policy id 3
ns(policy:3)-> set service ftp-get
ns(policy:3)-> set service https
ns(policy:3)-> set service ping
ns(policy:3)-> set attack CRITICAL:HTTP:SIGS
ns(policy:3)-> set attack HIGH:HTTP:ANOM
ns(policy:3)-> set attack HIGH:HTTP:SIGS
ns(policy:3)-> set attack MEDIUM:HTTP:ANOM
ns(policy:3)-> set attack MEDIUM:HTTP:SIGS
ns(policy:3)-> set attack CRITICAL:FTP:SIGS
ns(policy:3)-> exit
save
```

MAPPING CUSTOM SERVICES TO APPLICATIONS

When using a custom service in a policy with a Deep Inspection (DI) component, you must explicitly specify the application that is running on that service so that the DI module can function properly. For example, if you create a custom service for FTP running on the nonstandard port number 2121, you can reference that custom service in a policy as follows:

```
set service ftp-custom protocol tcp src-port 0-65535 dst-port 2121-2121
set policy id 1 from untrust to trust any ftp-srv1 custom-ftp permit
```

However, if you add a DI component to a policy that references a custom service, the DI module cannot recognize the application because it is using a nonstandard port number.

```
set policy id 1 from untrust to trust any ftp-srv1 custom-ftp permit attack
    CRITICAL:FTP:SIGS action close-server
```

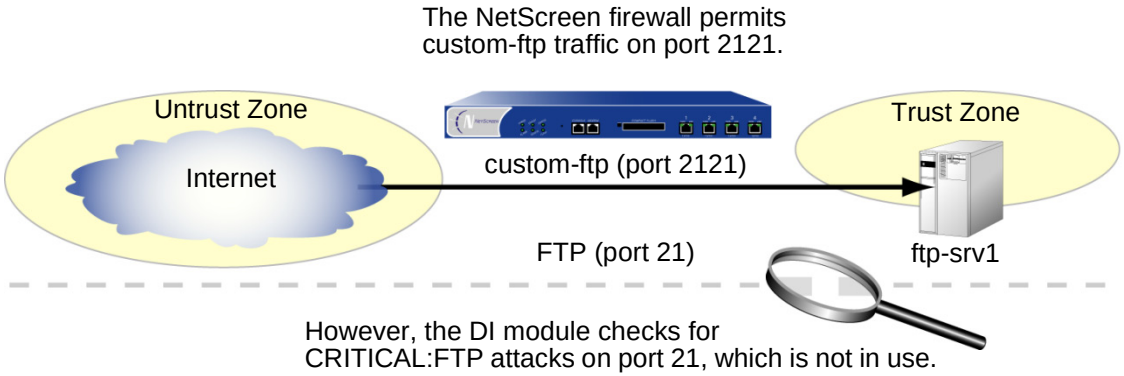

To avoid this problem, you must inform the DI module that the FTP application is running on port 2121. Essentially, you must map the protocol in the Application Layer to a specific port number in the Transport Layer. You can do this binding at the policy level:

```
set policy id 1 application ftp
```

When you map the FTP application to the custom service “custom-ftp” and configure DI to examine FTP traffic for the attacks defined in the CRITICAL:FTP:SIGS attack object group in a policy that references custom-ftp, the DI module perform its inspection on port 2121.

```
set policy id 1 from untrust to trust any ftp-srv1 custom-ftp permit attack
    CRITICAL:FTP:SIGS action close-server
set policy id 1 application ftp
```

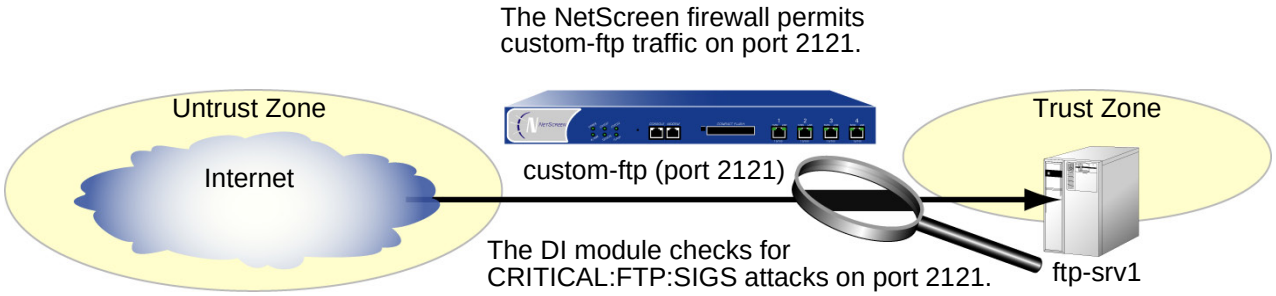

Example: Mapping an Application to a Custom Service

In this example, you define a custom service named “HTTP1” that uses destination port 8080. You map the HTTP application to the custom service for a policy permitting HTTP1 traffic from any address in the Untrust zone to a Web server named “server1” in the DMZ zone.

WebUI

1. Custom Service

Objects > Services > Custom > New: Enter the following, and then click **OK**:

Service Name: HTTP1

Transport Protocol: TCP (select)

Source Port Low: 0

Source Port High: 65535

Destination Port Low: 8080

Destination Port High: 8080

2. Address

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: server1

IP Address/Domain Name:

IP/Netmask: 1.2.2.5/32

Zone: DMZ

3. Policy

Policies > (From: Untrust, To: DMZ) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), server1

Service: HTTP1

Application: HTTP

Action: Permit

> Click **Deep Inspection**, enter the following, and then click **OK** to return to the basic policy configuration page:

Action: Close Server

Use the **<<** button to move the following attack groups from the Available Members column to the Selected Members column:

CRITICAL:HTTP:ANOM

CRITICAL:HTTP:SIGS

HIGH:HTTP:ANOM

HIGH:HTTP:SIGS

MEDIUM:HTTP:ANOM

MEDIUM:HTTP:SIGS

CLI

1. Custom Service

```
set service HTTP1 protocol tcp src-port 0-65535 dst-port 8080-8080
```

2. Address

```
set address dmz server1 1.2.2.5/32
```

3. Policy

```
ns-> set policy id 1 from untrust to dmz any server1 HTTP1 permit attack
    CRITICAL:HTTP:ANOM action close-server
ns-> set policy id 1
ns(policy:1)-> set attack CRITICAL:HTTP:SIGS
ns(policy:1)-> set attack HIGH:HTTP:ANOM
ns(policy:1)-> set attack HIGH:HTTP:SIGS
ns(policy:1)-> set attack MEDIUM:HTTP:ANOM
ns(policy:1)-> set attack MEDIUM:HTTP:SIGS
ns(policy:1)-> exit
ns-> set policy id 1 application http
save
```

CUSTOMIZED ATTACK OBJECTS AND GROUPS

You can define new attack objects and object groups to customize the Deep Inspection (DI) application to best meet your needs. The attack objects can be stateful signatures or—on the NetScreen-5000—TCP stream signatures.

User-Defined Stateful SIgnature Attack Objects

You can create a stateful signature attack object for FTP, HTTP, and SMTP. When creating an attack object, you perform the following steps:

- Name the attack object. (All user-defined attack objects must begin with “CS:”.)
- Set the context for the Deep Inspection search.
- Define the signature.
- Assign the attack object a severity level.

The following subsections examine the topics of contexts and signatures. For information on severity levels, which are used by NetScreen Security Manager 2004, see “Changing Severity Levels” on page 140.

Contexts

The context defines the location in the packet where the NetScreen DI module searches for a signature matching the attack object pattern. You can specify any of the following contexts:

- FTP Command: Sets the context as one of the FTP commands specified in RFC 959, “File Transfer Protocol (FTP)”
- FTP User Name: Sets the context as the name that a user enters when logging in to an FTP server
- HTTP URL Parsed: Sets the context as the “normalized” text string decoded from a unicode string
- SMTP Header From: Sets the context as the SMTP “From:” header
- SMTP Header To: Sets the context as the SMTP “To:” header
- SMTP Mail From: Sets the context as the SMTP ’MAIL FROM’ command line
- SMTP Recipient: Sets the context as the SMTP ’RCPT TO’ command line

You must then put a user-defined attack object in a user-defined attack object group for use in policies.

Note: *A user-defined attack object group can only contain user-defined attack objects. You cannot mix predefined and user-defined attack objects in the same attack object group.*

Signatures

When entering the text string for a signature, you can enter an alphanumeric string of ordinary characters to search for an exact character-to-character match, or you can use regular expressions to broaden the search for possible matches to sets of characters. ScreenOS supports the following regular expressions:

Purpose	Metacharacters	Example	Meaning
Direct binary match (octal)*	\O*octal_number*	\0162 Matches: 162	Exactly match this octal number: “162”.
Direct binary match (hexadecimal)†	\X*hexadecimal_number*\X	\X01 A5 00 00\X Matches: 01 A5 00 00	Exactly match these five hexadecimal numbers: “01 A5 00 00”.
Case-insensitive matches	\[*characters*\]	\[cat\] Matches: Cat, cAt, caT CAt, CaT, CAT cat, cAt	Match the characters in “cat” regardless of the case of each character.
Match any character	**.**	c **.** t Matches: cat, cbt, cct, … czt cAt, cBt, cCt, … cZt c1t, c2t, c3t, … c9t	Match “c-any character-t”.

Purpose	Metacharacters	Example	Meaning
Match the previous character 0 or more times, instead of only once	*	a*b+c Matches: bc bbc abc aaabbbbc	Match 0, 1, or multiple occurrences of “a”, followed by 1 or more occurrences of “b”, followed by one occurrence of “c”.
Match the previous character 1 or more times	+	a+b+c Matches: abc aabc aaabbbbc	Match 1 or more occurrences of “a”, followed by 1 or more occurrences of “b”, followed by one occurrence of “c”.
Match the previous character 0 times or 1 time	?	drop-?packet Matches: drop-packet droppacket	Match either “drop-packet” or “droppacket”.
Group expressions	( )		
Either the previous or the following character – typically used with ( )	|	(drop | packet) Matches: drop packet	Match either “drop” or “packet”.

Purpose	Metacharacters	Example	Meaning
Character range	[*start*-*end*]	[c-f]a(d | t) Matches: cad, cat dad, dat ead, eat fad, fat	Match everything that begins with “c”, “d”, “e”, or “f”, and has the middle letter “a” and the last letter “d” or “t”.
Negation of the following character	[^*character*]	[^0-9A-Z] Matches: a, b, c, d, e, … z	Match lowercase letters.

* Octal is a base-8 number system that uses only the digits 0-7.

† Hexadecimal is a base-16 number system that uses digits 0–9 as usual, and then the letters A–F representing hexadecimal digits with decimal values of 10-15.

Example: User-Defined Stateful Signature Attack Objects

In this example, you have an FTP server, a Web server, and a mail server in the DMZ zone. You define the following attack objects for the following uses:

Attack Object Name	You can use it to
cs:ftp-stor	stop someone from putting files on your FTP server.
cs:ftp-user-dm	deny FTP access to the user with the login name “dmartin”.
cs:url-index	block HTTP packets with a defined URL in any HTTP request.
cs:spammer	block e-mail from any e-mail address at “spam.com”.

You then organize them into a user-defined attack object group named “DMZ DI”, which you reference in a policy permitting traffic from the Untrust zone to the servers in the DMZ zone.

WebUI

1. Attack Object 1: ftp-stor

   Objects > Attacks > Custom > New: Enter the following, and then click **OK**:

   Attack Name: cs:ftp-stor

   Attack Context: FTP Command

   Attack Severity: Medium

   Attack Pattern: stor

2. Attack Object 2: ftp-user-dm

   Objects > Attacks > Custom > New: Enter the following, and then click **OK**:

   Attack Name: cs:ftp-user-dm

   Attack Context: FTP User Name

   Attack Severity: Low

   Attack Pattern: dmartin

3. Attack Object 3: url-index

Objects > Attacks > Custom > New: Enter the following, and then click **OK**:

Attack Name: cs:url-index

Attack Context: HTTP URL Parsed

Attack Severity: High

Attack Pattern: .*index.html.*

4. Attack Object 4: url-index

Objects > Attacks > Custom > New: Enter the following, and then click **OK**:

Attack Name: cs:spammer

Attack Context: SMTP From

Attack Severity: Info

Attack Pattern: .@spam.com

5. Attack Object Group

Objects > Attacks > Custom Group > New: Enter the following group name, move the following custom attack objects, and then click **OK**:

Group Name: CS:DMZ DI

Select **cs:ftp-stor** and use the **<<** button to move the address from the Selected Members column to the Selected Members column.

Select **cs:ftp-user-dm** and use the **<<** button to move the address from the Available Members column to the Selected Members column.

Select **cs:url-index** and use the **<<** button to move the address from the Available Members column to the Selected Members column.

Select **cs:spammer** and use the **<<** button to move the address from the Available Members column to the Selected Members column.

6. Policy

Policies > (From: Untrust, To: DMZ) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Any

Service: HTTP

> Click **Multiple**, select **FTP**, and then click **OK** to return to the basic policy configuration page.

Action: Permit

> Click **Deep Inspection**, enter the following, and then click **OK** to return to the basic policy configuration page:

Action: Close Server

Use the **<<** button to move the following attack groups from the Available Members column to the Selected Members column,

CS:DMZ DI

CLI

1. Attack Object 1: ftp-stor

```
set attack cs:ftp-stor ftp-command stor severity medium
```

2. Attack Object 2: ftp-user-dm

```
set attack cs:ftp-user-dm ftp-username dmartin severity low
```

3. Attack Object 3: url-index

```
set attack cs:url-index http-url-parsed index.html severity high
```

4. Attack Object 4: url-index

```
set attack cs:spammer smtp-from .@spam.com severity info
```

5. Attack Object Group

```
set attack group “CS:DMZ DI”
set attack group “CS:DMZ DI” add cs:ftp-stor
set attack group “CS:DMZ DI” add cs:ftp-user-dm
set attack group “CS:DMZ DI” add cs:url-index
set attack group “CS:DMZ DI” add cs:spammer
```

6. Policy

```
set policy id 1 from untrust to dmz any any http permit attack “CS:DMZ DI”
    action close-server
set policy id 1
ns(policy:1)-> set service ftp
ns(policy:1)-> exit
save
```

TCP Stream Signature Attack Objects

The stateful signatures are context-based within specific applications, such as an FTP user name or an SMTP header field. TCP stream signatures look for patterns anywhere in any kind of TCP traffic regardless of the application protocol in use.

Note: *You can define TCP stream signatures on NetScreen-5000 series systems only.*

Because there are no predefined TCP stream signature attack objects, you must define them. When creating a signature attack object, you define the following components:

- Attack object name (All user-defined attack objects must begin with “CS:”.)
- Object type (“stream”)
- Pattern definition
- Severity level

Example of a TCP Stream Signature Attack Object

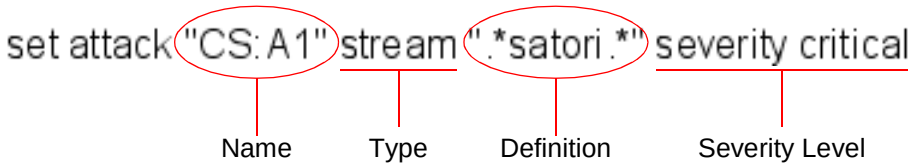

Example: User-Defined Stream Signature Attack Object

In this example, you define a stream signature object “.*satori.*”. You name it “CS:A1” and define its severity level as critical. Because a policy can reference only attack object groups, you create a group named “CS:Gr1”, and then add this object to it. Finally, you define a policy that references CS:Gr1 and that instructs the NetScreen device to sever the connection and send TCP RST to the client if the pattern appears in any traffic to which the policy applies.

WebUI

1. Stream Signature Attack Object

Objects > Attacks > Custom > New: Enter the following, and then click **OK**:

Attack Name: CS:A1

Attack Context: Stream

Attack Severity: Critical

Attack Pattern: .*satori.*

2. Stream Signature Attack Object Group

Objects > Attacks > Custom Group > New: Enter the following, and then click **OK**:

Group Name: CS:Gr1

Select **CS:A1** in the Available Members column and then click **<<** to move it to the Selected Members column.

3. Policy

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Any

Service: ANY

Action: Permit

> Click **Deep Inspection**, enter the following, and then click **OK** to return to the basic policy configuration page:

Action: Close Client

Select **CS:Gr1** in the Available Members column and then click **<<** to move it to the Selected Members column.

CLI

1. Stream Signature Attack Object

```
set attack “CS:A1” stream “.*satori.*” severity critical
```

2. Stream Signature Attack Group

```
set attack group “CS:Gr1”
set attack group “CS:Gr1” add “CS:A1”
```

3. Policy

```
set policy from trust to untrust any any any permit attack CS:Gr1 action
    close-client
save
```

GRANULAR BLOCKING OF HTTP COMPONENTS

A NetScreen device can selectively block ActiveX controls, Java applets, .zip files, and .exe files sent via HTTP. The danger that these components pose to the security of a network is that they provide a means for an untrusted party to load and then control an application on hosts in a protected network.

When you enable the blocking of one or more of these components in a security zone, the NetScreen device examines every HTTP header that arrives at an interface bound to that zone. It checks if the content type listed in the header indicates that any of the targeted components are in the packet payload. If the content type is Java, .exe, or .zip and you have configured the NetScreen device to block those HTTP component types, the NetScreen device blocks the packet. If the content type lists only “octet stream” instead of a specific component type, then the NetScreen device examines the file type in the payload. If the file type is Java, .exe, or .zip and you have configured the NetScreen device to block those component types, the NetScreen device blocks the packet.

When you enable the blocking of ActiveX controls, the NetScreen device blocks all HTTP packets containing any type of HTTP component in its payload—ActiveX controls, Java applets, .exe files, or .zip files.

Note: *When ActiveX-blocking is enabled, the NetScreen device blocks Java applets, .exe files, and .zip files whether they are contained within an ActiveX control or not.*

ActiveX Controls

Microsoft ActiveX technology provides a tool for Web designers to create dynamic and interactive Web pages. ActiveX controls are components that allow different programs to interact with each other. For example, ActiveX allows your Web browser to open a spreadsheet or display your personal account from a backend database. ActiveX components might also contain other components such as Java applets, or files such as .exe and .zip files.

When you visit an ActiveX-enabled Web site, the site prompts you to download ActiveX controls to your computer. Microsoft provides a pop-up message displaying the name of the company or programmer who authenticated the ActiveX code that is offered for download. If you trust the source of the code, you can proceed to download the controls. If you distrust the source, you can refuse them.

If you download an ActiveX control to your computer, it can then do whatever its creator designed it to do. If it is malicious code, it can now reformat your hard drive, delete all your files, send all your personal e-mail to your boss, and so on.

Java Applets

Serving a similar purpose as ActiveX, Java applets also increase the functionality of Web pages by allowing them to interact with other programs. You download Java applets to a Java Virtual Machine (VM) on your computer. In the initial version of Java, the VM did not allow the applets to interact with other resources on your computer. Starting with Java 1.1, some of these restrictions were relaxed to provide greater functionality. As a result, Java applets can now access local resources outside the VM. Because an attacker can program Java applets to operate outside the VM, they pose the same security threat as ActiveX controls do.

EXE Files

If you download and run an executable file (that is, a file with a .exe extension) obtained off the Web, you cannot guarantee that the file is uncontaminated. Even if you trust the site from which you downloaded it, it is possible that somebody sniffing download requests from that site has intercepted your request and responded with a doctored .exe file that contains malicious code.

ZIP Files

A zip file (that is, a file with a .zip extension) is a type of file containing one or more compressed files. The danger of downloading a .exe file presented in the previous section about .exe files applies to .zip files, because a .zip file can contain one or more .exe files.

Example: Blocking Java Applets and .exe Files

In this example, you block any HTTP traffic containing Java applets and .exe files in packets arriving at an Untrust zone interface.

WebUI

Screening > Screen (Zone: Untrust): Select **Block Java Component** and **Block EXE Component**, and then click **Apply**.

CLI

```
set zone untrust screen java
set zone untrust screen exe
save
```

1

Suspicious Packet Attributes

As shown in the other chapters in this volume, attackers can craft packets to perform reconnaissance or launch denial-of-service (DoS) attacks. Sometimes it is unclear what the intent of a crafted packet is, but the very fact that it is crafted suggests that its being put to some kind of insidious use. All of the SCREEN options presented in this chapter block suspicious packets that might contain hidden threats:

- "ICMP Fragments" on page 2
- "Large ICMP Packets" on page 4
- "Bad IP Options" on page 6
- "Unknown Protocols" on page 8
- "IP Packet Fragments" on page 10
- "SYN Fragments" on page 12

ICMP FRAGMENTS

Internet Control Message Protocol (ICMP) provides error reporting and network probe capabilities. Because ICMP packets contain very short messages, there is no legitimate reason for ICMP packets to be fragmented. If an ICMP packet is so large that it must be fragmented, something is amiss. When you enable the ICMP Fragment Protection SCREEN option, the NetScreen device blocks any ICMP packet with the More Fragments flag set, or with an offset value indicated in the offset field.

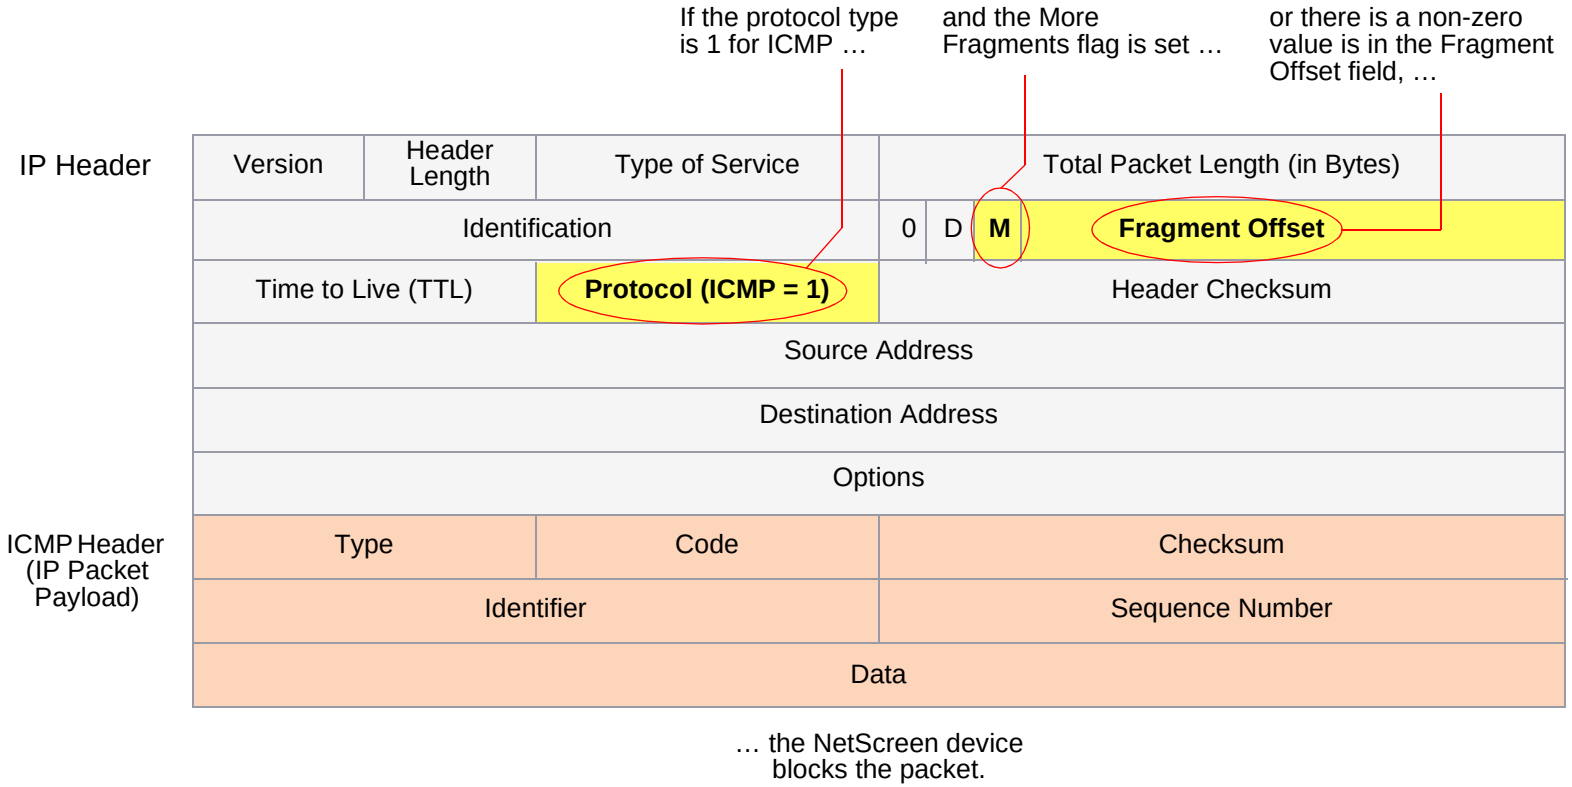

To block fragmented ICMP packets, do either of the following, where the specified security zone is the one from which the fragments originate:

WebUI

Screening > Screen (Zone: select a zone name): Select **ICMP Fragment Protection**, and then click **Apply**.

CLI

```
set zone zone screen icmp-fragment
```

LARGE ICMP PACKETS

As stated in the previous section "ICMP Fragments" on page 2, Internet Control Message Protocol (ICMP) provides error reporting and network probe capabilities. Because ICMP packets contain very short messages, there is no legitimate reason for large ICMP packets. If an ICMP packet is unusually large, something is wrong. For example, the Loki program uses ICMP as a channel for transmitting covert messages. The presence of large ICMP packets might expose a compromised machine acting as a Loki agent. It might also indicate some other kind of shifty activity.

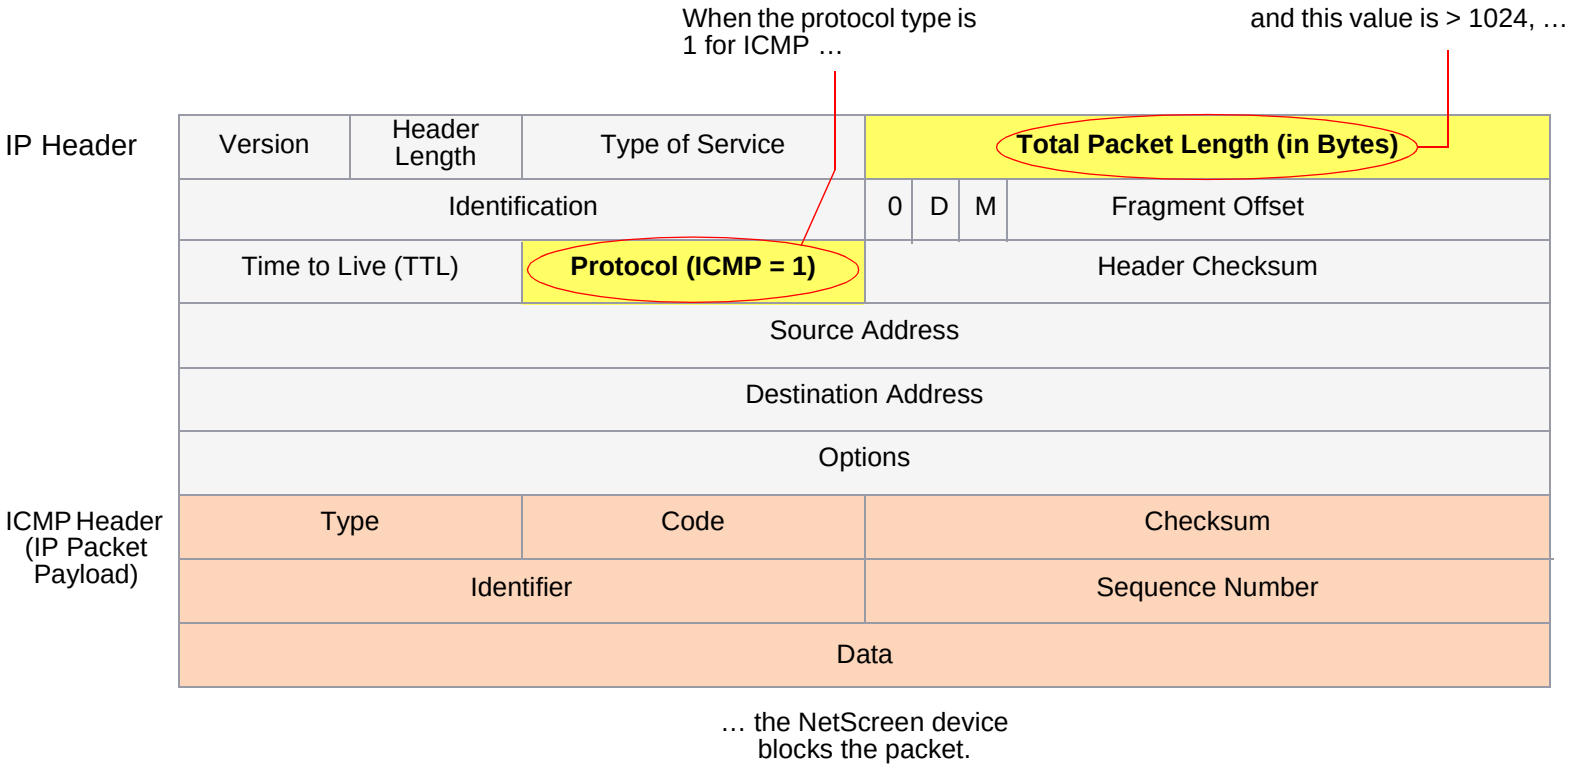

When you enable the Large Size ICMP Packet Protection SCREEN option, the NetScreen device checks drops ICMP packets with a length greater than 1024 bytes.

To block large ICMP packets, do either of the following, where the specified security zone is the one from which the ICMP packets originate:

WebUI

Screening > Screen (Zone: select a zone name): Select **Large Size ICMP Packet (Size > 1024) Protection**, and then click **Apply**.

CLI

```
set zone zone screen icmp-large
```

BAD IP OPTIONS

The Internet Protocol standard “RFC 791, Internet Protocol” specifies a set of eight options that provide special routing controls, diagnostic tools, and security. Although the original, intended uses for these options served worthy ends, people have figured out ways to twist these options to accomplish less commendable objectives. (For a summary of the exploits that attackers can wreak from IP options, see “Network Reconnaissance Using IP Options” on page 12.)

Either intentionally or accidentally, attackers sometimes misconfigure IP options, producing either incomplete or malformed fields. Regardless of the intentions of the person who crafted the packet, the misformatting is anomalous and potentially harmful to the intended recipient.

IP Header

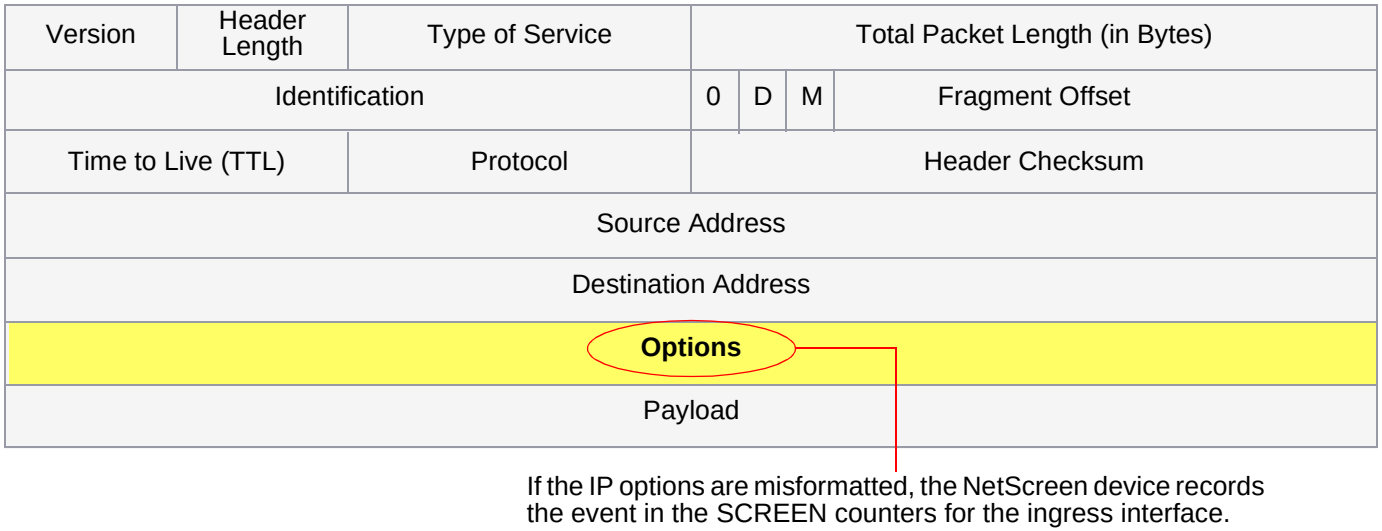

When you enable the Bad IP Option Protection SCREEN option, the NetScreen device blocks packets when any IP option in the IP packet header is incorrectly formatted. The NetScreen device records the event in the event log.

To detect and block IP packets with incorrectly formatted IP options, do either of the following, where the specified security zone is the one from which the packets originate:

WebUI

Screening > Screen (Zone: select a zone name): Select **Bad IP Option Protection**, and then click **Apply**.

CLI

```
set zone zone screen ip-bad-option
```

UNKNOWN PROTOCOLS

These protocol types with ID numbers of 135 or greater are reserved and undefined at this time. Precisely because these protocols are undefined, there is no way to know in advance if a particular unknown protocol is benign or malicious. Unless your network makes use of a non-standard protocol with an ID number of 135 or greater, a cautious stance is to block such unknown elements from entering your protected network.

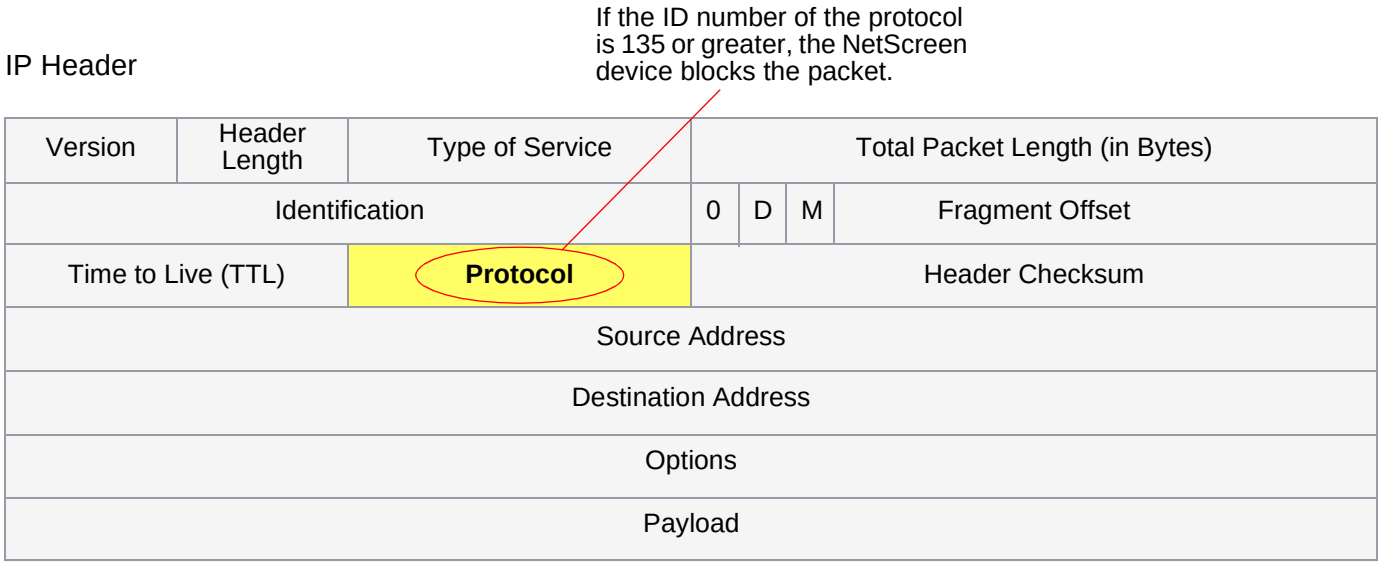

When you enable the Unknown Protocol Protection SCREEN option, the NetScreen device drops packets when the protocol field is contains a protocol ID number of 135 or greater.

To drop packets using an unknown protocol, do either of the following, where the specified security zone is the one from which the packets originate:

WebUI

Screening > Screen (Zone: select a zone name): Select **Unknown Protocol Protection**, and then click **Apply**.

CLI

```
set zone zone screen unknown-protocol
```

IP PACKET FRAGMENTS

As packets traverse different networks, it is sometimes necessary to break a packet into smaller pieces (fragments) based upon the maximum transmission unit (MTU) of each network. IP fragments might contain an attacker's attempt to exploit the vulnerabilities in the packet reassembly code of specific IP stack implementations. When the victim receives these packets, the results can range from processing the packets incorrectly to crashing the entire system.

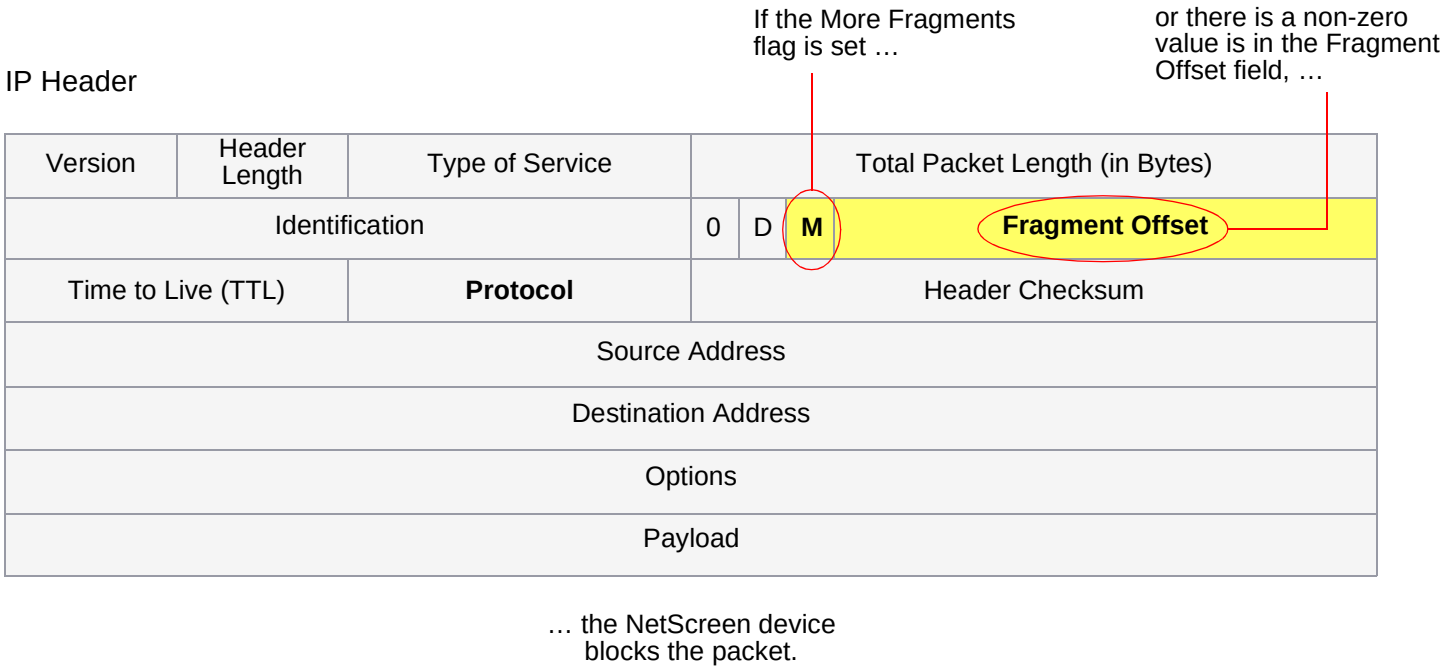

When you enable the NetScreen device to deny IP fragments on a security zone, the device blocks all IP packet fragments that it receives at interfaces bound to that zone.

To drop fragmented IP packets, do either of the following, where the specified security zone is the one from which the fragments originate:

WebUI

Screening > Screen (Zone: select a zone name): Select **Block Fragment Traffic**, and then click **Apply**.

CLI

```
set zone zone screen block-frag
```

SYN FRAGMENTS

The Internet Protocol (IP) encapsulates a Transmission Control Protocol (TCP) SYN segment in the IP packet that initiates a TCP connection. Because the purpose of this packet is to initiate a connection and invoke a SYN/ACK segment in response, the SYN segment typically does not contain any data. Because the IP packet is small, there is no legitimate reason for it to be fragmented. A fragmented SYN packet is anomalous, and as such suspect. To be cautious, block such unknown elements from entering your protected network.

When you enable the SYN Fragment Detection SCREEN option, the NetScreen device detects packets when the IP header indicates that the packet has been fragmented and the SYN flag is set in the TCP header. The NetScreen device records the event in the SCREEN counters list for the ingress interface.

To drop IP packets containing SYN fragments, do either of the following, where the specified security zone is the one from which the packets originate:

WebUI

Screening > Screen (Zone: select a zone name): Select **SYN Fragment Protection**, and then click **Apply**.

CLI

```
set zone zone screen syn-frag
```

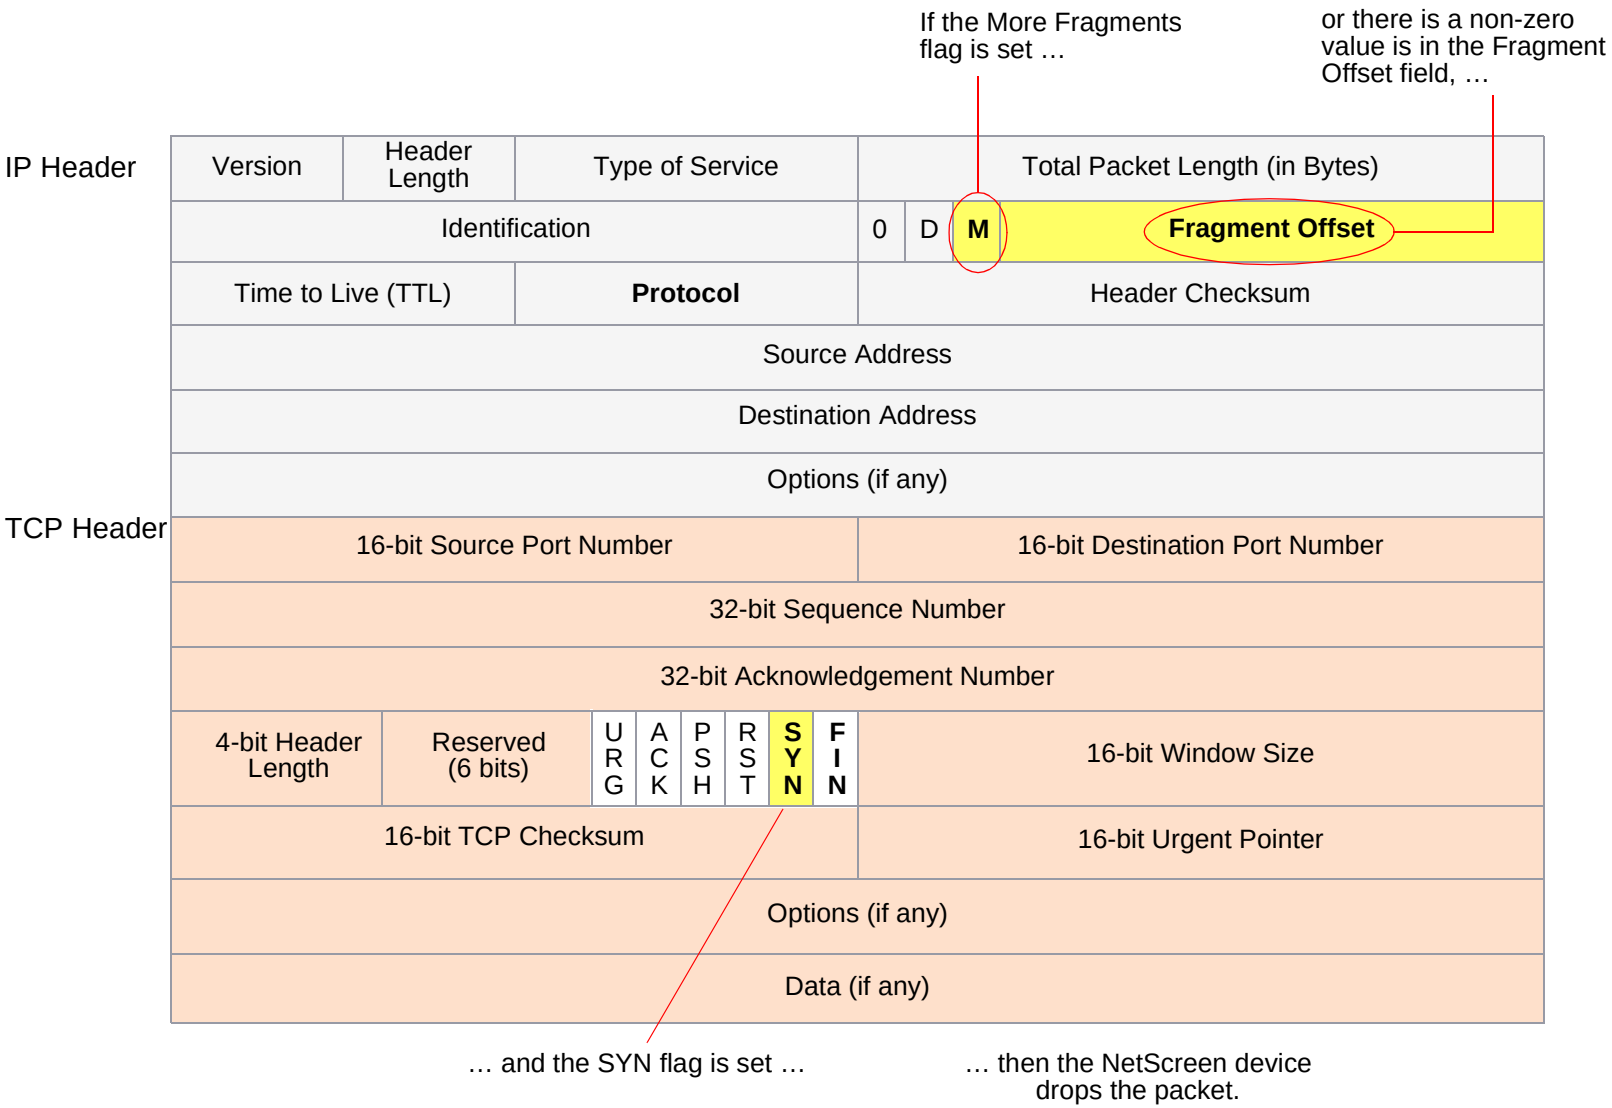
If the More Fragments flag is set …
or there is a non-zero value is in the Fragment Offset field, …
IP Header
Version
Header Length
Type of Service
Total Packet Length (in Bytes)
Identification
0
D
M
Fragment Offset
Time to Live (TTL)
Protocol
Header Checksum
Source Address
Destination Address
Options (if any)
TCP Header
16-bit Source Port Number
16-bit Destination Port Number
32-bit Sequence Number
32-bit Acknowledgement Number
4-bit Header Length
Reserved (6 bits)
URG
ACK
PSH
RST
SYN
FIN
16-bit Window Size
16-bit TCP Checksum
16-bit Urgent Pointer
Options (if any)
Data (if any)
… and the SYN flag is set …
… then the NetScreen device drops the packet.

Index

A

ActiveX controls, blocking 167
address sweep 8
aggressive aging 40–42
ALG 73
antivirus objects
  See AV objects
antivirus scanning
  See AV scanning
application layer gateway
  See ALG
attack actions 142–151
  close 142
  close client 142
  close server 142
  drop 142
  drop packet 142
  ignore 143
  none 143
attack object database 128–135
  auto notification and manual update 128, 132
  automatic update 128, 130
  changing the default URL 134
  immediate update 128, 129
  manual update 129, 134
attack object groups 140
  changing severity 140
  severity levels 140
attack objects 125
  protocol anomalies 139
  stateful signatures 138
  TCP stream signatures 164
attack protection
  policy level 5
  security zone level 5
attacks
  common objectives 1
  detection and defense options 3–5
  ICMP flood 59
  ICMP fragments 172
  IP packet fragments 180
  Land attack 63
  large ICMP packets 174
  Ping of Death 65
  stages of 2
  SYN flood 45–51
  SYN fragments 182–183
  Teardrop 67
  UDP flood 61
  unknown MAC addresses 51
  unknown protocols 178
  WinNuke 69
AV objects 93–101
  port number 94
  states 93
  timeout 94
AV scanning 76–112
  application 102
  AV objects 93–101
  decompression 85
  external AV scanner 90–112
  external, CSP resources 95
  external, HTTP 92
  external, SMTP 91
  fail-mode 95
  fail-mode threshold 96
  HTTP keep-alive 96
  HTTP trickling 97
  HTTP webmail 80
  internal AV scanner 77–89
  internal, HTTP 79
  internal, POP3 78
  internal, SMTP 77
  internal, subscription 81
  InterScan VirusWall 90
  max TCP connections 94
  multiple AV objects 106

C

character types, ScreenOS supported x
CLI
  conventions vi
content filtering 71–121
Content Scanning Protocol
  See CSP
conventions
  CLI vi
  illustration ix
  names x
  WebUI vii
CSP 90

D

DDoS 35
decompression, AV scanning 85
Deep Inspection 140–163
  attack actions 142–151
  attack object database 128–135
  attack object groups 140
  attack objects 125
  changing severity 140
  context 156
  custom attack objects 156
  custom services 152–155
  custom signatures 157–163
  protocol anomalies 139
  regular expressions 157–159
  stateful signatures 138
Denial-of-Service
  See DoS
DoS 35–70
  firewall 36–44
  network 45–63
  OS-specific 65–70
  session table flood 36
drop-no-rpf-route 23
dynamic packet filtering 3

E

evasion 22–33
exe files, blocking 168

exploits
See attacks

F

fail/pass mode, URL filtering 116
fail-mode 95
threshold 96
FIN scan 22
FIN without ACK flag 18
fragment reassembly 72–75

H

high-watermark threshold 41
HTTP
blocking components 167–169
keep-alive 96
session timeout 41
trickling 97

I

ICMP
fragments 172
large packets 174
ICMP flood 59
illustration
conventions ix
InterScan VirusWall 90
IP
packet fragments 180
IP options 12–14
attributes 12–14
incorrectly formatted 176
loose source route 13, 31–33
record route 13, 14
security 13, 14
source route 31
stream ID 13, 14
strict source route 14, 31–33
timestamp 14
IP spoofing 22–30
drop-no-rpf-route 23
Layer 2 24, 29
Layer 3 23, 25

J

Java applets, blocking 168

L

Land attack 63
loose source route IP option 13, 31–33
low-watermark threshold 41

M

malicious URL protection 72–75

N

names
conventions x

P

Ping of Death 65
policies
context 127
core section 126
URL filtering 118
port scan 10
probes
network 8
open ports 10
operating systems 16–20
protocol anomalies 139

R

reconnaissance 7–33
address sweep 8
FIN scan 22
IP options 12
port scan 10
SYN and FIN flags set 16
TCP packet without flags 20
record route IP option 13, 14
regular expressions 157–159

S

SCREEN
address sweep 8
bad IP options, drop 176
drop unknown MAC addresses 51
FIN with no ACK 22
FIN without ACK flag, drop 18
ICMP flood 59
ICMP fragments, block 172
IP options 12
IP packet fragments, block 180
IP spoofing 22–30
Land attack 63
large ICMP packets, block 174
loose source route IP option, detect 33
Ping of Death 65
port scan 10
source route IP option, deny 33
strict source route IP option, detect 33
SYN and FIN flags set 16
SYN flood 45–51
SYN fragments, detect 182–183
SYN-ACK-ACK proxy flood 43
TCP packet without flags, detect 20
Teardrop 67
UDP flood 61
unknown protocols, drop 178
VLAN and MGT zones 3
WinNuke attack 69
security IP option 13, 14
services
custom 152
session limits 36–40
destination based 37, 40
source based 36, 39
session table flood 36
session timeout
HTTP 41
TCP 41
UDP 41
stateful inspection 3
stateful signatures 138
definition 138
stream ID IP option 13, 14
strict source route IP option 14, 31–33
SYN and FIN flags set 16

SYN flood 45–51
  alarm threshold 49
  attack 45
  attack threshold 49
  destination threshold 50
  drop unknown MAC addresses 51
  queue size 51
  source threshold 50
  threshold 46
  timeout 51
SYN fragments 182–183
SYN-ACK-ACK proxy flood 43

T

TCP
  max simultaneous connections 94
  packet without flags 20
  session timeout 41
  stream signatures 164
Teardrop attack 67
three-way handshake 45
timestamp IP option 14
Transparent mode
  drop unknown MAC addresses 51

U

UDP
  session timeout 41
UDP flood 61
unknown protocols 178
URL filtering 113–121
  blocked URL message type 117
  communication timeout 116
  device-level activation 117
  fail/pass mode 116
  NetScreen blocked URL message 117
  policy-level application 118
  routing 119
  server status 118
  servers per vsys 115
  Websense server name 116
  Websense server port 116

W

Websense 113
WinNuke attack 69

Z

zip files, blocking 168
zombie agent 35, 37

NetScreen Concepts & Examples

ScreenOS Reference Guide

Volume 5: VPNs

ScreenOS 5.0.0

P/N 093-0928-000

Rev. B

Copyright Notice

NetScreen Technologies, Inc.
Building #3
805 11th Avenue
Sunnyvale, CA 94089
www.netscreen.com

FCC Statement

The following information is for FCC compliance of Class A devices: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. The equipment generates, uses, and can radiate radio-frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case users will be required to correct the interference at their own expense.

The following information is for FCC compliance of Class B devices: The equipment described in this manual generates and may radiate radio-frequency energy. If it is not installed in accordance with NetScreen's installation instructions, it may cause interference with radio and television reception. This equipment has been tested and found to comply with the limits for a Class B digital device in accordance with the specifications in part 15 of the FCC rules. These specifications are designed to provide reasonable protection against such interference in a residential installation. However, there is no guarantee that interference will not occur in a particular installation.

If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Consult the dealer or an experienced radio/TV technician for help.
- Connect the equipment to an outlet on a circuit different from that to which the receiver is connected.

Caution: Changes or modifications to this product could void the user's warranty and authority to operate this device.

Disclaimer

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR NETSCREEN REPRESENTATIVE FOR A COPY.

Contents

Preface v
Conventions vi
CLI Conventions vi
WebUI Conventions vii
Illustration Conventions ix
Naming Conventions and Character Types x
NetScreen Documentation xi
Chapter 1 IPSec 1
Introduction to VPNs 2
IPSec Concepts 3
Modes 4
Transport Mode 4
Tunnel Mode 5
Protocols 7
AH 7
ESP 8
Key Management 9
Manual Key 9
AutoKey IKE 9
Security Association 10
Tunnel Negotiation 11
Phase 1 11
Main Mode and Aggressive Mode 12
The Diffie-Hellman Exchange 13
Phase 2 13
Perfect Forward Secrecy 14
Replay Protection 14
Chapter 2 Public Key Cryptography 15
Introduction to Public Key Cryptography 16
PKI 18
Certificates and CRLs 21
Obtaining a Certificate Manually 22
Example: Requesting a Certificate Manually 23
Example: Loading Certificates and CRLs 26
Example: Configuring CRL Settings for a CA Certificate 28
Obtaining a Local Certificate Automatically 30
Example: Requesting a Local Certificate Automatically 31
Automatic Certificate Renewal 34
Key Pair Generation 35
Checking for Revocation Using OCSP 36
Configuring for OCSP 37
Specifying either CRL or OCSP for Revocation Checking 37
Displaying Certificate Revocation Status Attributes 37
Specifying the URL of an OCSP Responder for a Certificate 38
Removing Certificate Revocation Check Attributes 38
Chapter 3 VPN Guidelines 39
Cryptographic Options 40
Site-to-Site Cryptographic Options 41
Dialup VPN Options 50

Route- and Policy-Based Tunnels....58
Packet Flow: Site-to-Site VPN....60
Tunnel Configuration Tips....67
Chapter 4 Site-to-Site VPNs....69
Site-to-Site VPN Configurations....70
Site-to-Site Tunnel Configuration Steps....71
Example: Route-Based Site-to-Site VPN, AutoKey IKE....77
Example: Policy-Based Site-to-Site VPN, AutoKey IKE....91
Example: Route-Based Site-to-Site VPN, Dynamic Peer....102
Example: Policy-Based Site-to-Site VPN, Dynamic Peer....117
Example: Route-Based Site-to-Site VPN, Manual Key....131
Example: Policy-Based Site-to-Site VPN, Manual Key....142
FQDN for Dynamic IKE Gateways....151
Aliases....152
Example: AutoKey IKE Peer with FQDN....153
VPN Sites with Overlapping Addresses....168
Example: Tunnel Interface with NAT-Src and NAT-Dst....171
Transparent Mode VPN....186
Example: Transparent Mode, Policy-Based AutoKey IKE VPN....187
Chapter 5 Dialup VPNs....199
Dialup VPNs....200
Example: Policy-Based Dialup VPN, AutoKey IKE....201
Example: Route-Based Dialup VPN, Dynamic Peer....209
Example: Policy-Based Dialup VPN, Dynamic Peer....220
Bidirectional Policies for Dialup VPN Users....229
Example: Bidirectional Dialup VPN Policies....230
Group IKE ID....237
Group IKE ID with Certificates....238
Wildcard and Container ASN1-DN IKE ID Types....240
Example: Group IKE ID (Certificates)....243
Group IKE ID with Preshared Keys....250
Example: Group IKE ID (Preshared Keys)....252
Shared IKE IDs....259
Example: Shared IKE ID (Preshared Keys)....260
Chapter 6 L2TP....269
Introduction to L2TP....270
Packet Encapsulation and Decapsulation....274
Encapsulation....274
Decapsulation....275
L2TP Parameters....276
Example: Configuring an IP Pool and L2TP Default Settings....277
L2TP and L2TP-over-IPSec....279
Example: Configuring L2TP....280
Example: Configuring L2TP-over-IPSec....286
Chapter 7 Advanced VPN Features....299
IPSec NAT Traversal....301
Traversing a NAT Device....302
UDP Checksum....303

The Keepalive Frequency Value 303
IPSec NAT-Traversal and Initiator/Responder Symmetry 304
Example: Enabling NAT-Traversal 305
VPN Monitoring 307
Rekey and Optimization Options 307
Source Interface and Destination Address 308
Policy Considerations 310
Configuring the VPN Monitoring Feature 310
Example: Specifying Source and Destination Addresses for VPN Monitoring 312
Security Consideration for a Route-Based VPN Design 323
SNMP VPN Monitoring Objects and Traps 325
Multiple Tunnels per Tunnel Interface 326
Route-to-Tunnel Mapping 327
Remote Peers' Addresses 328
Manual and Automatic Table Entries 330
Manual Table Entries 330
Automatic Table Entries 331
Example: Multiple VPNs on One Tunnel Interface to Overlapping Subnets 333
Example: Automatic Route and NHTB Table Entries 364
Redundant VPN Gateways 382
VPN Groups 383
Monitoring Mechanisms 384
IKE Heartbeats 384
IKE Recovery Procedure 385
TCP SYN-Flag Checking 388
Example: Redundant VPN Gateways 389
Back-to-Back VPNs 401
Example: Back-to-Back VPNs 402
Hub-and-Spoke VPNs 412
Example: Hub-and-Spoke VPNs 413
Index ..... IX-I

Preface

A virtual private network (VPN) is a cost-effective and secure way for corporations to provide users dialup access to the corporate network and for remote networks to communicate with each other across the Internet. Secure private connections over the Internet are more cost-effective than dedicated private lines. NetScreen devices provide full VPN functions for secure site-to-site and dialup VPN applications.

Volume 5, “VPNs” describes the following VPN concepts and features that are available on NetScreen devices:

- Internet Protocol Security (IPsec) elements
- Certificates and certificate revocation lists (CRLs) within the context of Public Key Infrastructure (PKI)
- Site-to-site VPNs
- Dialup VPNs
- Layer 2 Tunneling Protocol (L2TP) and L2TP-over-IPSec
- Advanced VPN features such as binding multiple VPN tunnels to a single tunnel interface and redundant IKE gateways.

This volume also includes extensive examples for all the above features.

CONVENTIONS

This document contains several types of conventions, which are introduced in the following sections:

- "CLI Conventions"
- "WebUI Conventions" on page vii
- "Illustration Conventions" on page ix
- "Naming Conventions and Character Types" on page x

CLI Conventions

The following conventions are used when presenting the syntax of a command line interface (CLI) command:

- Anything inside square brackets [ ] is optional.
- Anything inside braces { } is required.
- If there is more than one choice, each choice is separated by a pipe (|). For example,

  ```
  set interface { ethernet1 | ethernet2 | ethernet3 } manage
  ```

  means "set the management options for the ethernet1, ethernet2, or ethernet3 interface".
- Variables appear in *italic.* For example:

  ```
  set admin user name password
  ```

When a CLI command appears within the context of a sentence, it is in **bold** (except for variables, which are always in *italic*). For example: "Use the **get system** command to display the serial number of a NetScreen device."

Note: *When typing a keyword, you only have to type enough letters to identify the word uniquely. For example, typing* ***set adm u joe j12fmt54*** *is enough to enter the command* ***set admin user joe j12fmt54****. Although you can use this shortcut when entering commands, all the commands documented here are presented in their entirety.*

WebUI Conventions

Throughout this book, a chevron (>) is used to indicate navigation through the WebUI by clicking menu options and links. For example, the path to the address configuration dialog box is presented as **Objects > Addresses > List > New**. This navigational sequence is shown below.

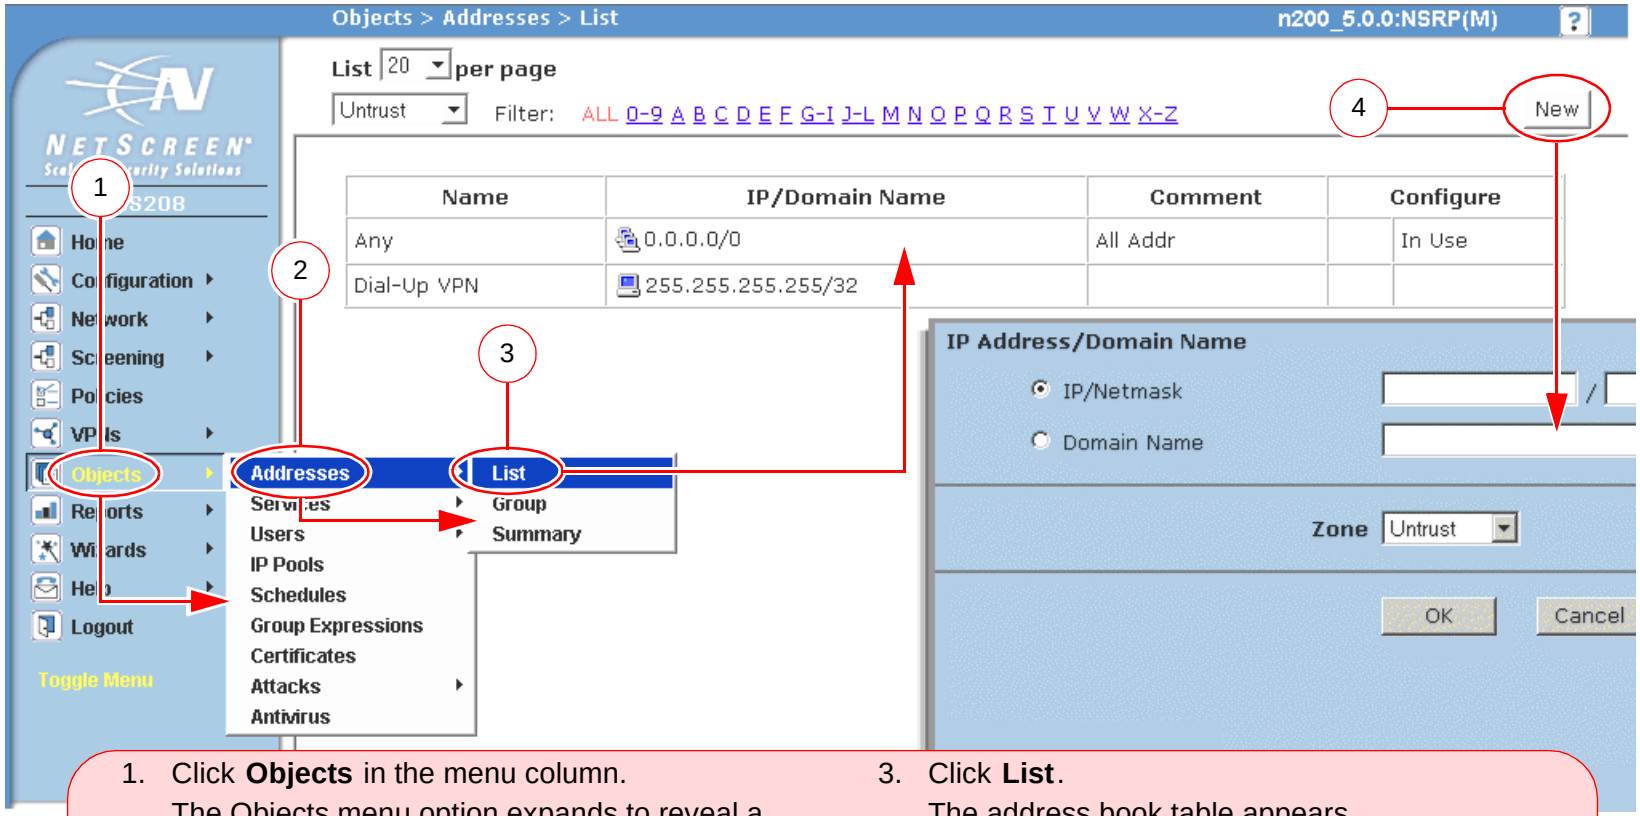

1. Click **Objects** in the menu column.
   The Objects menu option expands to reveal a subset of options for Objects.
2. (Applet menu) Hover the mouse over **Addresses**.
   (DHTML menu) Click **Addresses**.
   The Addresses option expands to reveal a subset of options for Addresses.
3. Click **List**.
   The address book table appears.
4. Click the **New** link.
   The new address configuration dialog box appears.

To perform a task with the WebUI, you must first navigate to the appropriate dialog box where you can then define objects and set parameters. The set of instructions for each task is divided into two parts: a navigational path and configuration details. For example, the following set of instructions includes the path to the address configuration dialog box and the settings for you to configure:

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: addr_1

IP Address/Domain Name:

IP/Netmask: (select), 10.2.2.5/32

Zone: Untrust

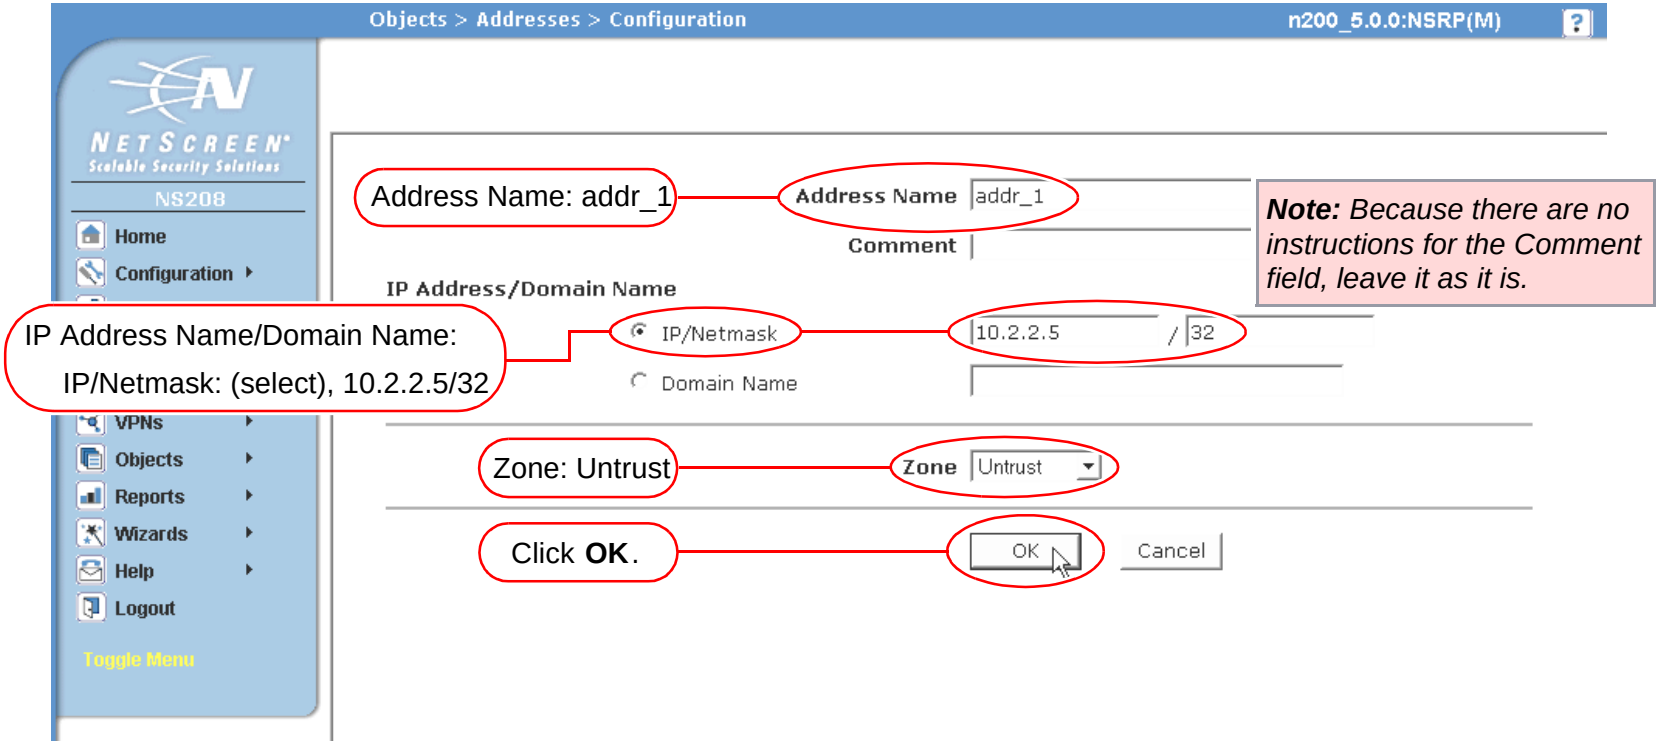

Illustration Conventions

The following graphics make up the basic set of images used in illustrations throughout this book:

Generic NetScreen Device

Virtual Routing Domain

Security Zone

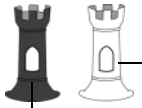

Security Zone Interfaces

White = Protected Zone Interface (example: Trust Zone)

Black = Outside Zone Interface (example: Untrust Zone)

Tunnel Interface

VPN Tunnel

Router Icon

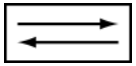

Switch Icon

Local Area Network (LAN) with a Single Subnet

(example: 10.1.1.0/24)

Internet

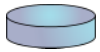

Dynamic IP (DIP) Pool

Desktop Computer

Laptop Computer

Generic Network Device

(examples: NAT server, Access Concentrator)

Server

Naming Conventions and Character Types

ScreenOS employs the following conventions regarding the names of objects—such as addresses, admin users, auth servers, IKE gateways, virtual systems, VPN tunnels, and zones—defined in ScreenOS configurations.

- If a name string includes one or more spaces, the entire string must be enclosed within double quotes ("); for example, **set address trust "local LAN" 10.1.1.0/24**.
- NetScreen trims any spaces leading or trailing text within a set of double quotes; for example, **" local LAN "** becomes **"local LAN"**.
- NetScreen treats multiple consecutive spaces as a single space.
- Name strings are case sensitive, although many CLI key words are case insensitive. For example, **"local LAN"** is different from **"local lan"**.

ScreenOS supports the following character types:

- Single-byte character sets (SBCS) and multiple-byte character sets (MBCS). Examples of SBCS are ASCII, European, and Hebrew. Examples of MBCS—also referred to as double-byte character sets (DBCS)—are Chinese, Korean, and Japanese.

  Note: *A console connection only supports SBCS. The WebUI supports both SBCS and MBCS, depending on the character sets that your Web browser supports.*

- ASCII characters from 32 (0x20 in hexidecimals) to 255 (0xff), except double quotes ("), which have special significance as an indicator of the beginning or end of a name string that includes spaces.

NETSCREEN DOCUMENTATION

To obtain technical documentation for any NetScreen product, visit www.netscreen.com/resources/manuals/.

To obtain the latest software version, visit www.netscreen.com/services/download_soft. Select a category of software product from the dropdown list, then follow the displayed instructions. (You must be a registered user to download NetScreen software.)

If you find any errors or omissions in the following content, please contact us at the e-mail address below:

techpubs@netscreen.com

1

IPSec

This chapter introduces the various elements of Internet Protocol Security (IPSec) and how they relate to virtual private network (VPN) tunneling. Following an “Introduction to VPNs” on page 2, the remainder of the chapter covers the following elements of IPSec:

- “IPSec Concepts” on page 3
  - “Modes” on page 4
  - “Protocols” on page 7
  - “Key Management” on page 9
  - “Security Association” on page 10
- “Tunnel Negotiation” on page 11
  - “Phase 1” on page 11
  - “Phase 2” on page 13

INTRODUCTION TO VPNS

A virtual private network (VPN) provides a means for securely communicating between remote computers across a public wide area network (WAN), such as the Internet.

A VPN connection can link two local area networks (LANs) or a remote dialup user and a LAN. The traffic that flows between these two points passes through shared resources such as routers, switches, and other network equipment that make up the public WAN. To secure VPN communication while passing through the WAN, the two participants create an IP Security (IPSec) tunnel[1].

An IPSec tunnel consists of a pair of unidirectional Security Associations (SAs)—one at each end of the tunnel—that specify the security parameter index (SPI), destination IP address, and security protocol (Authentication Header or Encapsulating Security Payload) employed.

Note: *For more information on SPIs, see "Security Association" on page 10. For more about the IPSec security protocols, see "Protocols" on page 7.*

Through the SA, an IPSec tunnel can provide the following security functions:

- Privacy (via encryption)
- Content integrity (via data authentication)
- Sender authentication and—if using certificates—nonrepudiation (via data origin authentication)

The security functions you employ depend on your needs. If you only need to authenticate the IP packet source and content integrity, you can authenticate the packet without applying any encryption. On the other hand, if you are only concerned with preserving privacy, you can encrypt the packet without applying any authentication mechanisms. Optionally, you can both encrypt and authenticate the packet. Most network security designers choose to encrypt, authenticate, and replay-protect their VPN traffic.

NetScreen supports IPSec technology for creating VPN tunnels with two kinds of key creation mechanisms:

- Manual Key
- AutoKey IKE with a preshared key or a certificate

1. The term "tunnel" does not denote either transport or tunnel mode (see "Modes" on page 4). It simply refers to the IPSec connection.

IPSec Concepts

IP Security (IPSec) is a suite of related protocols for cryptographically securing communications at the IP packet layer. IPSec consists of two modes and two main protocols:

- Transport and tunnel modes
- The Authentication Header (AH) protocol for authentication and the Encapsulating Security Payload (ESP) protocol for encryption (and authentication)

IPSec also provides methods for the manual and automatic negotiation of Security Associations (SAs) and key distribution, all the attributes for which are gathered in a Domain of Interpretation (DOI). See RFC 2407 and 2408.

Note: NetScreen does not support Transport Mode with AH.

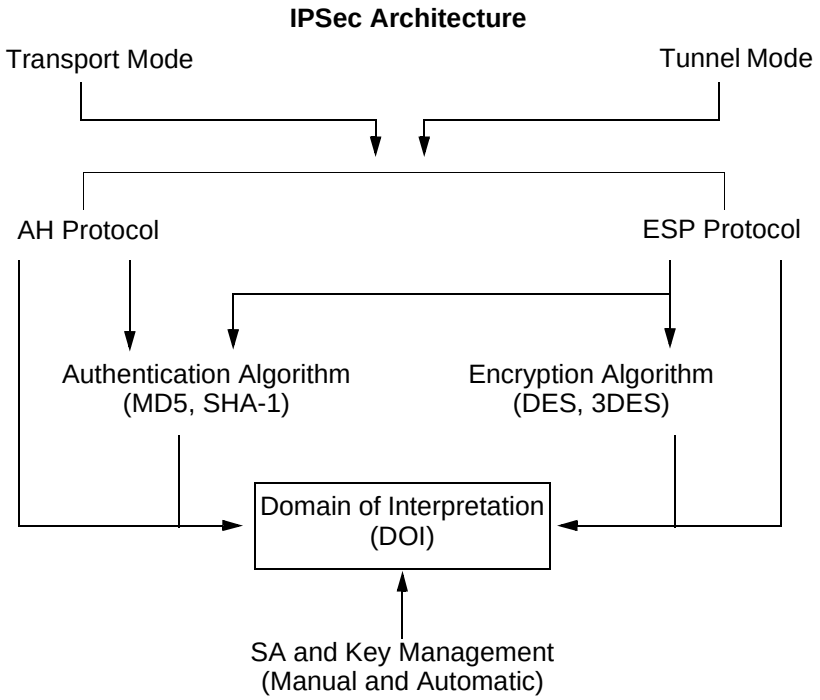

Note: *The IPSec Domain of Interpretation (DOI) is a document containing definitions for all the security parameters required for the successful negotiation of a VPN tunnel—essentially, all the attributes required for SA and IKE negotiations.*

Modes

IPSec operates in one of two modes: transport and tunnel. When both ends of the tunnel are hosts, you can use transport mode or tunnel mode. When at least one of the endpoints of a tunnel is a security gateway, such as a router or firewall, you must use tunnel mode. NetScreen devices always operate in tunnel mode for IPSec tunnels and transport mode for L2TP-over-IPSec tunnels.

Transport Mode

The original IP packet is not encapsulated within another IP packet. The entire packet can be authenticated (with AH), the payload can be encrypted (with ESP), and the original header remains in plaintext as it is sent across the WAN.

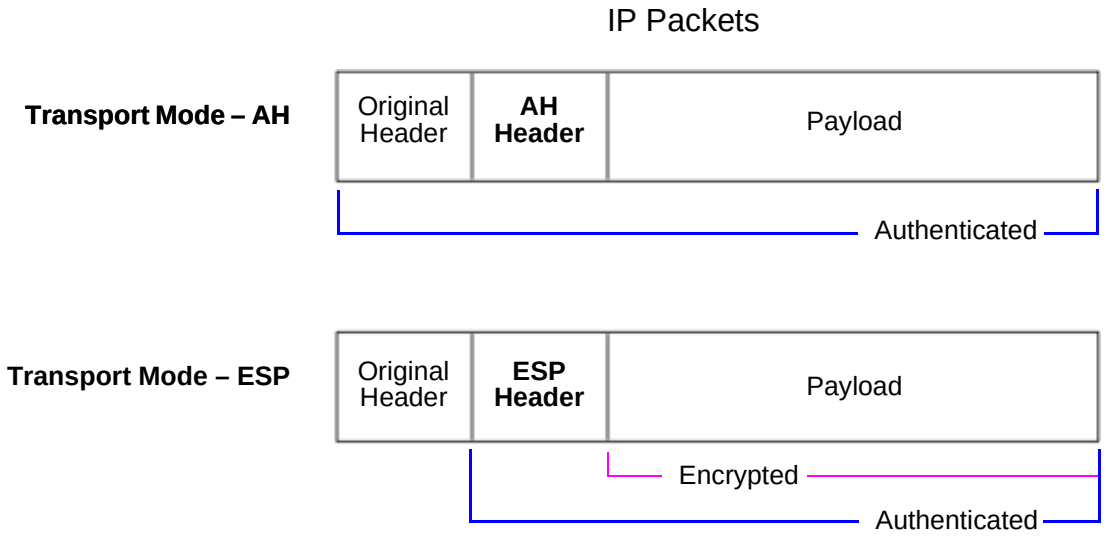

Tunnel Mode

The entire original IP packet—payload and header—is encapsulated within another IP payload and a new header appended to it. The entire original packet can be encrypted, authenticated, or both. With AH, the AH and new headers are also authenticated. With ESP, the ESP header can also be authenticated.

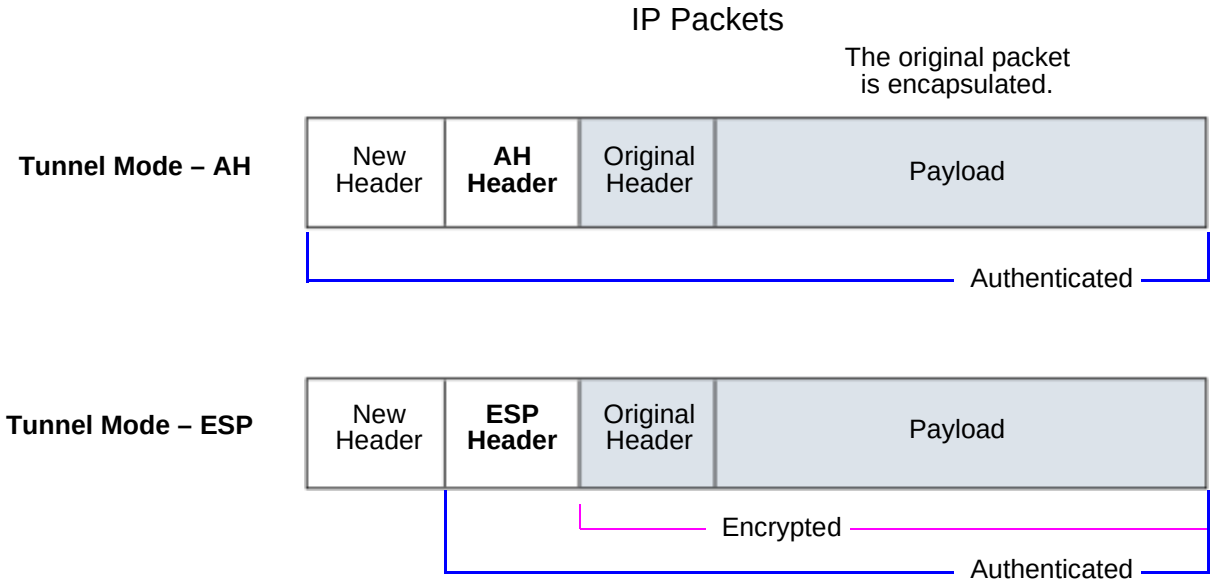

In a site-to-site VPN, the source and destination addresses used in the new header are the IP addresses of the outgoing interface (in NAT or Route mode) or the VLAN1 IP address (in Transparent mode); the source and destination addresses of the encapsulated packets are the addresses of the ultimate endpoints of the connection.

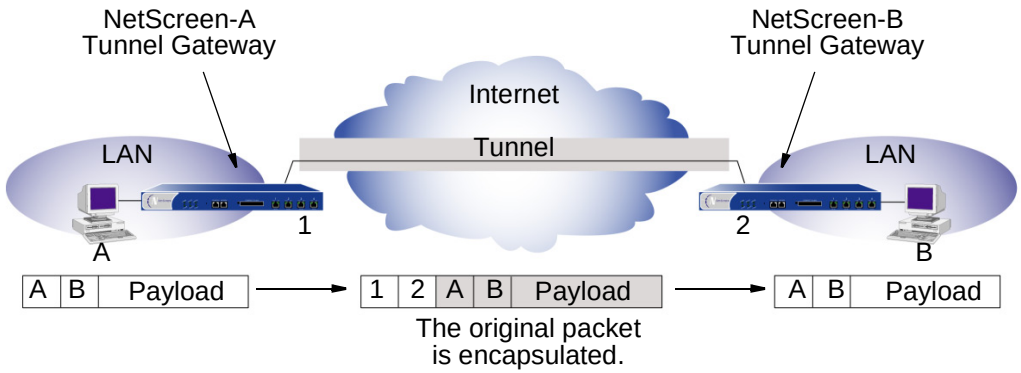

Site-to-Site VPN in Tunnel Mode

In a dialup VPN, there is no tunnel gateway on the VPN dialup client end of the tunnel; the tunnel extends directly to the client itself. In this case, on packets sent from the dialup client, both the new header and the encapsulated original header have the same IP address: that of the client's computer[2].

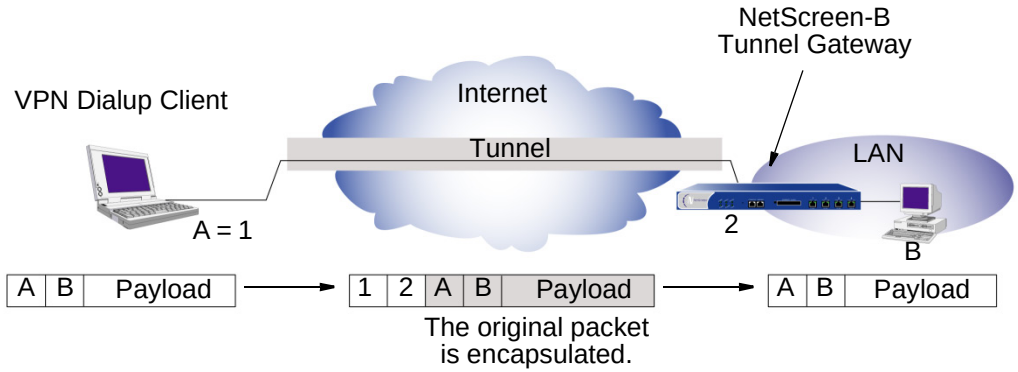

Dialup VPN in Tunnel Mode

2. Some VPN clients such as the NetScreen-Remote allow you to define a virtual inner IP address. In such cases, the virtual inner IP address is the source IP address in the original packet header of traffic originating from the client, and the IP address that the ISP dynamically assigns the dialup client is the source IP address in the outer header.

Protocols

IPSec uses two protocols to secure communications at the IP layer:

- Authentication Header (AH)—A security protocol for authenticating the source of an IP packet and verifying the integrity of its content
- Encapsulating Security Payload (ESP)—A security protocol for encrypting the entire IP packet (and authenticating its content)

AH

The Authentication Header (AH) protocol provides a means to verify the authenticity/integrity of the content and origin of a packet. You can authenticate the packet by the checksum calculated via a hash-based message authentication code (HMAC) using a secret key and either MD5 or SHA-1 hash functions.

Message Digest version 5 (MD5)—An algorithm that produces a 128-bit hash (also called a digital signature or message digest) from a message of arbitrary length and a 16-byte key. The resulting hash is used, like a fingerprint of the input, to verify content and source authenticity and integrity.

Secure Hash Algorithm-1 (SHA-1)—An algorithm that produces a 160-bit hash from a message of arbitrary length and a 20-byte key. It is generally regarded as more secure than MD5 because of the larger hashes it produces. Because the computational processing is done in the NetScreen ASIC, the performance cost is negligible.

Note: *For more information on MD5 and SHA-1 hashing algorithms, see the following RFCs: (MD5) 1321, 2403; (SHA-1) 2404. For information on HMAC, see RFC 2104.*

ESP

The Encapsulating Security Payload (ESP) protocol provides a means to ensure privacy (encryption), and source authentication and content integrity (authentication). ESP in tunnel mode encapsulates the entire IP packet (header and payload), and then appends a new IP header to the now encrypted packet. This new IP header contains the destination address needed to route the protected data through the network.

With ESP, you can encrypt and authenticate, encrypt only, or authenticate only. For encryption, you can choose either of the following encryption algorithms:

Data Encryption Standard (DES)—A cryptographic block algorithm with a 56-bit key.

Triple DES (3DES)—A more powerful version of DES in which the original DES algorithm is applied in three rounds, using a 168-bit key. DES provides a significant performance savings but is considered unacceptable for many classified or sensitive material transfers.

Advanced Encryption Standard (AES)—An emerging encryption standard which, when adopted by Internet infrastructures worldwide, will offer greater interoperability with other network security devices. NetScreen supports AES with 128-, 192-, and 256-bit keys.

For authentication, you can use either MD5 or SHA-1 algorithms.

For either the encryption or authentication algorithm you can select **NULL**; however, you cannot select **NULL** for both simultaneously.

Key Management

The distribution and management of keys are critical to successfully using VPNs. IPSec supports both manual and automatic key distribution methods.

Manual Key

With Manual Keys, administrators at both ends of a tunnel configure all the security parameters. This is a viable technique for small, static networks where the distribution, maintenance, and tracking of keys are not difficult. However, safely distributing Manual Key configurations across great distances poses security issues. Aside from passing the keys face-to-face, you cannot be completely sure that the keys have not been compromised while in transit. Also, whenever you want to change the key, you are faced with the same security issues as when you initially distributed it.

AutoKey IKE

When you need to create and manage numerous tunnels, you need a method that does not require you to configure every element manually. IPSec supports the automated generation and negotiation of keys and security associations using the Internet Key Exchange (IKE) protocol. NetScreen refers to such automated tunnel negotiation as AutoKey IKE and supports AutoKey IKE with preshared keys and AutoKey IKE with certificates.

AutoKey IKE with Preshared Keys

With AutoKey IKE using preshared keys to authenticate the participants in an IKE session, each side must configure and securely exchange the preshared key[3] in advance. In this regard, the issue of secure key distribution is the same as that with Manual Keys. However, once distributed, an AutoKey, unlike a Manual Key, can automatically change its keys at predetermined intervals using the IKE protocol. Frequently changing keys greatly improves security, and automatically doing so greatly reduces key-management responsibilities. However, changing keys increases traffic overhead; therefore, doing so too often can reduce data transmission efficiency.

3. A preshared key is a key for both encryption and decryption that both participants must have before initiating communication.

AutoKey IKE with Certificates

When using certificates to authenticate the participants during an AutoKey IKE negotiation, each side generates a public/private key pair (see Chapter 2, "Public Key Cryptography" on page 15) and acquires a certificate (see "Certificates and CRLs" on page 21). As long as the issuing certificate authority (CA) is trusted by both sides, the participants can retrieve the peer's public key and verify the peer's signature. There is no need to keep track of the keys and SAs; IKE does it automatically.

Note: *For examples of both Manual Key and AutoKey IKE tunnels, see Chapter 4, "Site-to-Site VPNs" on page 69.*

Security Association

A security association (SA) is a unidirectional agreement between the VPN participants regarding the methods and parameters to use in securing a communication channel. Full bidirectional communication requires at least two SAs, one for each direction.

An SA groups together the following components for securing communications:

- Security algorithms and keys
- Protocol mode (transport or tunnel)
- Key management method (Manual Key or AutoKey IKE)
- SA lifetime

For outbound VPN traffic, the policy invokes the SA associated with the VPN tunnel. For inbound traffic, the NetScreen device looks up the SA by using the following triplet: destination IP, security protocol (AH or ESP), and security parameter index (SPI) value.

TUNNEL NEGOTIATION

For a Manual Key IPSec tunnel, because all of the security association (SA) parameters have been previously defined, there is no need to negotiate which SAs to use. In essence, the tunnel has already been established. When traffic matches a policy using that Manual Key tunnel or when a route involves the tunnel, the NetScreen device simply encrypts and authenticates the data, as you determined, and forwards it to the destination gateway.

To establish an AutoKey IKE IPSec tunnel, two phases of negotiation are required:

- In Phase 1, the participants establish a secure channel in which to negotiate the IPSec SAs.
- In Phase 2, the participants negotiate the IPSec SAs for encrypting and authenticating the ensuing exchanges of user data.

Phase 1

Phase 1 of an AutoKey IKE tunnel negotiation consists of the exchange of proposals for how to authenticate and secure the channel. The exchange can be in one of two modes: Aggressive mode or Main mode (see below). Using either mode, the participants exchange proposals for acceptable security services such as:

- Encryption algorithms (DES and 3DES) and authentication algorithms (MD5 and SHA-1). For more information about these algorithms, see “Protocols” on page 7.
- A Diffie-Hellman Group (See “The Diffie-Hellman Exchange” on page 13.)
- Preshared Key or RSA/DSA certificates (see “AutoKey IKE” on page 9)

A successful Phase 1 negotiation concludes when both ends of the tunnel agree to accept at least one set of the Phase 1 security parameters proposed, and then process them. NetScreen devices support up to four proposals for Phase 1 negotiations, allowing you to define how restrictive a range of security parameters for key negotiation you will accept.

The predefined Phase 1 proposals that NetScreen provides are as follows:

- **Standard:** pre-g2-aes128-sha and pre-g2-3des-sha
- **Compatible:** pre-g2-3des-sha, pre-g2-3des-md5, pre-g2-des-sha, and pre-g2-des-md5
- **Basic:** pre-g1-des-sha and pre-g1-des-md5

You can also define custom Phase 1 proposals.

Main Mode and Aggressive Mode

Phase 1 can take place in either Main mode or Aggressive mode. The two modes are described below.

Main Mode: The initiator and recipient send three two-way exchanges (six messages total) to accomplish the following services:

- First exchange, (messages 1 and 2): Propose and accept the encryption and authentication algorithms.
- Second exchange, (messages 3 and 4): Execute a Diffie-Hellman exchange, and the initiator and recipient each provide a nonce (randomly generated number).
- Third exchange, (messages 5 and 6): Send and verify their identities.

The information transmitted in the third exchange of messages is protected by the encryption algorithm established in the first two exchanges. Thus, the participants' identities are not transmitted in the clear.

Aggressive Mode: The initiator and recipient accomplish the same objectives, but only in two exchanges, and a total of three messages:

- First message: The initiator proposes the SA, initiates a Diffie-Hellman exchange, and sends a nonce and its IKE identity.
- Second message: The recipient accepts the SA, authenticates the initiator, and sends a nonce, its IKE identity, and, if using certificates, the recipient's certificate.
- Third message: The initiator authenticates the recipient, confirms the exchange, and, if using certificates, sends the initiator's certificate.

Because the participants' identities are exchanged in the clear (in the first two messages), Aggressive mode does not provide identity protection.

Note: *When a dialup VPN user negotiates an AutoKey IKE tunnel with a preshared key, Aggressive mode must be used. Note also that a dialup VPN user can use an e-mail address, a fully qualified domain name (FQDN), or an IP address as its IKE ID. A dynamic peer can use either an e-mail address or FQDN, but not an IP address.*

The Diffie-Hellman Exchange

A Diffie-Hellman exchange allows the participants to produce a shared secret value. The strength of the technique is that it allows the participants to create the secret value over an unsecured medium without passing the secret value through the wire. There are five Diffie-Hellman (DH) groups (NetScreen supports groups 1, 2, and 5). The size of the prime modulus used in each group's calculation differs as follows:

- DH Group 1: 768-bit modulus[4]
- DH Group 2: 1024-bit modulus
- DH Group 5: 1536-bit modulus

The larger the modulus, the more secure the generated key is considered to be; however, the larger the modulus, the longer the key-generation process takes. Because the modulus for each DH group is a different size, the participants must agree to use the same group[5].

Phase 2

After the participants have established a secure and authenticated channel, they proceed through Phase 2, in which they negotiate the SAs to secure the data to be transmitted through the IPSec tunnel.

Like the process for Phase 1, the participants exchange proposals to determine which security parameters to employ in the SA. A Phase 2 proposal also includes a security protocol—either Encapsulating Security Payload (ESP) or Authentication Header (AH), and selected encryption and authentication algorithms. The proposal can also specify a Diffie-Hellman group, if Perfect Forward Secrecy (PFS) is desired.

Note: *For more about Diffie-Hellman groups, see "The Diffie-Hellman Exchange" above. For more about PFS, see "Perfect Forward Secrecy" on page 14.*

Regardless of the mode used in Phase 1, Phase 2 always operates in Quick mode and involves the exchange of three messages[5].

4. The strength of DH Group 1 security has depreciated, and NetScreen does not recommend its use.
5. If you configure multiple (up to four) proposals for Phase 1 negotiations, use the same Diffie-Hellman group in all proposals. The same guideline applies to multiple proposals for Phase 2 negotiations.

NetScreen devices support up to four proposals for Phase 2 negotiations, allowing you to define how restrictive a range of tunnel parameters you will accept. NetScreen also provides a replay protection feature. Use of this feature does not require negotiation because packets are always sent with sequence numbers. You simply have the option of checking the sequence numbers or not. (For more information about replay protection, see below.)

The predefined Phase 2 proposals that NetScreen provides are as follows:

- **Standard:** g2-esp-3des-sha and g2-esp-aes128-sha
- **Compatible:** nopfs-esp-3des-sha, nopfs-esp-3des-md5, nopfs-esp-des-sha, and nopfs-esp-des-md5
- **Basic:** nopfs-esp-des-sha and nopfs-esp-des-md5

You can also define custom Phase 2 proposals.

In Phase 2, the peers also exchange proxy IDs. A proxy ID is a three-part tuple consisting of local IP address–remote IP address–service. The proxy ID for both peers must match, which means that the service specified in the proxy ID for both peers must be the same, and the local IP address specified for one peer must be the same as the remote IP address specified for the other peer.

Perfect Forward Secrecy

Perfect Forward Secrecy (PFS) is a method for deriving Phase 2 keys independent from and unrelated to the preceding keys. Alternatively, the Phase 1 proposal creates the key (the SKEYID_d key) from which all Phase 2 keys are derived. The SKEYID_d key can generate Phase 2 keys with a minimum of CPU processing. Unfortunately, if an unauthorized party gains access to the SKEYID_d key, all your encryption keys are compromised.

PFS addresses this security risk by forcing a new Diffie-Hellman key exchange to occur for each Phase 2 tunnel. Using PFS is thus more secure, although the rekeying procedure in Phase 2 might take slightly longer with PFS enabled.

Replay Protection

A replay attack occurs when somebody intercepts a series of packets and uses them later either to flood the system, causing a denial-of-service (DoS), or to gain entry to the trusted network. The replay protection feature enables NetScreen devices to check every IPSec packet to see if it has been received before. If packets arrive outside a specified sequence range, the NetScreen device rejects them.

2

Public Key Cryptography

This chapter provides an introduction to public key cryptography and the use of certificates and certificate revocation lists (CRLs) within the context of Public Key Infrastructure (PKI). The material is organized into the following sections:

- "Introduction to Public Key Cryptography" on page 16
- "PKI" on page 18
- "Certificates and CRLs" on page 21
  - "Obtaining a Certificate Manually" on page 22
  - "Obtaining a Local Certificate Automatically" on page 30
  - "Automatic Certificate Renewal" on page 34
- "Checking for Revocation Using OCSP" on page 36
  - "Configuring for OCSP" on page 37

Introduction to Public Key Cryptography

In public key cryptography, a public/private key pair is used to encrypt and decrypt data. Data encrypted with a public key, which the owner makes available to the public, can only be decrypted with the corresponding private key, which the owner keeps secret and protected. For example, if Alice wants to send Bob an encrypted message, Alice can encrypt it with Bob's public key and send it to him. Bob then decrypts the message with his private key.

The reverse is also useful; that is, encrypting data with a private key and decrypting it with the corresponding public key. This is known as creating a digital signature. For example, if Alice wants to present her identity as the sender of a message, she can encrypt the message with her private key and send the message to Bob. Bob then decrypts the message with Alice's public key, thus verifying that Alice is indeed the sender.

Public/private key pairs also play an important role in the use of digital certificates. The procedure for signing a certificate (by a CA) and then verifying the signature works as follows (by the recipient):

Signing a Certificate

1. The Certificate Authority (CA) that issues a certificate hashes the certificate by using a hash algorithm (MD5 or SHA-1) to generate a digest.
2. The CA then "signs" the certificate by encrypting the digest with its private key. The result is a digital signature.
3. The CA then sends the digitally signed certificate to the person who requested it.

Verifying a Digital Signature

1. When the recipient gets the certificate, he or she also generates another digest by applying the same hash algorithm (MD5 or SHA-1) on the certificate file.
2. The recipient uses the CA's public key to decrypt the digital signature.
3. The recipient compares the decrypted digest with the digest he or she just generated. If the two digests match, the recipient can confirm the integrity of the CA's signature and, by extension, the integrity of the accompanying certificate.

1. Using either the MD5 or SHA-1 hash algorithm, the CA makes digest A from the certificate.
2. Using the its private key, the CA encrypts digest A. The result is digest B, the digital signature.
3. The CA sends the digitally signed certificate to the person who requested it.

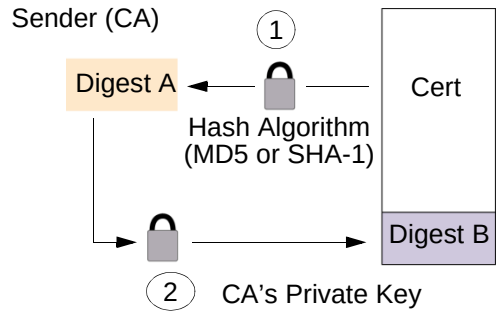

1. Using either MD5 or SHA-1, the recipient makes digest A from the certificate.
2. Using the CA's public key, the recipient decrypts digest B.
3. The recipient compares digest A with digest B. If they match, the recipient knows that the certificate has not been tampered with.

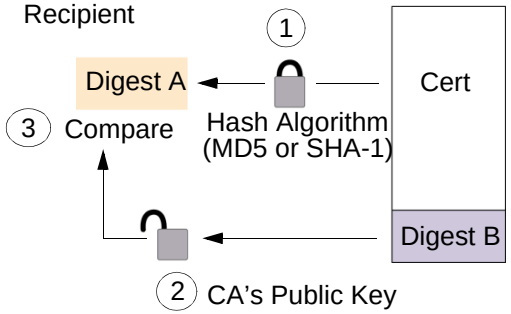

The procedure for digitally signing messages sent between two participants in an IKE session works very similarly, with the following differences:

- Instead of making a digest from the CA certificate, the sender makes it from the data in the IP packet payload.
- Instead of using the CA's public/private key pair, the participants use the sender's public/private key pair.

PKI

The term Public Key Infrastructure (PKI) refers to the hierarchical structure of trust required for the successful implementation of public key cryptography. To verify the trustworthiness of a certificate, you must be able to track a path of certified CAs from the one issuing your local certificate back to a root authority of a CA domain.

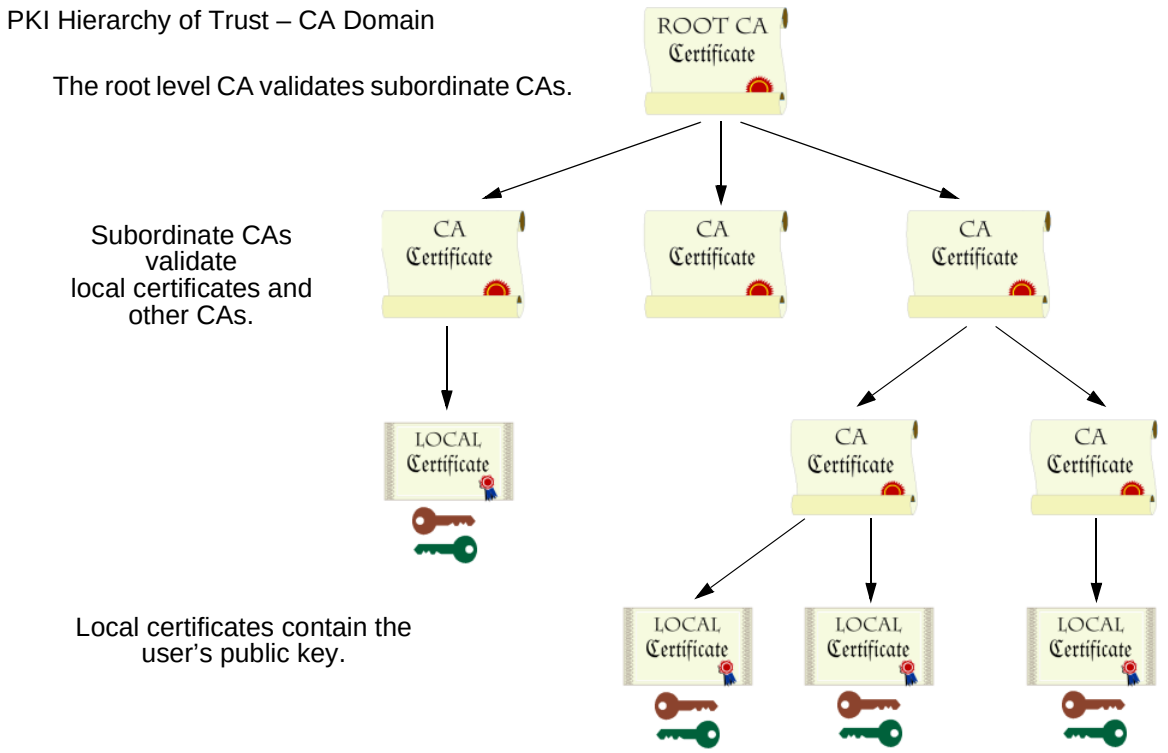

If certificates are used solely within an organization, that organization can have its own CA domain within which a company CA issues and validates certificates among its employees. If that organization later wants its employees to be able to exchange their certificates with those from another CA domain (for example, with employees at another organization that also has its own CA domain), the two CAs can develop cross-certification; that is, they can agree to trust the authority of each other. In this case, the PKI structure does not extend vertically but horizontally.

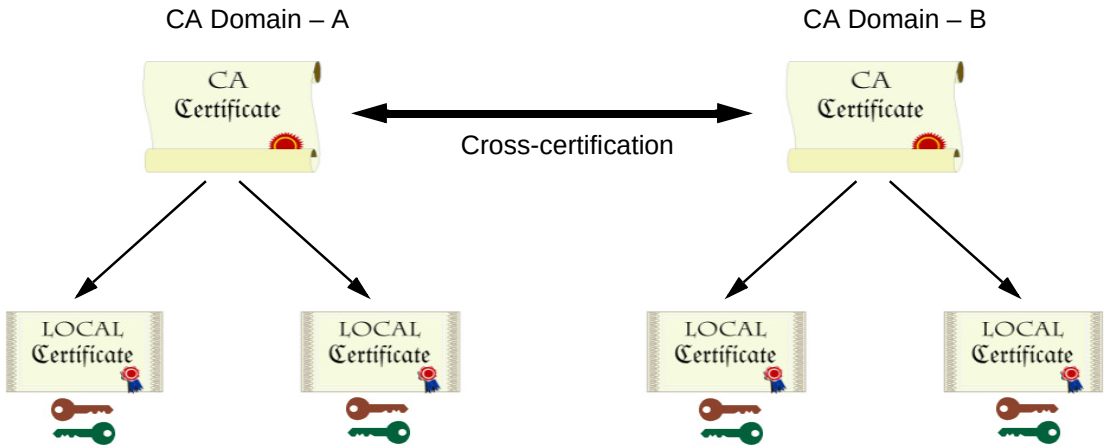

Users in CA domain A can use their certificates and key pairs with users in CA domain B because the CAs have cross-certified each other.

For convenience and practicality, PKI must be transparently managed and implemented. Toward this goal, the NetScreen ScreenOS does the following:

1. Generates a public/private key pair when you create a certificate request.
2. Supplies that public key as part of the certificate request in the form of a text file for transmission to a Certificate Authority (CA) for certificate enrollment (PKCS10 file).

3. Supports loading the local certificate, the CA certificate, and the certificate revocation list (CRL)[1] into the unit.

   You can also specify an interval for refreshing the CRL online. For more information on CRLs, see “Certificates and CRLs” on page 21.

4. Provides certificate delivery when establishing an IPSec tunnel.
5. Supports certificate path validation upward through eight levels of CA authorities in the PKI hierarchy.
6. Supports the PKCS #7 cryptographic standard, which means the NetScreen device can accept X.509 certificates and CRLs packaged within a PKCS #7 envelope[2]. PKCS #7 support allows you to submit multiple X.509 certificates within a single PKI request. You can now configure PKI to validate all the submitted certificates from the issuing CA at one time.
7. Supports online CRL retrieval via LDAP or HTTP.

1. The Certificate Authority usually provides a CRL. Although you can load a CRL into the NetScreen device, you cannot view it once loaded.
2. NetScreen supports a PKCS #7 file size of up to 7 Kilobytes.

CERTIFICATES AND CRLS

A digital certificate is an electronic means for verifying your identity through the word of a trusted third party, known as a Certificate Authority (CA). The CA server you use can be owned and operated by an independent CA[3], or by your own organization, in which case you become your own CA. If you use an independent CA, you must contact them for the addresses of their CA and CRL servers (for obtaining certificates and certificate revocation lists), and for the information they require when submitting personal certificate requests. When you are your own CA, you determine this information yourself.

Note: *ScreenOS contains a CA certificate for authenticating downloads from the antivirus (AV) pattern file server and the Deep Inspection (DI) attack object database server. For more information about the AV pattern file server, see "Enabling Internal AV Scanning" on page* ***4****-81. For more information about the DI attack object database server, see "Attack Object Database Server" on page* ***4****-128.*

To use a digital certificate to authenticate your identity when establishing a secure VPN connection, you must first do the following:

- Generate a key in the NetScreen device, send it to a CA to obtain a personal certificate (also known as a *local* certificate), and load the certificate in the NetScreen device.
- Obtain a CA certificate for the CA that issued the personal certificate (basically verifying the identity of the CA verifying you), and load the CA certificate in the NetScreen device. You can perform this task manually, or automatically using Simple Certificate Enrollment Protocol (SCEP).
- If the certificate does not contain a certificate distribution point (CDP) extension, and you cannot automatically retrieve the CRL via LDAP or HTTP, you can retrieve a CRL manually and load that in the NetScreen device.

During the course of business, there are several events that make it necessary to revoke a certificate. You might wish to revoke a certificate if you suspect that it has been compromised or when a certificate holder leaves a company. Managing certificate revocations and validation can be accomplished locally (which is a limited solution) or by referencing a CA's CRL, which you can automatically access online at daily, weekly, or monthly intervals, or at the default interval set by the CA.

3. NetScreen supports the following CAs: Baltimore, Entrust, Microsoft, Netscape, RSA Keon, and Verisign.

Obtaining a Certificate Manually

To obtain a signed digital certificate using the manual method, you must complete several tasks in the following order:

1. Generate a public/private key pair.
2. Fill out the Certificate Request.
3. Submit your request to your CA of choice.
4. After you receive your signed certificate, you must load it into the NetScreen device along with the CA certificate.

You now have the following items for the following uses:

- A local certificate for the NetScreen device, to authenticate your identity with each tunnel connection
- A CA Certificate (their public key), to be used to verify the peer's certificate
- If the Certificate Revocation List (CRL) was included with the CA certificate[4], a CRL to identify invalid certificates

When you receive these files (the certificate files typically have the extension .cer, and the CRL typically has the extension .crl), load them into your NetScreen using the procedure described in the following section.

Note: *If you are planning to use e-mail to submit a PKCS10 file to obtain your certificates, you must properly configure your NetScreen settings so that you can send e-mail to your system administrator. You have to set your primary and secondary DNS servers and specify the SMTP server and e-mail address settings.*

4. A CRL might accompany a CA certificate and be stored in the NetScreen database. Alternatively, the CA certificate might contain the CRL URL (either LDAP or HTTP) for a CRL that is stored in the CA's database. If the CRL is unobtainable by either method, you can manually enter the default server settings for the CRL URL in the NetScreen device, as explained in "Example: Configuring CRL Settings for a CA Certificate" on page 28.

Example: Requesting a Certificate Manually

When you request a certificate, the NetScreen device generates a key pair. The public key becomes incorporated in the request itself and, eventually, in the digitally signed local certificate you receive from the CA.

In the following example, the security administrator is making a certificate request for Michael Zhang in the Development department at NetScreen Technologies in Santa Clara, California. The certificate is going to be used for a NetScreen device at IP address 10.10.5.44. The administrator instructs the NetScreen device to send the request via e-mail to the security administrator at *admin@netscreen.com.* The security administrator then copies and pastes the request in the certificate request text field at the CA's certificate enrollment site. After the enrollment process is complete, the CA usually sends the certificates via e-mail back to the security administrator.

Note: *Before generating a certificate request, make sure that you have set the system clock and assigned a host name and domain name to the NetScreen device. (If the NetScreen device is in an NSRP cluster, replace the host name with a cluster name. For more information, see "Cluster Name" on page* ***8****-17.)*

WebUI

1. Certificate Generation

Objects > Certificates > New: Enter the following, and then click **Generate**:

Name: Michael Zhang

Phone: 408-730-6000

Unit/Department: Development

Organization: NetScreen Technologies

County/Locality: Santa Clara

State: CA

Country: US

E-mail: mzhang@netscreen.com[5]

IP Address: 10.10.5.44

Write to file: (select)

RSA: (select)

Create new key pair of *1024*[6] length: (select)

The NetScreen generates a PKCS #10 file and prompts you to send the file via e-mail, save the file to disk, or automatically enroll via the Simple Certificate Enrollment Protocol (SCEP).

Select the **E-mail to** option, type **admin@netscreen.com**, and then click **OK**[7].

2. Certificate Request

The security administrator opens the file, and copies its contents, taking care to copy the entire text but not any blank spaces before or after the text. (Start at "-----BEGIN CERTIFICATE REQUEST-----", and end at "-----END CERTIFICATE REQUEST-----".)

The security administrator then follows the certificate request directions at the CA's Web site, pasting the PKCS #10 file in the appropriate field when required.

3. Certificate Retrieval

When the security administrator receives the certificate from the CA via e-mail, he forwards it to you. Copy it to a text file, and save it to your workstation (to be loaded to the NetScreen device later through the WebUI) or to a TFTP server (to be loaded later through the CLI).

---

5. Some CAs do not support an e-mail address in a certificate. If you do not include an e-mail address in the local certificate request, you cannot use an e-mail address as the local IKE ID when configuring the NetScreen device as a dynamic peer. Instead, you can use a fully-qualified domain name (if it is in the local certificate), or you can leave the local ID field empty. By default the NetScreen device sends its hostname.domainname. If you do not specify a local ID for a dynamic peer, enter the hostname.domainname of that peer on the device at the other end of the IPSec tunnel in the peer ID field.
6. The value 1024 indicates the bit length of the key pair. If you are using the certificate for SSL (see "Secure Sockets Layer" on page **3**-7), be sure to use a bit length that your Web browser also supports.
7. Using the e-mail address assumes that you have already configured the IP address for your SMTP server: **set admin mail server-name** { *ip_addr* **|** *dom_name* }.

CLI

1. Certificate Generation

```
set pki x509 dn country-name US
set pki x509 dn email mzhang@netscreen.com
set pki x509 dn ip 10.10.5.44
set pki x509 dn local-name “Santa Clara”
set pki x509 dn name “Michael Zhang”
set pki x509 dn org-name “NetScreen Technologies”
set pki x509 dn org-unit-name Development
set pki x509 phone 408-730-6000
set pki x509 dn state-name CA
set pki x509 default send-to admin@netscreen.com[8]
exec pki rsa new-key 1024
```

The certificate request is sent via e-mail to admin@netscreen.com.

2. Certificate Request

The security administrator opens the file, and copies its contents, taking care to copy the entire text but not any blank spaces before or after the text. (Start at “-----BEGIN CERTIFICATE REQUEST-----”, and end at “-----END CERTIFICATE REQUEST-----”.)

The security administrator then follows the certificate request directions at the CA’s Web site, pasting the PKCS #10 file in the appropriate field when required.

3. Certificate Retrieval

When the security administrator receives the certificate from the CA via e-mail, he forwards it to you. Copy it to a text file, and save it to your workstation (to be loaded to the NetScreen device later through the WebUI) or to a TFTP server (to be loaded later through the CLI).

8. Using the e-mail address assumes that you have already configured the IP address for your SMTP server: **set admin mail server-name** { *ip_addr* **|** *dom_name* }.

Example: Loading Certificates and CRLs

The CA returns the following three files to you for loading onto the NetScreen device:

- A CA certificate, which contains the CA's public key
- A local certificate that identifies your local machine (your public key)
- A CRL, which lists any certificates revoked by the CA

For the WebUI example, you have downloaded the files to a directory named C:\certs\ns on the administrator's workstation. For the CLI example, you have downloaded the TFTP root directory on a TFTP server with IP address 198.168.1.5.

Note: *NetScreen devices, including virtual systems, configured with ScreenOS 2.5 or later support loading multiple local certificates from different CAs.*

This example illustrates how to load two certificate files named auth.cer (CA certificate) and local.cer (your public key), and the CRL file named distrust.crl.

WebUI

1. Objects > Certificates: Select **Load Cert**, and then click **Browse**.
2. Navigate to the C:\certs directory, select **auth.cer**, and then click **Open**.

   The directory path and file name (C:\certs\ns\auth.cer) appear in the File Browse field.
3. Click **Load**.

   The auth.cer certificate file loads.
4. Objects > Certificates: Select **Load Cert**, and then click **Browse**.
5. Navigate to the C:\certs directory, select **local.cer**, and then click **Open**.

   The directory path and file name (C:\certs\ns\local.cer) appear in the File Browse field.

6. Click **Load**.

   The auth.cer certificate file loads.

7. Objects > Certificates: Select **Load CRL**, and then click **Browse**.
8. Navigate to the C:\certs directory, select **distrust.crl**, and then click **Open**.
9. Click **Load**.

   The distrust.crl CRL file loads.

CLI

```
exec pki x509 tftp 198.168.1.5 cert-name auth.cer
exec pki x509 tftp 198.168.1.5 cert-name local.cer
exec pki x509 tftp 198.168.1.5 crl-name distrust.crl
```

Example: Configuring CRL Settings for a CA Certificate

In Phase 1 negotiations, participants check the CRL list to see if certificates received during an IKE exchange are still valid. If a CRL did not accompany a CA certificate and is not loaded in the NetScreen database, the NetScreen device tries to retrieve the CRL through the LDAP or HTTP[9] CRL location defined within the CA certificate itself. If there is no URL address defined in the CA certificate, the NetScreen device uses the URL of the server that you define for that CA certificate. If you do not define a CRL URL for a particular CA certificate, the NetScreen device refers to the CRL server at the default CRL URL address.

Note: *With ScreenOS 2.5 and later, you can disable the checking of a CRL's digital signature when you load the CRL. However, disabling CRL certificate checking compromises the security of your NetScreen device.*

In this example, you first configure the Entrust CA server to check the CRL daily by connecting to the LDAP server at 2.2.2.121 and locating the CRL file. You then configure default certificate validation settings to use the company's LDAP server at 10.1.1.200, also checking the CRL on a daily basis.

Note: *The index (IDX) number for the Entrust CA certificate is 1. To view a list of the IDX numbers for all the CA certificates loaded on a NetScreen device, use the following CLI command:* ***get pki x509 list ca-cert****.*

WebUI

Objects > Certificates (Show: CA) > Server Settings (for NetScreen): Enter the following, and the click **OK**:

X509 Cert_Path Validation Level: Full

CRL Settings:

URL Address: ldap:///CN=Entrust,CN=en2001,CN=PublicKeyServices, CN=Services,CN=Configuration,DC=EN2001,DC=com?CertificateRevocationList?base?objectclass=CRLDistributionPoint

LDAP Server: 2.2.2.121

Refresh Frequency: Daily

9. The CRL distribution point extension (.cdp) in an X509 certificate can be either an HTTP URL or an LDAP URL.

Objects > Certificates > Default Cert Validation Settings: Enter the following, and then click **OK**:

X509 Certificate Path Validation Level: Full

Certificate Revocation Settings:

Check Method: CRL

URL Address: ldap:///CN=NetScreen,CN=safecert,CN=PublicKeyServices,CN=Services,CN=Configuration,DC=SAFECERT,DC=com?CertificateRevocationList?base?objectclass=CRLDistributionPoint

LDAP Server: 10.1.1.200

CLI

```
set pki authority 1 cert-path full
set pki authority 1 cert-status crl url “ldap:///CN=Entrust,CN=en2001,
    CN=PublicKeyServices,CN=Services,CN=Configuration,DC=EN2000,DC=com?Certific
    ateRevocationList?base?objectclass=CRLDistributionPoint”
set pki authority 1 cert-status crl server-name 2.2.2.121
set pki authority 1 cert-status crl refresh daily
set pki authority default cert-path full
set pki authority default cert-status crl url “ldap:///CN=NetScreen,
    CN=safecert,CN=PublicKeyServices,CN=Services,CN=Configuration,DC=SAFECERT,
    DC=com?CertificateRevocationList?base?objectclass=CRLDistributionPoint”
set pki authority default cert-status crl server-name 10.1.1.200
set pki authority default cert-status crl refresh daily
save
```

Obtaining a Local Certificate Automatically

To use a digital certificate to authenticate your identity when establishing a secure VPN connection, you must first do the following:

- Obtain a certificate authority (CA) certificate from which you intend to obtain a personal certificate, and then load the CA certificate in the NetScreen device.
- Obtain a local certificate (also known as a personal certificate) from the CA whose CA certificate you have previously loaded, and then load the local certificate in the NetScreen device. You can perform this task manually, or automatically using Simple Certificate Enrollment Protocol (SCEP).

Because the manual method of requesting local certificates has steps requiring you to copy information from one certificate to another, it can be a somewhat lengthy process. To bypass these steps, you can use the automatic method.

Note: *Before using SCEP, you must perform the following tasks:*

- *Configure and enable DNS (see "Domain Name System Support" on page **2**-495).*
- *Set the system clock (see "System Clock" on page **2**-541).*
- *Assign a host name and domain name to the NetScreen device. (If the NetScreen device is in an NSRP cluster, replace the host name with a cluster name. For more information, see "Cluster Name" on page **8**-17.)*

Example: Requesting a Local Certificate Automatically

In this example, you use the automatic method to request a certificate using SCEP from the Verisign CA. You set the following CA settings:

- Full certificate path validation
- RA CGI: http://ipsec.verisign.com/cgi-bin/pkiclient.exe[10]
- CA CGI: http://ipsec.verisign.com/cgi-bin/pkiclient.exe
- Automatic integrity confirmation of CA certificates
- CA ID, which identifies a SCEP server, where Verisign SCEP server uses a domain name, such as netscreen.com or a domain set up by Verisign for your company
- Challenge password
- Automatic certificate polling every 30 minutes (the default is no polling)

You then generate an RSA key pair, specifying a key length of 1024 bits, and initiate the SCEP operation to request a local certificate from the Verisign CA with the above CA settings.

When using the WebUI, you refer to CA certificates by name. When using the CLI, you refer to CA certificates by index (IDX) number. In this example, the IDX number for the Verisign CA is “1”. To see the IDX numbers for CA certificates, use the following command: **get pki x509 list ca-cert**. The output displays an IDX number and an ID number for each certificate. Note the IDX number and use that when referencing the CA certificate in commands.

10. The Common Gateway Interface (CGI) is a standard way for a web server to pass a user request to an application program, and to receive data back. CGI is part of the Hypertext Transfer Protocol (HTTP). You must specify an RA CGI path even if the RA does not exist. If the RA does not exist, use the value specified for the CA CGI.

WebUI

1. CA Server Settings

Objects > Certificates > Show CA > Server Settings (for Verisign): Enter the following, and then click **OK**:

X509 certificate path validation level: Full

SCEP Settings:

RA CGI: http://ipsec.verisigncom/cgi-bin/pkiclient.exe

CA CGI: http://ipsec.verisigncom/cgi-bin/pkiclient.exe

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic CA Server Settings configuration page:

Polling Interval: 30

Certificate Authentication: Auto

Certificate Renew: 14

2. Local Certificate Request

Objects > Certificates > New: Enter the following, and then click **Generate**:

Name: Michael Zhang

Phone: 408-730-6000

Unit/Department: Development

Organization: NetScreen Technologies

County/Locality: Santa Clara

State: CA

Country: US

Email: mzhang@netscreen.com

IP Address: 10.10.5.44

Create new key pair of *1024*[11] length: (select)

Issue the **get pki x509 pkcs** CLI command to have the NetScreen device generate a PKCS #10 file and then, do one of the following:

- Send the PKCS #10 certificate request file to an e-mail address
- Save it to disk
- Automatically enroll by sending the file to a CA that supports the Simple Certificate Enrollment Protocol (SCEP)

3. Automatic Enrollment

Select the **Automatically enroll to** option, select the **Existing CA server settings** option, and then select **Verisign** from the drop-down list.

Contact Verisign to inform them of your certificate request. They must authorize the certificate request before you can download the certificate.

CLI

1. CA Server Settings

```
set pki authority 1 cert-path full
set pki authority 1 scep ca-cgi “http://ipsec.verisign.com/cgi-bin
    /pkiclient.exe”[12]
set pki authority 1 scep ra-cgi “http://ipsec.verisign.com/cgi-bin
    /pkiclient.exe”[13]
set pki authority 1 scep polling-int 30
set pki authority 1 scep renew-start 14
```

11. The value 1024 indicates the bit length of the key pair. If you are using the certificate for SSL, be sure to use a bit length that your Web browser also supports.
12. The Common Gateway Interface (CGI) is a standard way for a web server to pass a user request to an application program, and to receive data back. CGI is part of the Hypertext Transfer Protocol (HTTP).
13. You must specify an RA CGI path even if the RA does not exist. If the RA does not exist, use the value specified for the CA CGI.

2. Local Certificate Request

```
set pki x509 dn country-name US
set pki x509 dn email mzhang@netscreen.com
set pki x509 dn ip 10.10.5.44
set pki x509 dn local-name “Santa Clara”
set pki x509 dn name “Michael Zhang”
set pki x509 dn org-name “NetScreen Technologies”
set pki x509 dn org-unit-name Development
set pki x509 phone 408-730-6000
set pki x509 dn state-name CA
exec pki rsa new 1024
```

3. Automatic Enrollment

```
exec pki x509 scep 1
```

If this is the first certificate request from this CA, a prompt appears presenting a fingerprint value for the CA certificate. You must contact the CA to confirm that this is the correct CA certificate.

Contact Verisign to inform them of your certificate request. They must authorize the certificate request before you can download the certificate.

Automatic Certificate Renewal

You can enable the NetScreen device to automatically renew certificates it acquired through SCEP (Simple Certificate Enrollment Protocol). This feature saves you from having to remember to renew certificates on the NetScreen device before they expire, and by the same token, helps maintain valid certificates at all times.

This feature is disabled by default. You can configure the NetScreen device to automatically send out a request to renew a certificate before it expires. You can set the time when you want the NetScreen device to send out the certificate renewal request in number of days and minutes before the expiration date. By setting different times for each certificate, you prevent the NetScreen device from having to renew all certificates at the same time.

For this feature to work, the NetScreen device must be able to reach the SCEP server, and the certificate must be present on the NetScreen device during the renewal process. Furthermore, for this feature to work, you must also ensure that the CA issuing the certificate can do the following:

- Support automatic approval of certificate requests.
- Return the same DN (Domain Name). In other words, the CA cannot modify the subject name and SubjectAltName extension in the new certificate.

You can enable and disable the automatic SCEP certificate renewal feature for all SCEP certificates or on a per-certificate basis.

Key Pair Generation

A NetScreen device holds pre-generated keys in memory. The number of pre-generated keys depends on the device model. During normal operation, the NetScreen device can manage to have enough keys available to renew certificates by generating a new key every time it uses one. The process of generating a key usually goes unnoticed as the device has time to generate a new key before one is needed. In the event that the NetScreen device renews a great number of certificates at once, thus using up keys rapidly, it might run out of pre-generated keys and have to generate them promptly for each new request. In this case, the generation of keys might affect the performance of the NetScreen device. Especially in a HA (High Availability) environment where the performance of the NetScreen device might slow down for a number of minutes.

The number of pre-generated key pairs on a NetScreen device depends on the model. For more information, refer to the specification sheet for your NetScreen product.

CHECKING FOR REVOCATION USING OCSP

When a NetScreen device performs an operation that uses a certificate, it may be necessary to check the certificate for premature revocation. The default way to check the revocation status of a digital certificate is to use CRL.

Online Certificate Status Protocol (OCSP) is an alternative way to check the status of a digital certificate. OCSP may provide additional information about the certificate. It may also provide the certificate status in a more timely manner.

When a NetScreen device uses OCSP, it is referred to as the *OCSP client* (or *requester*). This client sends a verification request to a server device called the *OCSP responder*. ScreenOS supports Verisign and Valicert as OCSP responders. The client's request contains the identity of the certificate to check. Before the NetScreen device can perform any OCSP operation, you must configure it to recognize the location of the OCSP responder.

After receiving the request, the OCSP responder confirms that the status information for the certificate is available, then returns the current status to the NetScreen device. The response of the OCSP responder contains the certificate's revocation status, the name of the responder, and the validity interval of the response. Unless the response is an error message, the responder signs the response using the responder's private key. The NetScreen device verifies the validity of the responder's signature by using the certificate of the responder. The certificate of the responder may either be embedded in the OCSP response, or stored locally and specified in the OCSP configuration. If the certificate is stored locally, use the following command to specify the locally stored certificate:

```
set pki authority id_num1 cert-status ocsp cert-verify id id_num2
```

id_num1 identifies the CA certificate that issued the certificate being verified, and *id_num2* identifies the locally stored certificate the device uses to verify the signature on the OCSP response.

If the certificate of the responder is not embedded in the OCSP response or stored locally, then the NetScreen device verifies the signature by using the CA certificate that issued the certificate in question.

Configuring for OCSP

You can use CLI commands to configure a NetScreen device to support OCSP operation. Most of these commands use an identification number to associate the revocation reference URL with the CA certificate. You can obtain this ID number using the following CLI command:

```
get pki x509 list ca-cert
```

Note: *The NetScreen device dynamically assigns the ID number to the CA certificate when you list the CA certificates. This number might change after you modify the certificate store.*

Specifying either CRL or OCSP for Revocation Checking

To specify the revocation check method (CRL, OCSP, or none) for a certificate of a particular CA, use the following CLI syntax:

```
set pki authority id_num cert-status revocation-check { CRL | OCSP | none }
```

where *id_num* is the identification number for the certificate.

The following example specifies OCSP revocation checking.

```
set pki authority 3 cert-status revocation-check ocsp
```

The ID number 3 identifies the certificate of the CA.

Displaying Certificate Revocation Status Attributes

To display the revocation check attributes for a particular CA, use the following CLI syntax:

```
get pki authority id_num cert-status
```

where *id_num* is the identification number for the certificate issued by the CA.

To display the revocation status attributes for the CA that issued certificate 7:

```
get pki authority 7 cert-status
```

Specifying the URL of an OCSP Responder for a Certificate

To specify the URL string of an OCSP responder for a particular certificate, use the following CLI syntax:

set pki authority *id_num* **cert-status ocsp url** *url_str*

To specify the URL string of an OCSP responder (http:\\192.168.10.10) for the CA with certificate at index 5, use the following CLI syntax:

```
set pki authority 5 cert-status ocsp url http:\\192.168.10.10
```

To remove the URL (http:\\192.168.2.1) of a CRL server for a certificate 5:

```
unset pki authority 5 cert-status ocsp url http:\\192.168.2.1
```

Removing Certificate Revocation Check Attributes

To remove all attributes related to a certificate revocation check for a CA that issued a particular certificate, use the following syntax:

unset pki authority *id_num* **cert-status**

To remove all revocation attributes related to certificate 1:

```
unset pki authority 1 cert-status
```

3

VPN Guidelines

NetScreen offers a variety of cryptographic options when you configure a VPN tunnel. Even when configuring a simple tunnel, you must make choices. The goal of the first half of this chapter is to summarize all the choices for a basic site-to-site VPN and a basic dialup VPN, and to present one or more reasons for choosing one option or another.

In the second half of the chapter, we explore the difference between policy-based and route-based VPN tunnels. We then examine the packet flow for a route-based and policy-based site-to-site AutoKey IKE VPN tunnel to see the outbound and inbound processing stages that a packet undergoes. The chapter concludes with some useful VPN configuration tips to keep in mind when configuring a tunnel.

The chapter is organized as follows:

- "Cryptographic Options" on page 40
  - "Site-to-Site Cryptographic Options" on page 41
  - "Dialup VPN Options" on page 50
- "Route- and Policy-Based Tunnels" on page 58
- "Packet Flow: Site-to-Site VPN" on page 60
- "Tunnel Configuration Tips" on page 67

CRYPTOGRAPHIC OPTIONS

When configuring a VPN, you must make many decisions about the cryptography you want to use. Questions arise about which Diffie-Hellman group is the right one to choose, which encryption algorithm provides the best balance between security and performance, and so on. This section presents all the cryptographic options required to configure a basic site-to-site VPN tunnel and a basic dialup VPN tunnel, and explains one or more benefits about each one to help you make your decisions.

The first decision that you must make is whether the tunnel is for a site-to-site VPN tunnel (between two NetScreen devices) or whether it is for a dialup VPN (from the NetScreen-Remote VPN client to a NetScreen device). Although this is a networking decision, the distinction between the two types of tunnels affects some cryptographic options. Therefore, the options are presented in two different decision trees:

- "Site-to-Site Cryptographic Options" on page 41
- "Dialup VPN Options" on page 50

After you decide whether you are going to configure a site-to-site tunnel or a dialup tunnel, you can refer to the appropriate decision tree for guidance. Each tree presents the cryptographic choices that you must make while configuring the tunnel. Following each tree are reasons for choosing each option that appears in the tree.

Note: *Examples for configuring both kinds of tunnels are in Chapter 4, "Site-to-Site VPNs" and Chapter 5, "Dialup VPNs".*

Site-to-Site Cryptographic Options

When configuring a basic site-to-site VPN tunnel, you must choose among the cryptographic options in the decision tree below. Advantages for each option follow.

Note: *Options highlighted in* ***purple*** *indicate NetScreen-recommended options. For background information about the different IPSec options, see Chapter 1, "IPSec".*

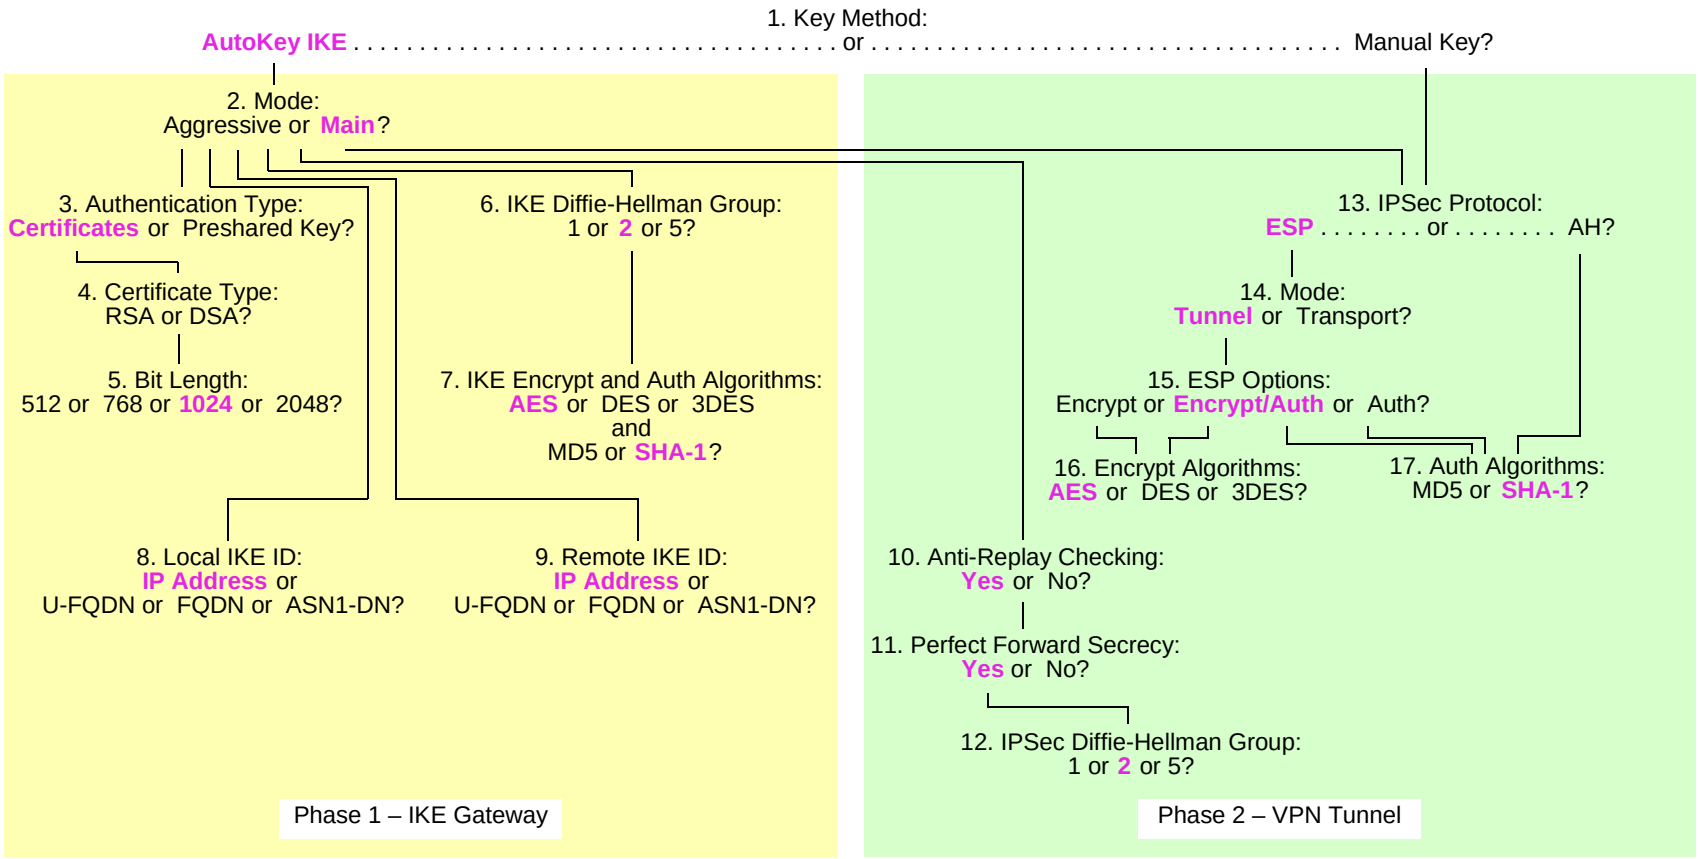

1. Key Method: Manual Key of AutoKey IKE?

AutoKey IKE

– Provides automatic key renewal and key freshness, thereby increasing security

Manual Key

– Useful for debugging IKE problems
– Eliminates IKE negotiation delays when establishing a tunnel

2. Mode: Aggressive or Main?

Aggressive

– Required when the IP address of one of the IPSec peers is dynamically assigned and a preshared key is used

Main

– Provides identity protection
– Can be used when the dialup user has a static IP address or if certificates are used for authentication

3. Authentication Type: Preshared Key or Certificates?

Certificates

– Greater security than provided by preshared keys because you can validate certificates with a certificate authority (CA). (For more information, see Chapter 2, “Public Key Cryptography”.)

Preshared Key

– Easier to use and faster to set up because it does not require a Public Key Infrastructure (PKI)

4. Certificate Type: RSA or DSA?

This depends on the CA from whom you get your certificates. There is no advantage of one certificate type over the other.

5. Bit Length: 512 or 768 or 1024 or 2048?

512

- Incurs the least processing overhead

768

- Provides more security than 512 bits
- Incurs less processing overhead than 1024 and 2048 bits

1024

- Provides more security than 512 and 768 bits
- Incurs less processing overhead than 2048 bits

2048

- Provides the most security

6. IKE Diffie-Hellman Group: 1 or 2 or 5?

Diffie-Hellman Group 1

- Incurs less processing overhead than Diffie-Hellman Groups 2 and 5
- Processing acceleration provided in NetScreen hardware

Diffie-Hellman Group 2

- Incurs less processing overhead than Diffie-Hellman Group 5
- Provides more security than Diffie-Hellman Group 1
- Processing acceleration provided in NetScreen hardware

Diffie-Hellman Group 5

- Provides the most security

7. IKE Encrypt and Auth Algorithms: AES or DES or 3DES and MD5 or SHA-1?

AES

- Cryptographically stronger than DES and 3DES if key lengths are all equal
- Processing acceleration provided in NetScreen hardware
- Approved encryption algorithm for Federal Information Processing Standards (FIPS) and Common Criteria EAL4 standards

DES

- Incurs less processing overhead than 3DES and AES
- Useful when the remote peer does not support AES

3DES

- Provides more cryptographic security than DES
- Processing acceleration provided in NetScreen hardware

MD5

- Incurs less processing overhead than SHA-1

SHA-1

- Provides more cryptographic security than MD5
- The only authentication algorithm that FIPS accepts

8. Local IKE ID: IP Address (Default) or U-FQDN or FQDN or ASN1-DN?

IP Address

- Can only be used if the local NetScreen device has a static IP address
- Default IKE ID when using a preshared key for authentication
- Can be used with a a certificate if the IP address appears in the SubjectAltName field

U-FQDN

- User-Fully Qualified Domain Name (U-FQDN—an e-mail address): Can be used with a preshared key or a certificate if the U-FQDN appears in the SubjectAltName field

FQDN

- Fully Qualified Domain Name (FQDN): Can be used with a preshared key or a certificate if the FQDN appears in the SubjectAltName field
- Useful for VPN gateways that have dynamic IP addresses
- Default IKE ID when using RSA or DSA certificates for authentication

ASN1-DN

- Can be used only with certificates
- Useful if the CA does not support the SubjectAltName field in the certificates it issues

9. Remote IKE ID: IP Address (Default) or U-FQDN or FQDN or ASN1-DN?

IP Address

- Does not require you to enter a remote IKE ID for a peer at a static IP address when using preshared keys for authentication and the peer is a NetScreen device
- Can be used for a device with a static IP address
- Can be used with a preshared key or a certificate if the IP address appears in the SubjectAltName field

U-FQDN

- User-Fully Qualified Domain Name (U-FQDN—an e-mail address): Can be used with a preshared key or a certificate if the U-FQDN appears in the SubjectAltName field

FQDN

- Fully Qualified Domain Name (FQDN): Can be used with a preshared key or a certificate if the FQDN appears in the SubjectAltName field
- Useful for VPN gateways that have dynamic IP addresses
- Does not require you to enter a remote IKE ID when using certificates for authentication and the peer is a NetScreen device

ASN1-DN

- Can be used only with certificates
- Useful if the CA does not support the SubjectAltName field in the certificates it issues

10. Anti-Replay Checking:

No or Yes?

Yes

- Enables the recipient to check sequence numbers in packet headers to prevent Denial-of-Service (DoS) attacks caused when a malefactor resends intercepted IPSec packets

No

- Disabling this might resolve compatibility issues with third-party peers

11. Perfect Forward Secrecy: No of Yes?

Yes

- Perfect Forward Secrecy (PFS): Provides increased security because the peers perform a second Diffie-Hellman exchange to produce the key used for IPSec encryption/decryption

No

- Provides faster tunnel setup
- Incurs less processing during Phase 2 IPSec negotiations

12. IPSec Diffie-Hellman Group: 1 or 2 or 5?

Diffie-Hellman Group 1

- Incurs less processing overhead than Diffie-Hellman Groups 2 and 5
- Processing acceleration provided in NetScreen hardware

Diffie-Hellman Group 2

- Incurs less processing overhead than Diffie-Hellman Group 5
- Provides more security than Diffie-Hellman Group 1
- Processing acceleration provided in NetScreen hardware

Diffie-Hellman Group 5

- Provides the most security

13. IPSec Protocol:

ESP or AH?

ESP

- Encapsulating Security Payload (ESP): Can provide both confidentiality through encryption and encapsulation of the original IP packet and integrity through authentication
- Can provide either encryption alone or authentication alone

AH

- Authentication Header (AH): Provides authentication of the entire IP packet, including the IPSec header and outer IP header

14. Mode: Tunnel or Transport?

Tunnel

- Conceals the original IP header, thereby increasing privacy

Transport

- Required for L2TP-over-IPSec tunnel support

15. ESP Options: Encrypt or Encrypt/Auth or Auth?

Encrypt

- Provides faster performance and incurs less processing overhead than using encrypt/auth
- Useful when you require confidentiality but do not require authentication

Encrypt/Auth

- Useful when you want confidentiality and authentication

Auth

- Useful when you want authentication but do not require confidentiality. Perhaps when the information is not secret, but it is important to establish that the information truly comes from the person who claims to send it and that nobody tampered with the content while in transit.

16. Encrypt Algorithms: AES or DES or 3DES?

AES

- Cryptographically stronger than DES and 3DES if key lengths are all equal
- Processing acceleration provided in NetScreen hardware
- Approved encryption algorithm for FIPS and Common Criteria EAL4 standards

DES

- Incurs less processing overhead than 3DES and AES
- Useful when the remote peer does not support AES

3DES

- Provides more cryptographic security than DES
- Processing acceleration provided in NetScreen hardware

17. Auth Algorithms: MD5 or SHA-1?

MD5

– Incurs less processing overhead than SHA-1

SHA-1

– Provides more cryptographic security than MD5

Using the recommended options from the above list, a generic site-to-site VPN configuration between two NetScreen devices with static IP addresses would consist of the following components:

- AutoKey IKE
- Main mode
- 1024-bit certificates (RSA or DSA)
- Phase 1 Diffie-Hellman Group 2
- Encryption = AES
- Authentication = SHA-1
- IKE ID = IP address (default)
- Anti-replay protection = yes
- Perfect Forward Secrecy (PFS) = yes
- Phase 2 Diffie-Hellman Group 2
- Encapsulating Security Payload (ESP)
- Tunnel mode
- Encryption/Authentication
- Encryption = AES
- Authentication = SHA-1

Dialup VPN Options

When configuring a basic dialup VPN tunnel, you must choose among the cryptographic options in the decision tree below. Advantages for each option follow.

Note: *Options highlighted in* ***purple*** *indicate NetScreen-recommended options. For background information about the different IPSec options, see Chapter 1, "IPSec".*

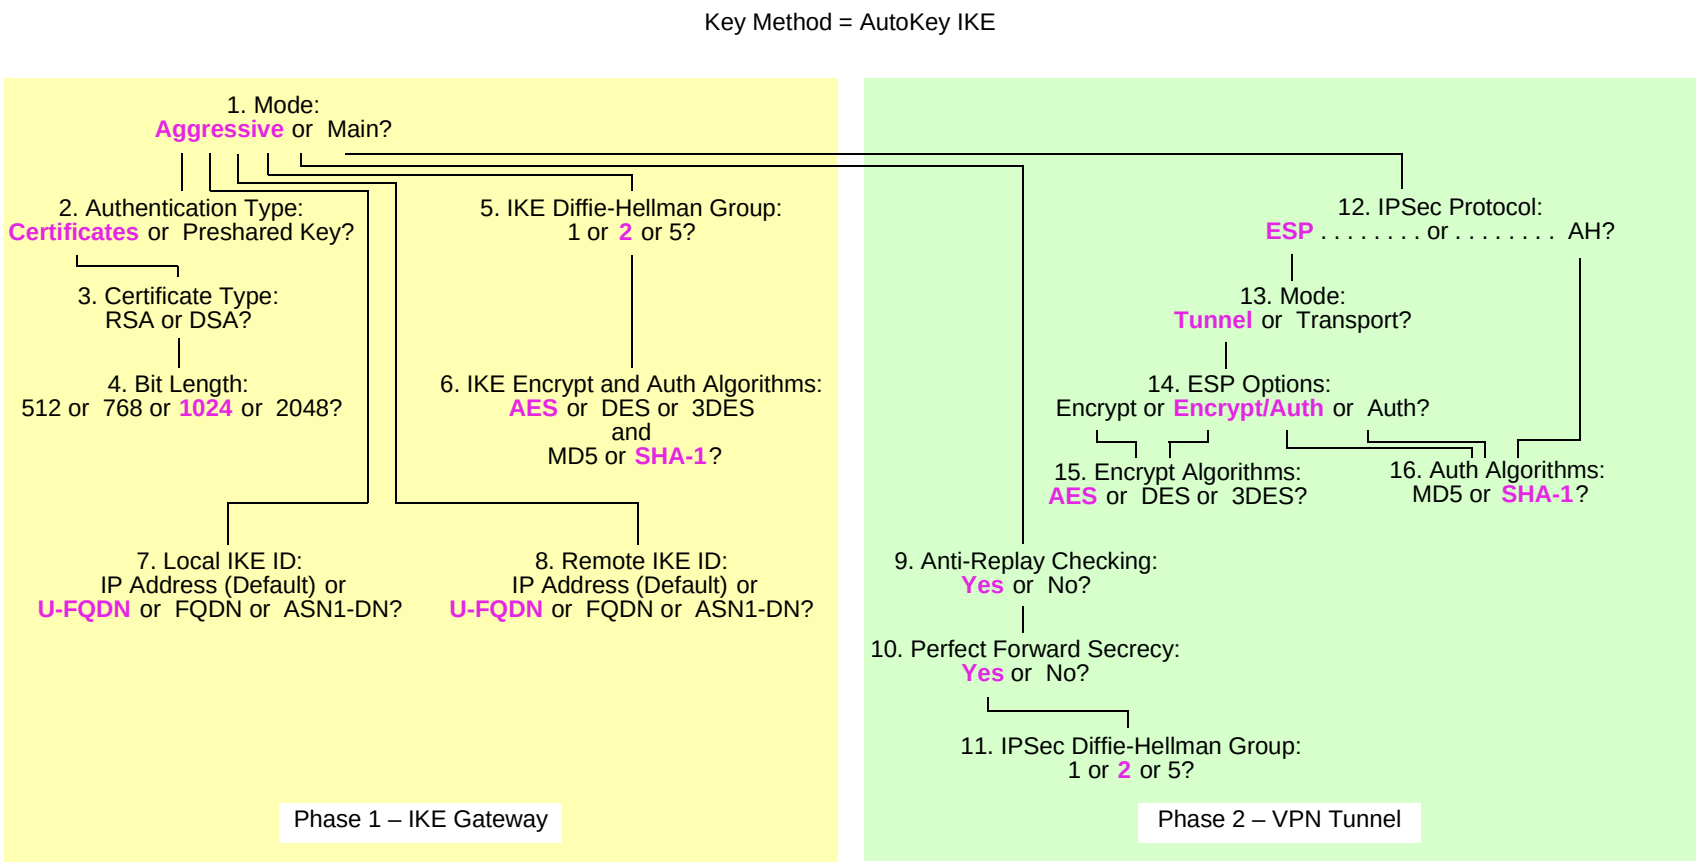

1. Mode: Aggressive or Main?

Aggressive

- Required when the IP address of one of the IPSec peers is dynamically assigned and a preshared key is used
- Can be used with either certificates or preshared keys for authentication

Main

- Provides identity protection

2. Authentication Type: Preshared Key or Certificates?

Certificates

- Greater security than provided by preshared keys because you can validate certificates with a certificate authority (CA). (For more information, see Chapter 2, “Public Key Cryptography”.)

Preshared Key

- Easier to use and faster to set up because it does not require a Public Key Infrastructure (PKI)

3. Certificate Type: RSA or DSA?

This depends on the CA from whom you get your certificates. There is no advantage of one certificate type over the other.

4. Bit Length: 512 or 768 or 1024 or 2048?

512

- Incurs the least processing overhead

768

- Provides more security than 512 bits
- Incurs less processing overhead than 1024 and 2048 bits

1024

- Provides more security than 512 and 768 bits
- Incurs less processing overhead than 2048 bits

2048

- Provides the most security

5. IKE Diffie-Hellman Group: 1 or 2 or 5?

Diffie-Hellman Group 1

- Incurs less processing overhead than Diffie-Hellman Groups 2 and 5
- Processing acceleration provided in NetScreen hardware

Diffie-Hellman Group 2

- Incurs less processing overhead than Diffie-Hellman Group 5
- Provides more security than Diffie-Hellman Group 1
- Processing acceleration provided in NetScreen hardware

Diffie-Hellman Group 5

- Provides the most security

6. IKE Encrypt and Auth Algorithms: AES or DES or 3DES and MD5 or SHA-1?

AES

- Cryptographically stronger than DES and 3DES if key lengths are all equal
- Processing acceleration provided in NetScreen hardware
- Approved encryption algorithm for FIPS and Common Criteria EAL4 standards

DES

- Incurs less processing overhead than 3DES and AES
- Useful when the remote peer does not support AES

3DES

- Provides more cryptographic security than DES
- Processing acceleration provided in NetScreen hardware

MD5

- Incurs less processing overhead than SHA-1

SHA-1

- Provides more cryptographic security than MD5

7. Local IKE ID: IP Address (Default) or U-FQDN or FQDN or ASN1-DN?

IP Address (Default)

- Does not require you to enter an IKE ID for a device with a static IP address
- Can be used for a device with a static IP address
- Can be used with a preshared key or a certificate if the IP address appears in the SubjectAltName field

U-FQDN

- User-Fully Qualified Domain Name (U-FQDN—an e-mail address): Can be used with a preshared key or a certificate if the U-FQDN appears in the SubjectAltName field

FQDN

- Fully Qualified Domain Name (FQDN): Can be used with a preshared key or a certificate if the FQDN appears in the SubjectAltName field
- Useful for VPN gateways that have dynamic IP addresses

ASN1-DN

- Can be used only with certificates
- Useful if the CA does not support the SubjectAltName field in the certificates it issues

8. Remote IKE ID: IP Address (Default) or U-FQDN or FQDN or ASN1-DN?

IP Address (Default)

- Does not require you to enter an IKE ID for a device with a static IP address
- Can be used for a device with a static IP address
- Can be used with a preshared key or a certificate if the IP address appears in the SubjectAltName field

U-FQDN

- User-Fully Qualified Domain Name (U-FQDN—an e-mail address): Can be used with a preshared key or a certificate if the U-FQDN appears in the SubjectAltName field

FQDN

- Fully Qualified Domain Name (FQDN): Can be used with a preshared key or a certificate if the FQDN appears in the SubjectAltName field
- Useful for VPN gateways that have dynamic IP addresses

ASN1-DN

- Can be used only with certificates
- Useful if the CA does not support the SubjectAltName field in the certificates it issues

9. Anti-Replay Checking: No or Yes?

Yes

- Enables the recipient to check sequence numbers in packet headers to prevent Denial-of-Service (DoS) attacks caused when a malefactor resends intercepted IPSec packets

No

- Disabling this might resolve compatibility issues with third-party peers

10. Perfect Forward Secrecy: No of Yes?

Yes

- Perfect Forward Secrecy (PFS): Provides increased security because the peers perform a second Diffie-Hellman exchange to produce the key used for IPSec encryption/decryption

No

- Provides faster tunnel setup
- Incurs less processing during Phase 2 IPSec negotiations

11. IPSec Diffie-Hellman Group: 1 or 2 or 5?

Diffie-Hellman Group 1

- Incurs less processing overhead than Diffie-Hellman Groups 2 and 5
- Processing acceleration provided in NetScreen hardware

Diffie-Hellman Group 2

- Incurs less processing overhead than Diffie-Hellman Group 5
- Provides more security than Diffie-Hellman Group 1
- Processing acceleration provided in NetScreen hardware

Diffie-Hellman Group 5

- Provides the most security

12. IPSec Protocol: ESP or AH?

ESP

- Encapsulating Security Payload (ESP): Can provide both confidentiality through encryption and encapsulation of the original IP packet and integrity through authentication
- Can provide either encryption alone or authentication alone

AH

- Authentication Header (AH): Provides authentication of the entire IP packet, including the IPSec header and outer IP header

13. Mode: Tunnel or Transport?

Tunnel

- Conceals the original IP header, thereby increasing privacy

Transport

- Required for L2TP-over-IPSec tunnel support

14. ESP Options: Encrypt or Encrypt/Auth or Auth?

Encrypt

- Provides faster performance and incurs less processing overhead than using encrypt/auth
- Useful when you require confidentiality but do not require authentication

Encrypt/Auth

- Useful when you want confidentiality and authentication

Auth

- Useful when you want authentication but do not require confidentiality. Perhaps when the information is not secret, but it is important to establish that the information truly comes from the person who claims to send it and that nobody tampered with the content while in transit.

15. Encrypt Algorithms: AES or DES or 3DES?

AES

- Cryptographically stronger than DES and 3DES if key lengths are all equal
- Processing acceleration provided in NetScreen hardware
- Approved encryption algorithm for FIPS and Common Criteria EAL4 standards

DES

- Incurs less processing overhead than 3DES and AES
- Useful when the remote peer does not support AES

3DES

- Provides more cryptographic security than DES
- Processing acceleration provided in NetScreen hardware

16. Auth Algorithms: MD5 or SHA-1?

MD5

- Incurs less processing overhead than SHA-1

SHA-1

- Provides more cryptographic security than MD5

Using the recommended options from the above list, a generic dialup VPN configuration between two NetScreen devices with static IP addresses would consist of the following components:

- Aggressive mode
- 1024-bit certificates (RSA or DSA)
- Phase 1 Diffie-Hellman Group 2
- Encryption = AES
- Authentication = SHA-1
- IKE ID = U-FQDN (e-mail address)
- Anti-replay protection = yes
- Perfect Forward Secrecy (PFS) = yes
- Phase 2 Diffie-Hellman Group 2
- Encapsulating Security Payload (ESP)
- Tunnel mode
- Encryption/Authentication
- Encryption = AES
- Authentication = SHA-1

ROUTE- AND POLICY-BASED TUNNELS

The configuration of a NetScreen device for VPN support is particularly flexible. You can create route-based and policy-based VPN tunnels. Additionally, each type of tunnel can use Manual Key or AutoKey IKE to manage the keys used for encryption and authentication.

With policy-based VPN tunnels, a tunnel is treated as an object (or a building block) that together with source, destination, service, and action, comprises a policy that permits VPN traffic. (Actually, the VPN policy action is *tunnel*, but the action *permit* is implied, if unstated). In a policy-based VPN configuration, a policy specifically references a VPN tunnel by name.

With route-based VPNs, the policy does not specifically reference a VPN tunnel. Instead, the policy references a destination address. When the NetScreen device does a route lookup to find the interface through which it must send traffic to reach that address, it finds a route via a tunnel interface, which is bound to a specific VPN tunnel[1].

Thus, with a policy-based VPN tunnel, you can consider a tunnel as an element in the construction of a policy. With a route-based VPN tunnel, you can consider a tunnel as a means for delivering traffic, and the policy as a method for either permitting or denying the delivery of that traffic.

The number of policy-based VPN tunnels that you can create is limited by the number of policies that the device supports. The number of route-based VPN tunnels that you create is limited by the number of route entries or the number of tunnel interfaces that the device supports—whichever number is lower.

A route-based VPN tunnel configuration is a good choice when you want to conserve tunnel resources while setting granular restrictions on VPN traffic. Although you can create numerous policies referencing the same VPN tunnel, each policy creates an individual IPSec security association (SA) with the remote peer, each of which counts as an individual VPN tunnel. With a route-based approach to VPNs, the regulation of traffic is not coupled to the means of its delivery. You can configure dozens of policies to regulate traffic flowing through a single VPN tunnel between two sites, and there is just one IPSec SA at work. Also, a route-based VPN configuration allows you to create policies referencing a destination reached through a VPN tunnel in which the action is *deny*, unlike a policy-based VPN configuration, in which—as stated earlier—the action must be *tunnel*, implying *permit*.

1. Typically, a tunnel interface is bound to a single tunnel. You can also bind a tunnel interface to multiple tunnels. For more information, see "Multiple Tunnels per Tunnel Interface" on page 326.

Another advantage that route-based VPNs offer is the exchange of dynamic routing information through VPN tunnels. You can enable an instance of a dynamic routing protocol, such as Border Gateway Protocol (BGP), on a tunnel interface that is bound to a VPN tunnel. The local routing instance exchanges routing information through the tunnel with a neighbor enabled on a tunnel interface bound to the other end.

When a tunnel does not connect large networks running dynamic routing protocols and you do not need to conserve tunnels or define various policies to filter traffic through the tunnel, a policy-based tunnel makes sense. Also, because there is no network beyond a dialup VPN client, policy-based VPN tunnels are good choices for dialup VPN configurations.

PACKET FLOW: SITE-TO-SITE VPN

To better understand how the various components comprising the creation of an IPSec tunnel work in relation to each other, this section looks at the processing of a packet flow through a tunnel—both when a NetScreen device sends outbound VPN traffic and when it receives inbound VPN traffic. The processing for a route-based VPN is presented, followed by an addendum noting the two places in the flow that differ for a policy-based VPN.

A company based in Tokyo has just opened a branch office in Paris and needs to connect the two sites through an IPSec tunnel. The tunnel uses AutoKey IKE, the ESP protocol, AES for encryption, SHA-1 for authentication using a preshared key, and has anti-replay checking enabled. The NetScreen devices protecting each site are in NAT mode, and all the zones are in the trust-vr routing domain. The addresses are as follows:

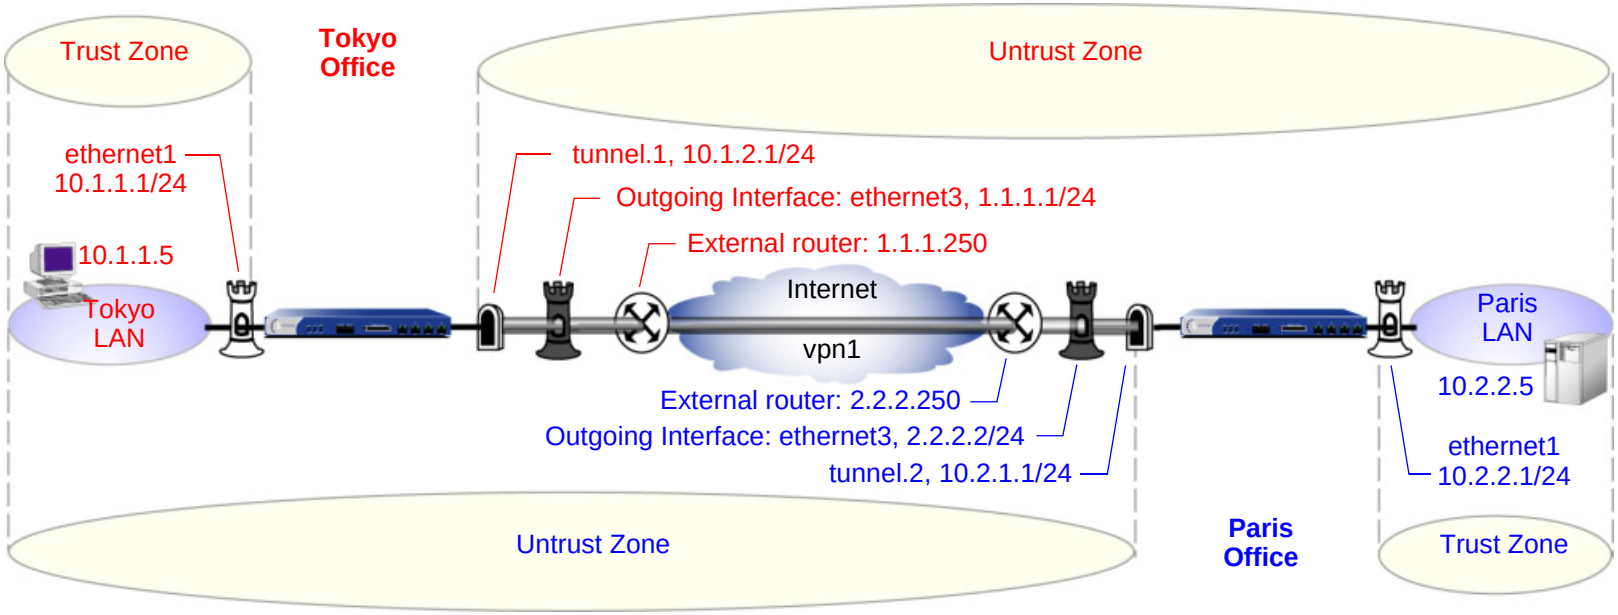

The path of a packet coming from 10.1.1.5/32 in the Tokyo LAN and going to 10.2.2.5/32 in the Paris LAN through an IPSec tunnel proceeds as described in the following subsections.

Tokyo (Initiator)

1. The host at 10.1.1.5 sends a packet destined for 10.2.2.5 to 10.1.1.1, which is the IP address ethernet1 and is the default gateway configured in the TCP/IP settings of host.
2. The packet arrives at ethernet1, which is bound to the Trust zone.
3. If you have enabled SCREEN options such as IP spoof detection for the Trust zone, the NetScreen device activates the SCREEN module at this point. SCREEN checking can produce one of the following three results:
   - If a SCREEN mechanism detects anomalous behavior for which it is configured to block the packet, the NetScreen device drops the packet and makes an entry in the event log.
   - If a SCREEN mechanism detects anomalous behavior for which it is configured to record the event but not block the packet, the NetScreen device records the event in the SCREEN counters list for ethernet1 and proceeds to the next step.
   - If the SCREEN mechanisms detect no anomalous behavior, the NetScreen device proceeds to the next step.

   If you have not enabled any SCREEN options for the Trust zone, the NetScreen device immediately proceeds to the next step.
4. The session module performs a session lookup, attempting to match the packet with an existing session.

   If the packet does not match an existing session, the NetScreen device performs First Packet Processing, a procedure involving the remaining steps.

   If the packet matches an existing session, the NetScreen device performs Fast Processing, using the information available from the existing session entry to process the packet. Fast Processing bypasses the route and policy lookups because the information generated by the bypassed steps has already been obtained during the processing of the first packet in the session.
5. The address-mapping module checks if a mapped IP (MIP) configuration uses the destination IP address 10.2.2.5. Because 10.2.2.5 is not used in a MIP configuration, the NetScreen device proceeds to the next step. (For information about packet processing when MIPs, VIPs, or destination address translation [NAT-dst] is involved, see “Packet Flow for Destination Translation” on page **2**-278.)

6. To determine the destination zone, the route module does a route lookup for 10.2.2.5. (The route module uses the ingress interface to determine which virtual router to use for the route lookup.) It finds a route entry directing traffic to 10.2.2.5 through the tunnel.1 interface bound to a VPN tunnel named “vpn1”. The tunnel interface is in the Untrust zone. By determining the ingress and egress interfaces, the NetScreen device has thereby determined the source and destination zones and can now do a policy lookup.
7. The policy engine does a policy lookup between the Trust and Untrust zones (as determined by the corresponding ingress and egress interfaces). The action specified in the policy matching the source address and zone, destination address and zone, and service is permit.
8. The IPSec module checks if an active Phase 2 security association (SA) exists with the remote peer. The Phase 2 SA check can produce either of the following results:
   – If the IPSec module discovers an active Phase 2 SA with that peer, it proceeds to step 10.
   – If the IPSec module does not discover an active Phase 2 SA with that peer, it drops the packet and triggers the IKE module.
9. The IKE module checks if an active Phase 1 SA exists with the remote peer. The Phase 1 SA check can produce either of the following results:
   – If the IKE module discovers an active Phase 1 SA with the peer, it uses this SA to negotiate a Phase 2 SA.
   – If the IKE module does not discover an active Phase 1 SA with that peer, it begins Phase 1 negotiations in Main mode, and then Phase 2 negotiations.
10. The IPSec module puts an ESP header and then an outer IP header on the packet. Using the address specified as the outgoing interface, it puts 1.1.1.1 as the source IP address in the outer header. Using the address specified for remote gateway, it puts 2.2.2.2 as the destination IP address in the outer header. Next, it encrypts the packet from the payload to the next header field in the original IP header. Then, it authenticates the packet from the ESP trailer to the ESP header.
11. The NetScreen device sends the encrypted and authenticated packet destined for 2.2.2.2 through the outgoing interface (ethernet3) to the external router at 1.1.1.250.

Paris (Recipient)

1. The packet arrives at 2.2.2.2, which is the IP address of ethernet3, an interface bound to the Untrust zone.
2. Using the SPI, destination IP address, and IPSec protocol contained in the outer packet header, the IPSec module attempts to locate an active Phase 2 SA with the initiating peer along with the keys to authenticate and decrypt the packet. The Phase 2 SA check can produce one of the following three results:
   - If the IPSec module discovers an active Phase 2 SA with the peer, it proceeds to step 4.
   - If the IPSec module does not discover an active Phase 2 SA with the peer but it can match an inactive Phase 2 SA using the source IP address but not the SPI, it drops the packet, makes an event log entry, and sends a notification that it received a bad SPI to the initiating peer.
   - If the IPSec module does not discover an active Phase 2 SA with that peer, it drops the packet and triggers the IKE module.
3. The IKE module checks if an active Phase 1 SA exists with the remote peer. The Phase 1 SA check can produce either of the following results:
   - If the IKE module discovers an active Phase 1 SA with the peer, it uses this SA to negotiate a Phase 2 SA.
   - If the IKE module does not discover an active Phase 1 SA with that peer, it begins Phase 1 negotiations in Main mode, and then Phase 2 negotiations.
4. The IPSec module performs an anti-replay check. This check can produce one of two results:
   - If the packet fails the anti-replay check, because it detects a sequence number that the NetScreen device has already received, the NetScreen device drops the packet.
   - If the packet passes the anti-replay check, the NetScreen device proceeds to the next step.
5. The IPSec module attempts to authenticate the packet. The authentication check can produce one of two results:
   - If the packet fails the authentication check, the NetScreen device drops the packet.
   - If the packet passes the authentication check, the NetScreen device proceeds to the next step.
6. Using the Phase 2 SA and keys, the IPSec module decrypts the packet, uncovering its original source address (10.1.1.5) and its ultimate destination (10.2.2.5). It learns that the packet came through vpn1, which is bound to tunnel.1. From this point forward, the NetScreen device treats the packet as if its ingress interface is tunnel.1 instead of ethernet3. It also adjusts the anti-replay sliding window at this point.

7. If you have enabled SCREEN options for the Untrust zone, the NetScreen device activates the SCREEN module at this point. SCREEN checking can produce one of the following three results:
   - If a SCREEN mechanism detects anomalous behavior for which it is configured to block the packet, the NetScreen device drops the packet and makes an entry in the event log.
   - If a SCREEN mechanism detects anomalous behavior for which it is configured to record the event but not block the packet, the NetScreen device records the event in the SCREEN counters list for ethernet3 and proceeds to the next step.
   - If the SCREEN mechanisms detect no anomalous behavior, the NetScreen device proceeds to the next step.
8. The session module performs a session lookup, attempting to match the packet with an existing session. It then either performs First Packet Processing or Fast Processing.

   If the packet matches an existing session, the NetScreen device performs Fast Processing, using the information available from the existing session entry to process the packet. Fast Processing bypasses all but the last two steps (encrypting the packet and forwarding it) because the information generated by the bypassed steps has already been obtained during the processing of the first packet in the session.
9. The address-mapping module checks if a mapped IP (MIP) or virtual IP (VIP) configuration uses the destination IP address 10.2.2.5. Because 10.2.2.5 is not used in a MIP or VIP configuration, the NetScreen device proceeds to the next step.
10. The route module first uses the ingress interface to determine the virtual router to use for the route lookup; in this case, the trust-vr. It then performs a route lookup for 10.2.2.5 in the trust-vr and discovers that it is accessed through ethernet1. By determining the ingress interface (tunnel.1) and the egress interface (ethernet1), the NetScreen device can thereby determine the source and destination zones. The tunnel.1 interface is bound to the Untrust zone, and ethernet1 is bound to the Trust zone. The NetScreen device can now do a policy lookup.
11. The policy engine checks its policy list from the Untrust zone to the Trust zone and finds a policy that grants access.
12. The NetScreen device forwards the packet through ethernet1 to its destination at 10.2.2.5.

Addendum: Policy-Based VPN

The packet flow for a policy-based VPN configuration differs from that of a route-based VPN configuration at two points: the route lookup and the policy lookup.

Tokyo (Initiator)

The first stages of the outbound packet flow are the same for both route-based and policy-based VPN configurations until the route lookup and subsequent policy lookup occur:

Route Lookup: To determine the destination zone, the route module does a route lookup for 10.2.2.5. Not finding an entry for that specific address, the route module resolves it to a route through ethernet3, which is bound to the Untrust zone. By determining the ingress and egress interfaces, the NetScreen device has thereby determined the source and destination zones, and can now perform a policy lookup.

Policy Lookup: The policy engine does a policy lookup between the Trust and Untrust zones. The lookup matches the source address and zone, destination address and zone, and service and finds a policy that references a VPN tunnel named vpn1.

The NetScreen device then forwards the packet through ethernet1 to its destination at 10.2.2.5.

Paris (Recipient)

Most stages of the inbound packet flow on the recipient's end are the same for both route-based and policy-based VPN configurations except that the tunnel is not bound to a tunnel interface, but to a tunnel zone. The NetScreen device learns that the packet came through vpn1, which is bound to the Untrust-Tun tunnel zone, whose carrier zone is the Untrust zone. Unlike route-based VPNs, the NetScreen device considers ethernet3 to be the ingress interface of the decrypted packet—not tunnel.1.

The flow changes after packet decryption is complete. At this point, the route and policy lookups differ:

Route Lookup: The route module performs a route lookup for 10.2.2.5 and discovers that it is accessed through ethernet1, which is bound to the Trust zone. By learning that the Untrust zone is the source zone (because vpn1 is bound to the Untrust-Tun tunnel zone, whose carrier zone is the Untrust zone) and by determining the destination zone based on the egress interface (ethernet1 is bound to

the Trust zone), the NetScreen device can now check for a policy from the Untrust to the Trust zones that references vpn1.

Policy Lookup: The policy engine checks its policy list from the Untrust zone to the Trust zone and finds a policy that references a VPN tunnel named vpn1 and that grants access to 10.2.2.5.

The NetScreen device then forwards the packet to its destination.

TUNNEL CONFIGURATION TIPS

This section offers some guidelines, or tips, to keep in mind when configuring VPN tunnels. When configuring an IPSec VPN tunnel, keep the following points in mind:

- NetScreen supports up to four proposals for Phase 1 negotiations and up to four proposals for Phase 2 negotiations. A peer must be configured to accept at least one Phase 1 proposal and one Phase 2 proposal proffered by the other peer. For information about Phase 1 and Phase 2 IKE negotiations, see "Tunnel Negotiation" on page 11.
- If you want to use certificates for authentication and there is more than one local certificate loaded on the NetScreen device, you must specify which certificate you want each VPN tunnel configuration to use. For more information about certificates, see Chapter 2, "Public Key Cryptography" on page 15.
- For a basic policy-based VPN:
  - Use user-defined addresses in the policy, not the pre-defined address "Any".
  - The addresses and service specified in policies configured at both ends of the VPN must match.
  - Use symmetric policies for bidirectional VPN traffic.
- The proxy ID for both peers must match, which means that the service specified in the proxy ID for both peers is the same, and the local IP address specified for one peer is the same as the remote IP address specified for the other peer[2].
  - For a route-based VPN configuration, the proxy ID is user configurable.
  - For a policy-based VPN configuration, the NetScreen device—by default—derives the proxy ID from the source address, destination address, and service specified in the policy that references that VPN tunnel in the policy list. You can also define a proxy ID for a policy-based VPN that supersedes the derived proxy ID.

  The simplest way to ensure that the proxy IDs match is to use 0.0.0.0/0 for the local address, 0.0.0.0/0 for the remote address, and "any" for the service. Instead of using the proxy ID for access control, you use policies to control the traffic to and from the VPN. For examples of VPN configurations with user-configurable proxy IDs, see the route-based VPN examples in Chapter 4, "Site-to-Site VPNs".

2. The proxy ID is a three-part tuple consisting of local IP address-remote IP address-service.

- As long as the peers' proxy ID settings match, it does not matter if one peer defines a route-based VPN and the other defines a policy-based VPN. If peer-1 uses a policy-based VPN configuration and peer-2 uses a route-based VPN configuration, then peer-2 must define a proxy ID that matches the proxy ID derived from peer-1's policy[3]. If peer-1 performs source network address translation (NAT-src) using a DIP pool, use the address and netmask for the DIP pool as the remote address in peer-2's proxy ID. For example:

When the DIP pool is:	Use this in the proxy ID:
1.1.1.8 – 1.1.1.8	1.1.1.8/32
1.1.1.20 – 1.1.1.50	1.1.1.20/26
1.1.1.100 – 1.1.1.200	1.1.1.100/25
1.1.1.0 – 1.1.1.255	1.1.1.0/24

  For more information about proxy IDs when used with NAT-src and NAT-dst, see "VPN Sites with Overlapping Addresses" on page 168.

- Because proxy IDs support either a single service or all services, the service in a proxy ID derived from a policy-based VPN referencing a service group is considered as "any".
- When both peers have static IP addresses, they can each use the default IKE ID, which is their IP addresses. When a peer or dialup user has a dynamically assigned IP address, that peer or user must use another type of IKE ID. An FQDN is a good choice for a dynamic peer and a U-FQDN (e-mail address) is a good choice for a dialup user. You can use both FQDN and U-FQDN IKE ID types with preshared keys and certificates (if the FQDN or U-FQDN appears in the SubjectAltName field in the certificate). If you use certificates, the dynamic peer or dialup user can also use all or part of the ASN1-DN as the IKE ID.

3. Peer-1 can also define a proxy ID that matches peer-2's proxy ID. Peer-1's user-defined proxy ID supersedes the proxy ID that the NetScreen device derives from the policy components.

4

Site-to-Site VPNs

This chapter explains how to configure a site-to-site virtual private network (VPN) tunnel between two NetScreen devices. It examines route-based and policy-based VPN tunnels, presents the various elements that you must consider when setting up a tunnel, and offers several examples.

- "Site-to-Site VPN Configurations" on page 70
  - "Site-to-Site Tunnel Configuration Steps" on page 71
  - "Example: Route-Based Site-to-Site VPN, AutoKey IKE" on page 77
  - "Example: Policy-Based Site-to-Site VPN, AutoKey IKE" on page 91
  - "Example: Route-Based Site-to-Site VPN, Dynamic Peer" on page 102
  - "Example: Policy-Based Site-to-Site VPN, Dynamic Peer" on page 117
  - "Example: Route-Based Site-to-Site VPN, Manual Key" on page 131
  - "Example: Policy-Based Site-to-Site VPN, Manual Key" on page 142
- "FQDN for Dynamic IKE Gateways" on page 151
  - "Example: AutoKey IKE Peer with FQDN" on page 153
- "VPN Sites with Overlapping Addresses" on page 168
  - "Example: Tunnel Interface with NAT-Src and NAT-Dst" on page 171
- "Transparent Mode VPN" on page 186
  - "Example: Transparent Mode, Policy-Based AutoKey IKE VPN" on page 187

SITE-TO-SITE VPN CONFIGURATIONS

An IPSec VPN tunnel exists between two gateways, and each gateway needs an IP address. When both gateways have static IP addresses, you can configure the following kinds of tunnels:

- Site-to-Site VPN, AutoKey IKE tunnel (with a preshared key or certificates)
- Site-to-Site VPN, Manual Key tunnel

When one gateway has a static address and the other has a dynamically assigned address, you can configure the following kind of tunnel:

- Dynamic Peer Site-to-Site VPN, AutoKey IKE tunnel (with a preshared key or certificates)

As used here, a static site-to-site VPN involves an IPSec tunnel connecting two sites, each with a NetScreen device operating as a secure gateway. The physical interface or subinterface used as the outgoing interface on both devices has a fixed IP address, and the internal hosts also have static IP addresses. If the NetScreen device is in Transparent mode, it uses the VLAN1 address as the IP address for the outgoing interface. With a static site-to-site VPN, hosts at either end of the tunnel can initiate the VPN tunnel setup because the IP address of the remote gateway remains constant and thus reachable.

If the outgoing interface of one of the NetScreen devices has a dynamically assigned IP address, that device is termed a dynamic peer and the VPN is configured differently. With a dynamic peer site-to-site VPN, only hosts behind the dynamic peer can initiate the VPN tunnel setup because only their remote gateway has a fixed IP address and is thus reachable from their local gateway. However, after a tunnel is established between a dynamic peer and a static peer, hosts behind either gateway can initiate VPN traffic if the destination hosts have fixed IP addresses.

Note: *For background information about the available VPN options, see Chapter 1, "IPSec". For guidance when choosing among the various options, see Chapter 3, "VPN Guidelines".*

Site-to-Site Tunnel Configuration Steps

The configuration of a site-to-site VPN tunnel requires the coordination of the tunnel configuration with that of other settings—interfaces, addresses, routes, and policies. The three example VPN configurations in this section are set in the following context: an office in Tokyo wants to communicate securely with an office in Paris through an IPSec VPN tunnel.

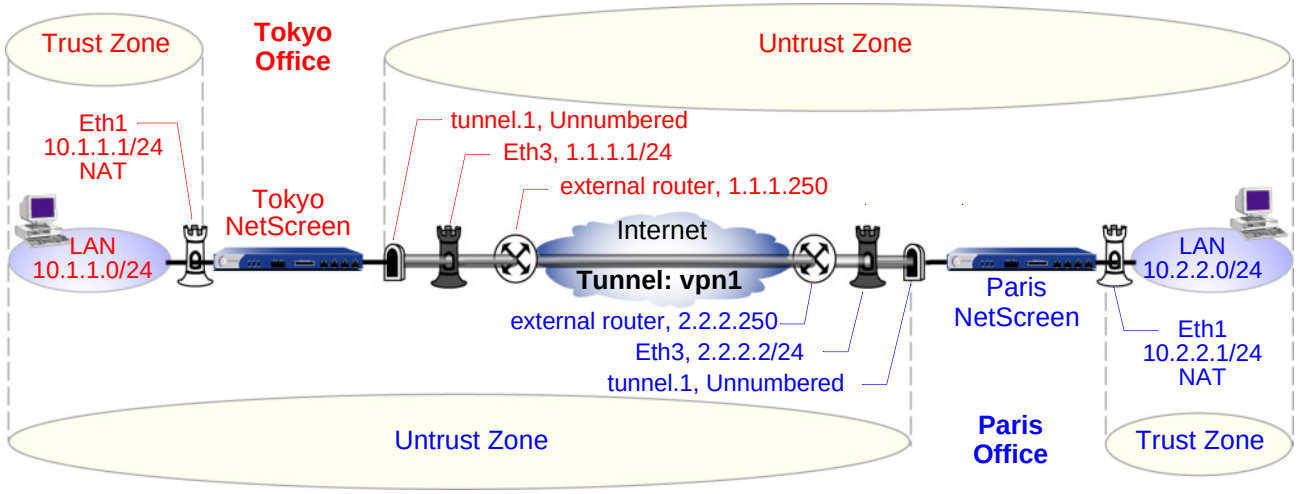

The administrators in both offices configure the following settings:

- Interfaces – Security Zones and Tunnel
- Addresses
- VPN (one of the following)
  - AutoKey IKE
  - Dynamic Peer
  - Manual Key
- Routes
- Policies

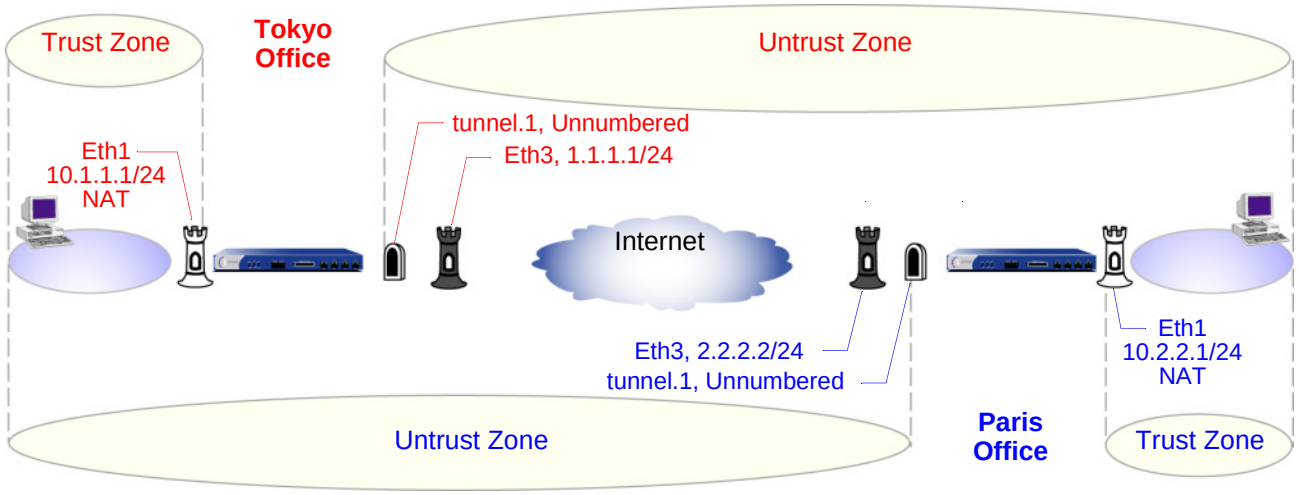

1. Interfaces – Security Zones and Tunnel

The admin at the Tokyo office configures the security zone and tunnel interfaces with the settings that appear in red in the above illustration. The admin at the Paris office does likewise with the settings that appear in blue.

Ethernet3 is going to be the outgoing interface for VPN traffic and the remote gateway for VPN traffic sent from the other end of the tunnel.

Ethernet1 is in NAT mode so each admin can assign IP addresses to all the internal hosts, yet when traffic passes from the Trust zone to the Untrust zone, the NetScreen device translates the source IP address in the packet headers to the address of the Untrust zone interface, ethernet3—1.1.1.1 for Tokyo, and 2.2.2.2 for Paris.

For a route-based VPN, each admin binds the tunnel interface tunnel.1 to the VPN tunnel vpn1. By defining a route to the address space of the remote office LAN, the NetScreen device can direct all traffic bound for that LAN to the tunnel.1 interface and thus through the tunnel to which tunnel.1 is bound.

Because policy-based NAT services are not needed, a route-based VPN configuration does not require tunnel.1 to have an IP address/netmask, and a policy-based VPN configuration does not even require a tunnel interface.

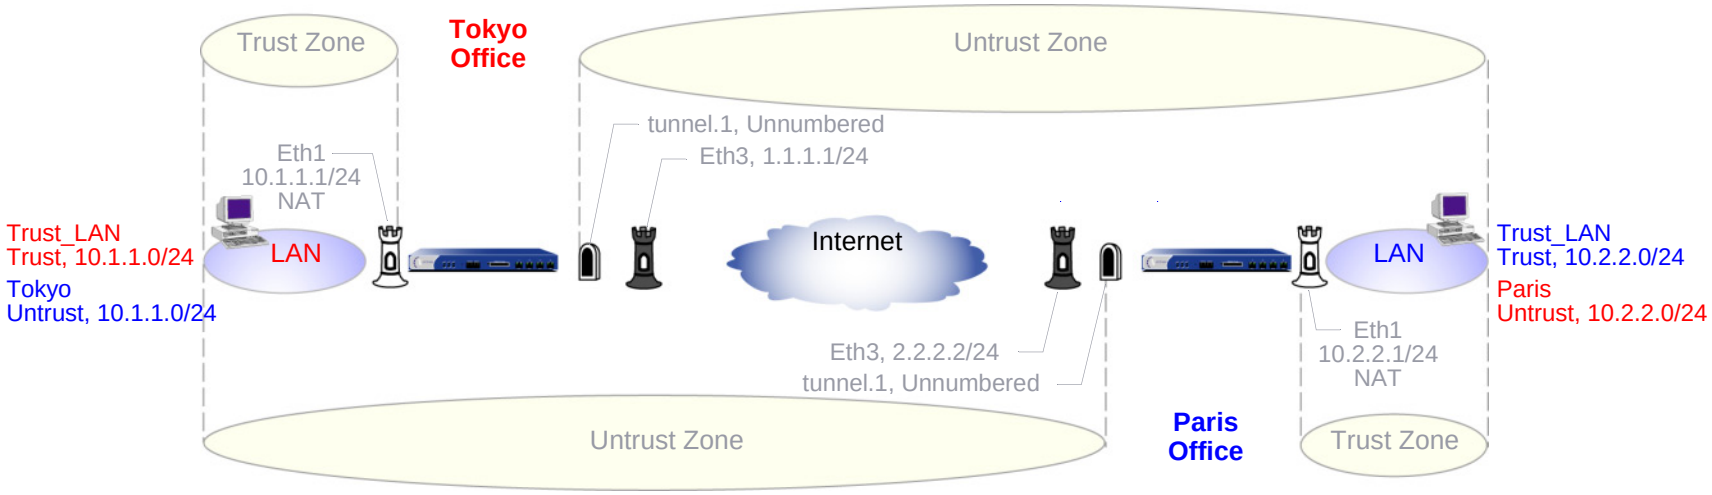

2. Addresses

The admins define addresses for later use in inbound and outbound policies. The admin at the Tokyo office defines the addresses that appear in red in the above illustration. The admin at the Paris office does likewise with the addresses that appear in blue.

For policy-based VPNs, the NetScreen device derives proxy IDs from policies[1]. Because the proxy IDs used by the NetScreen devices at both ends of the VPN tunnel must match perfectly, you cannot use the predefined address “ANY”, whose IP address is 0.0.0.0/0, at one end of the tunnel if you use a more specific address at the other end. For example,

If the proxy ID in Tokyo is …		and the proxy ID in Paris is …	then the proxy IDs do not
From: 0.0.0.0/0	✗	To: 10.1.1.0/24	match and IKE
To: 10.2.2.0/24	✓	From: 10.2.2.0/24	negotiations will fail.
Service: ANY	✓	Service: ANY	

For route-based VPNs, you can use “0.0.0.0/0–0.0.0.0/0–any” to define the local and remote IP addresses and service type for a proxy ID. You can then use more restrictive policies to filter the inbound and outbound VPN traffic by source address, destination address, and service type.

1. In ScreenOS 5.0.0, you can also define proxy IDs for VPN tunnels referenced in policy-based VPN configurations.

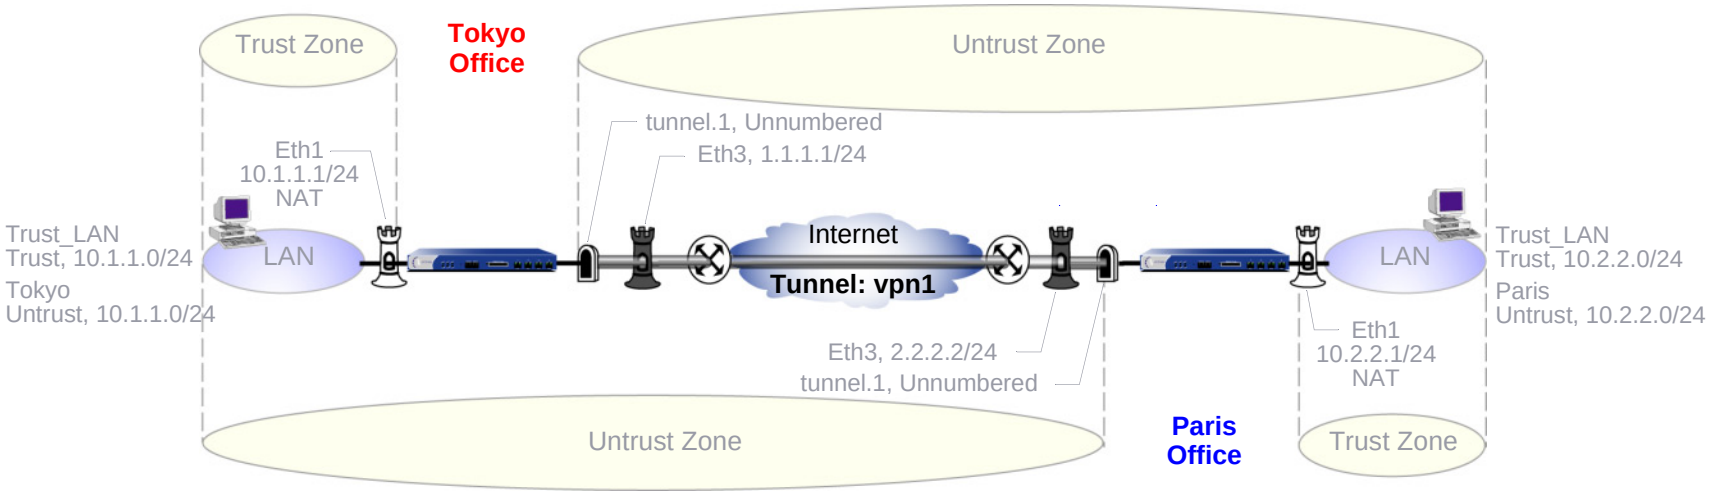

3. VPN

You can configure one of the following three VPNs:

- AutoKey IKE

  The AutoKey IKE method uses a preshared key or a certificate to refresh—that is, change—the encryption and authentication keys automatically at user-defined intervals (known as key lifetimes). Essentially, frequently updating these keys strengthens security, although excessively short lifetimes might reduce overall performance.

- Dynamic Peer

  A dynamic peer is a remote gateway that has a dynamically assigned IP address. Because the IP address of the remote peer might be different each time IKE negotiations begin, hosts behind the peer must initiate VPN traffic. Also—if using a preshared key for authentication—the peer must send an IKE ID during the first message of Phase 1 negotiations in aggressive mode to identify itself.

- Manual Key

  The Manual Key method requires you to set and update the encryption and authentication keys manually. This method is a viable option for a small set of VPN tunnels.

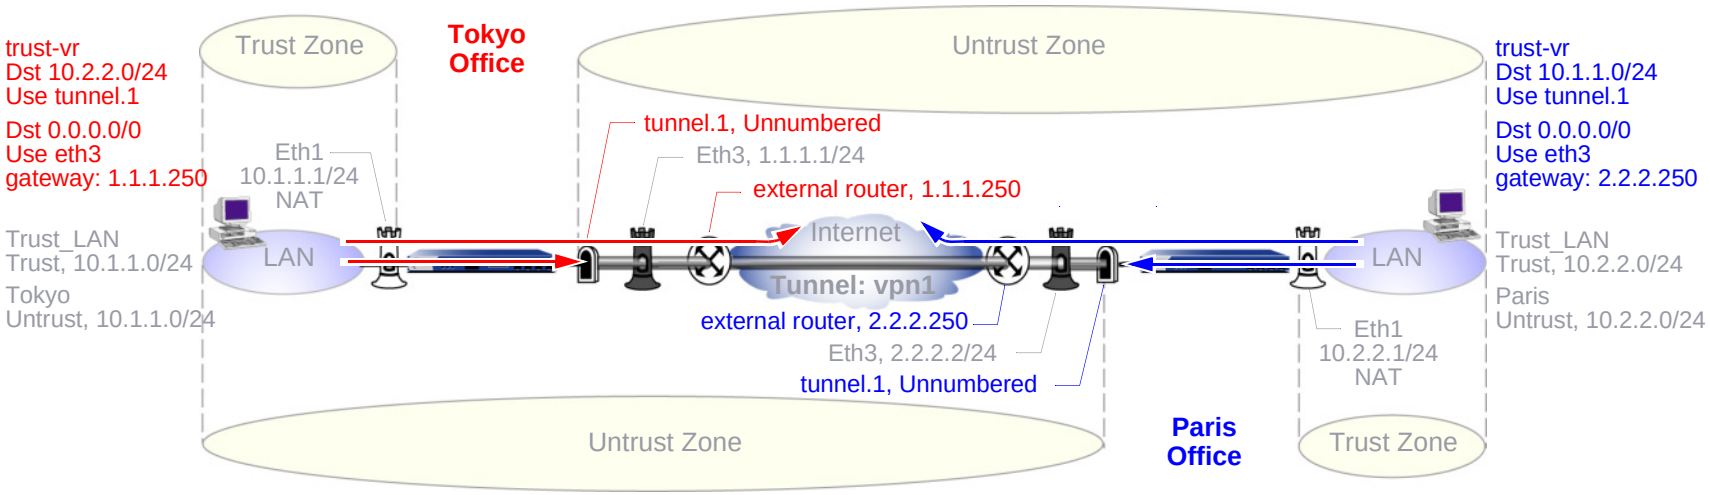

4. Routes

The admins at each site must configure at least the following two routes:

- A route for traffic to reach an address on the remote LAN to use tunnel.1
- A default route for all other traffic, including the outer VPN tunnel traffic, to reach the internet via ethernet3 and then the external router beyond it—1.1.1.250 for the Tokyo office and 2.2.2.250 for Paris[2]. The external router is the default gateway to which the NetScreen device forwards any traffic for which it does not have a specific route in its routing table.

2. If the NetScreen device at the Tokyo office receives its external IP address dynamically from its ISP (that is, from the point of view of the Paris office, the NetScreen device at the Tokyo office is its dynamic peer), then the ISP automatically provides the Tokyo NetScreen with its default gateway IP address.

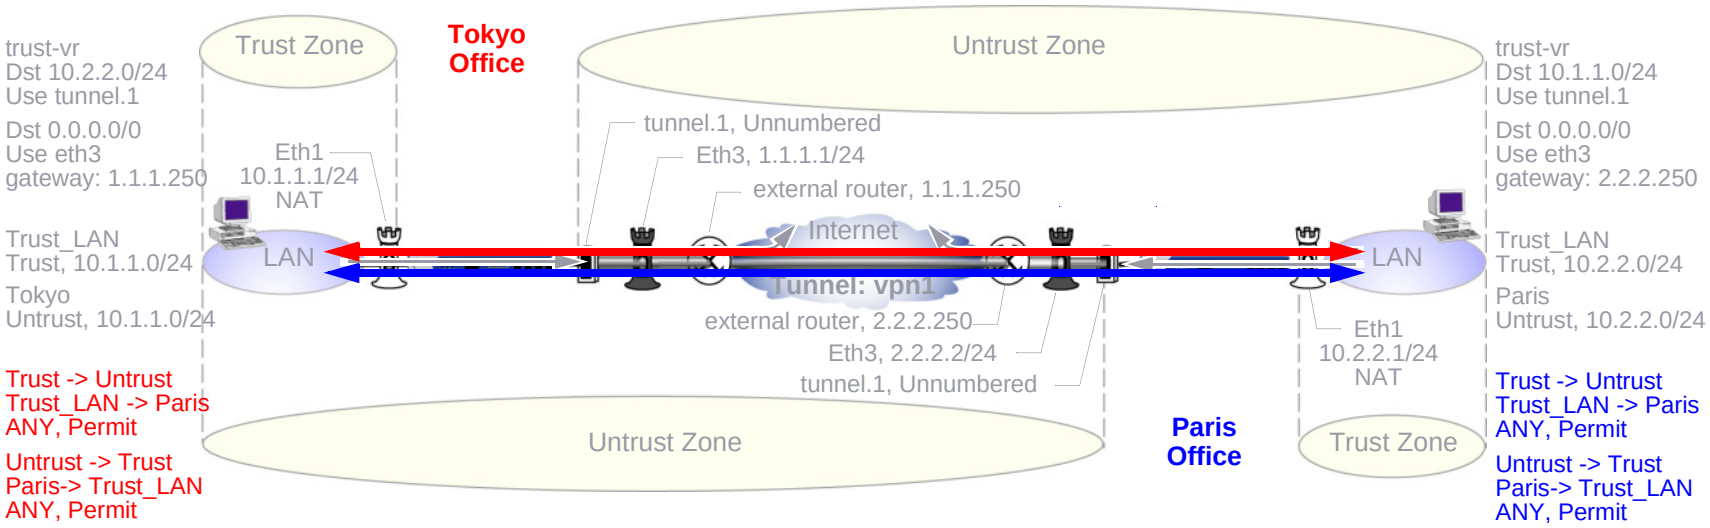

5. Policies

The admins at each site define policies to permit traffic between the two offices:

- A policy permitting any kind of traffic from “Trust_LAN” in the Trust zone to “Paris” or “Tokyo” in the Untrust zone
- A policy permitting any kind of traffic from “Paris” or “Tokyo” in the Untrust zone to “Trust_LAN” in the Trust zone

Because the route to the remote site specifies tunnel.1, which is bound to the VPN tunnel vpn1, the policy does not need to reference the VPN tunnel.

Example: Route-Based Site-to-Site VPN, AutoKey IKE

In this example, an AutoKey IKE tunnel using either a preshared secret or a pair of certificates (one at each end of the tunnel) provides the secure connection between the Tokyo and Paris offices. For the Phase 1 and 2 security levels, you specify one Phase 1 proposal—either pre-g2-3des-sha for the preshared key method or rsa-g2-3des-sha for certificates—and select the predefined "Compatible" set of proposals for Phase 2. All zones are in the trust-vr.

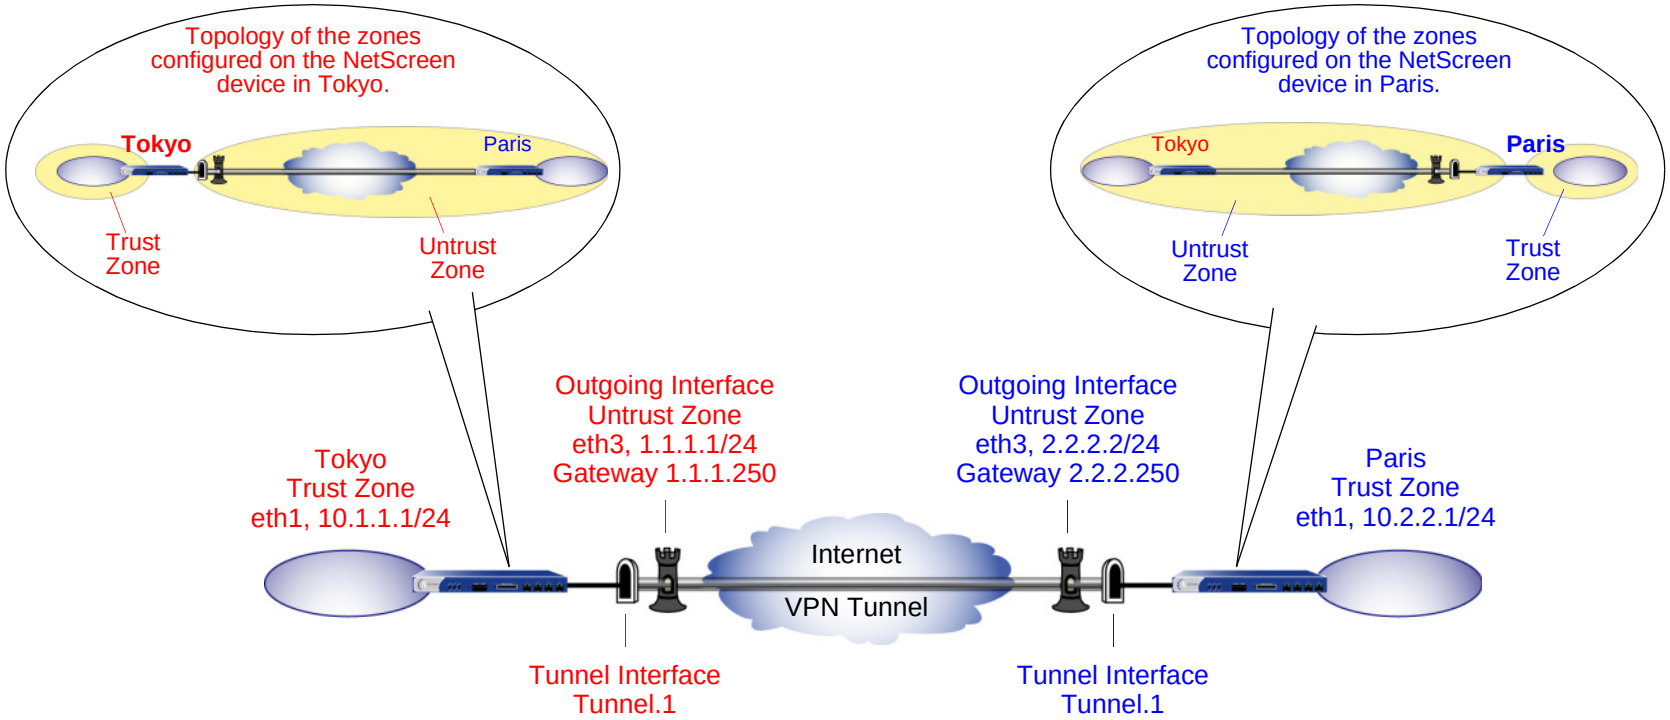

Setting up a route-based AutoKey IKE tunnel using either a preshared secret or certificates involves the following steps:

1. Assign IP addresses to the physical interfaces bound to the security zones and to the tunnel interface.
2. Configure the VPN tunnel, designate its outgoing interface in the Untrust zone, bind it to the tunnel interface, and configure its proxy-ID.

3. Enter the IP addresses for the local and remote endpoints in the address books for the Trust and Untrust zones.
4. Enter a default route to the external router in the trust-vr, and a route to the destination via the tunnel interface.
5. Set up policies for VPN traffic to pass between each site.

In the following examples, the preshared key is h1p8A24nG5. It is assumed that both participants already have RSA certificates and are using Entrust as the certificate authority (CA). (For information about obtaining and loading certificates, see "Certificates and CRLs" on page 21.)

WebUI (Tokyo)

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

Network > Interfaces > New Tunnel IF: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.1

Zone (VR): Untrust (trust-vr)

Unnumbered: (select)

Interface: ethernet3 (trust-vr)

2. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Trust_LAN

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.0/24

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Paris_Office

IP Address/Domain Name:

IP/Netmask: (select), 10.2.2.0/24

Zone: Untrust

3. VPN

VPNs > AutoKey Advanced > Gateway > New: Enter the following, and then click **OK**:

Gateway Name: To_Paris

Security Level: Custom

Remote Gateway Type:

Static IP Address: (select), IP Address/Hostname: 2.2.2.2

Preshared Key

Preshared Key: h1p8A24nG5

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Gateway configuration page:

Security Level: Custom

Phase 1 Proposal (for Custom Security Level): pre-g2-3des-sha

Mode (Initiator): Main (ID Protection)

(or)

Certificates

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Gateway configuration page:

Security Level: Custom

Phase 1 Proposal (for Custom Security Level): rsa-g2-3des-sha

Preferred certificate (optional)

Peer CA: Entrust

Peer Type: X509-SIG

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: Tokyo_Paris

Security Level: Compatible

Remote Gateway:

Predefined: (select), To_Paris

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Security Level: Compatible

Bind to: Tunnel Interface, tunnel.1

Proxy-ID: (select)

Local IP / Netmask: 10.1.1.0/24

Remote IP / Netmask: 10.2.2.0/24

Service: ANY

4. Routes

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 1.1.1.250

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 10.2.2.0/24

Gateway: (select)

Interface: Tunnel.1

Gateway IP Address: 0.0.0.0

5. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Name: To Paris

Source Address: Trust_LAN

Destination Address: Paris_Office

Service: ANY

Action: Permit

Position at Top: (select)

Policies > (From: Untrust, To: Trust) > New: Enter the following, and then click **OK**:

Name: From Paris

Source Address: Paris_Office

Destination Address: Trust_LAN

Service: ANY

Action: Permit

Position at Top: (select)

WebUI (Paris)

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.2.2.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 2.2.2.2/24

Network > Interfaces > New Tunnel IF: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.1

Zone (VR): Untrust (trust-vr)

Unnumbered: (select)

Interface: ethernet3 (trust-vr)

2. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Trust_LAN

IP Address/Domain Name:

IP/Netmask: (select), 10.2.2.0/24

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Tokyo_Office

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.0/24

Zone: Untrust

3. VPN

VPNs > AutoKey Advanced > Gateway > New: Enter the following, and then click **OK**:

Gateway Name: To_Tokyo

Security Level: Custom

Remote Gateway Type:

Static IP Address: (select), IP Address/Hostname: 1.1.1.1

Preshared Key

Preshared Key: h1p8A24nG5

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Gateway configuration page:

Security Level: Custom

Phase 1 Proposal (For Custom Security Level):
pre-g2-3des-sha

Mode (Initiator): Main (ID Protection)

(or)

Certificates

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Gateway configuration page:

Security Level: Custom

Phase 1 Proposal (for Custom Security Level): rsa-g2-3des-sha

Preferred certificate (optional)

Peer CA: Entrust

Peer Type: X509-SIG

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

Name: Paris_Tokyo

Security Level: Custom

Remote Gateway:

Predefined: (select), To_Tokyo

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Security Level: Compatible

Bind to: Tunnel Interface, tunnel.1

Proxy-ID: (select)

Local IP / Netmask: 10.2.2.0/24

Remote IP / Netmask: 10.1.1.0/24

Service: ANY

4. Routes

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 2.2.2.250

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 10.1.1.0/24

Gateway: (select)

Interface: Tunnel.1

Gateway IP Address: 0.0.0.0

5. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Name: To Tokyo

Source Address:

Address Book Entry: (select), Trust_LAN

Destination Address:

Address Book Entry: (select), Tokyo_Office

Service: ANY

Action: Permit

Position at Top: (select)

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Name: From Tokyo

Source Address:

Address Book Entry: (select), Tokyo_Office

Destination Address:

Address Book Entry: (select), Trust_LAN

Service: ANY

Action: Permit

Position at Top: (select)

CLI (Tokyo)

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24

set interface tunnel.1 zone untrust
set interface tunnel.1 ip unnumbered interface ethernet3
```

2. Addresses

```
set address trust Trust_LAN 10.1.1.0/24
set address untrust Paris_Office 10.2.2.0/24
```

3. VPN

Preshared Key

```
set ike gateway To_Paris address 2.2.2.2 main outgoing-interface ethernet3
    preshare h1p8A24nG5 proposal pre-g2-3des-sha
set vpn Tokyo_Paris gateway To_Paris sec-level compatible
set vpn Tokyo_Paris bind interface tunnel.1
set vpn Tokyo_Paris proxy-id local-ip 10.1.1.0/24 remote-ip 10.2.2.0/24 any
```

(or)

Certificate

```
set ike gateway To_Paris address 2.2.2.2 main outgoing-interface ethernet3
    proposal rsa-g2-3des-sha
set ike gateway To_Paris cert peer-ca 1[3]
set ike gateway To_Paris cert peer-cert-type x509-sig
set vpn Tokyo_Paris gateway To_Paris sec-level compatible
set vpn Tokyo_Paris bind interface tunnel.1
set vpn Tokyo_Paris proxy-id local-ip 10.1.1.0/24 remote-ip 10.2.2.0/24 any
```

4. Routes

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
set vrouter trust-vr route 10.2.2.0/24 interface tunnel.1
```

5. Policies

```
set policy top name “To Paris” from trust to untrust Trust_LAN Paris_Office any
    permit
set policy top name “From Paris” from untrust to trust Paris_Office Trust_LAN
    any permit
save
```

3. The number 1 is the CA ID number. To discover the CA’s ID number, use the following command: **get ike ca**.

CLI (Paris)

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.2.2.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 2.2.2.2/24

set interface tunnel.1 zone untrust
set interface tunnel.1 ip unnumbered interface ethernet3
```

2. Addresses

```
set address trust Trust_LAN 10.2.2.0/24
set address untrust Tokyo_Office 10.1.1.0/24
```

3. VPN

Preshared Key

```
set ike gateway To_Tokyo address 1.1.1.1 main outgoing-interface ethernet3
    preshare h1p8A24nG5 proposal pre-g2-3des-sha
set vpn Paris_Tokyo gateway To_Tokyo sec-level compatible
set vpn Paris_Tokyo bind interface tunnel.1
set vpn Paris_Tokyo proxy-id local-ip 10.2.2.0/24 remote-ip 10.1.1.0/24 any
```

(or)

Certificate

```
set ike gateway To_Tokyo address 1.1.1.1 main outgoing-interface ethernet3
    proposal rsa-g2-3des-sha
set ike gateway To_Tokyo cert peer-ca 1
set ike gateway To_Tokyo cert peer-cert-type x509-sig
set vpn Paris_Tokyo gateway To_Tokyo sec-level compatible
set vpn Paris_Tokyo bind interface tunnel.1
set vpn Paris_Tokyo proxy-id local-ip 10.2.2.0/24 remote-ip 10.1.1.0/24 any
```

4. Routes

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 2.2.2.250
set vrouter trust-vr route 10.1.1.0/24 interface tunnel.1
```

5. Policies

```
set policy top name “To Tokyo” from trust to untrust Trust_LAN Tokyo_Office any
    permit
set policy top name “From Tokyo” from untrust to trust Tokyo_Office Trust_LAN
    any permit
save
```

Example: Policy-Based Site-to-Site VPN, AutoKey IKE

In this example, an AutoKey IKE tunnel using either a preshared secret or a pair of certificates (one at each end of the tunnel) provides the secure connection between the Tokyo and Paris offices. For the Phase 1 and 2 security levels, you specify one Phase 1 proposal—either pre-g2-3des-sha for the preshared key method or rsa-g2-3des-sha for certificates—and select the predefined “Compatible” set of proposals for Phase 2. All zones are in the trust-vr.

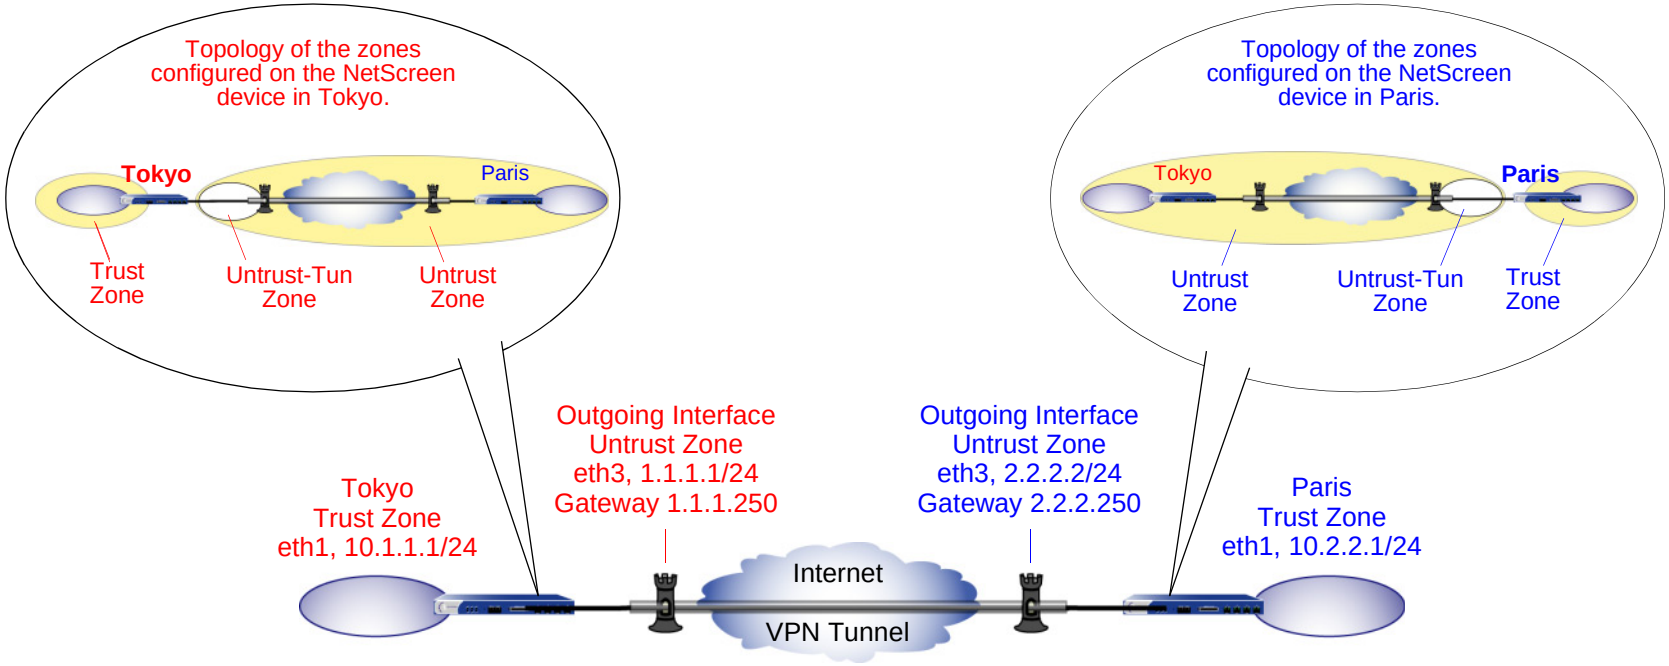

Setting up the AutoKey IKE tunnel using AutoKey IKE with either a preshared secret or certificates involves the following steps:

1. Define security zone interface IP addresses.
2. Make address book entries for the local and remote end entities.

3. Define the remote gateway and key exchange mode, and specify either a preshared secret or a certificate.
4. Create the Autokey IKE VPN.
5. Set a default route to the external router.
6. Configure policies.

In the following examples, the preshared key is h1p8A24nG5. It is assumed that both participants already have RSA certificates and are using Entrust as the certificate authority (CA). (For information about obtaining and loading certificates, see "Certificates and CRLs" on page 21.)

WebUI (Tokyo)

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

2. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Trust_LAN
IP Address/Domain Name:
IP/Netmask: (select), 10.1.1.0/24
Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Paris_Office
IP Address/Domain Name:
IP/Netmask: (select), 10.2.2.0/24
Zone: Untrust

3. VPN

VPNs > AutoKey Advanced > Gateway > New: Enter the following, and then click **OK**:

Gateway Name: To_Paris
Security Level: Custom
Remote Gateway Type:
Static IP Address: (select), IP Address/Hostname: 2.2.2.2

Preshared Key

Preshared Key: h1p8A24nG5
Outgoing Interface: ethernet3
> Advanced: Enter the following advanced settings, and then click **OK** to return to the basic Gateway configuration page:
Security Level: Custom
Phase 1 Proposal (For Custom Security Level):
pre-g2-3des-sha
Mode (Initiator): Main (ID Protection)

(or)

Certificates

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **OK** to return to the basic Gateway configuration page:

Security Level: Custom

Phase 1 Proposal (For Custom Security Level): rsa-g2-3des-sha

Mode (Initiator): Main (ID Protection)

Preferred certificate (optional)

Peer CA: Entrust

Peer Type: X509-SIG

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: Tokyo_Paris

Security Level: Compatible

Remote Gateway: Predefined: (select), To_Paris

4. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 1.1.1.250

5. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Name: To/From Paris

Source Address:

Address Book Entry: (select), Trust_LAN

Destination Address:

Address Book Entry: (select), Paris_Office

Service: ANY

Action: Tunnel

Tunnel VPN: Tokyo_Paris

Modify matching bidirectional VPN policy: (select)

Position at Top: (select)

WebUI (Paris)

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.2.2.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 2.2.2.2/24

2. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Trust_LAN

IP Address/Domain Name:

IP/Netmask: (select), 10.2.2.0/24

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Tokyo_Office

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.0/24

Zone: Untrust

3. VPN

VPNs > AutoKey Advanced > Gateway > New: Enter the following, and then click **OK**:

Gateway Name: To_Tokyo

Security Level: Custom

Remote Gateway Type:

Static IP Address: (select), IP Address/Hostname: 1.1.1.1

Preshared Key

Preshared Key: h1p8A24nG5

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Gateway configuration page:

Security Level: Custom

Phase 1 Proposal (For Custom Security Level): pre-g2-3des-sha

Mode (Initiator): Main (ID Protection)

(or)

Certificates

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Gateway configuration page:

Security Level: Custom

Phase 1 Proposal (For Custom Security Level): rsa-g2-3des-sha

Mode (Initiator): Main (ID Protection)

Preferred certificate (optional)

Peer CA: Entrust

Peer Type: X509-SIG

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: Paris_Tokyo

Security Level: Compatible

Remote Gateway: Predefined: (select), To_Tokyo

4. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 2.2.2.250

5. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Name: To/From Tokyo

Source Address:

Address Book Entry: (select), Trust_LAN

Destination Address:

Address Book Entry: (select), Tokyo_Office

Service: ANY

Action: Tunnel

Tunnel VPN: Paris_Tokyo

Modify matching bidirectional VPN policy: (select)

Position at Top: (select)

CLI (Tokyo)

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
```

2. Addresses

```
set address trust Trust_LAN 10.1.1.0/24
set address untrust paris_office 10.2.2.0/24
```

3. VPN

Preshared Key

```
set ike gateway to_paris address 2.2.2.2 main outgoing-interface ethernet3
    preshare h1p8A24nG5 proposal pre-g2-3des-sha
set vpn tokyo_paris gateway to_paris sec-level compatible
```

(or)

Certificates

```
set ike gateway to_paris address 2.2.2.2 main outgoing-interface ethernet3
    proposal rsa-g2-3des-sha
set ike gateway to_paris cert peer-ca 1[4]
set ike gateway to_paris cert peer-cert-type x509-sig
set vpn tokyo_paris gateway to_paris sec-level compatible
```

4. The number 1 is the CA ID number. To discover the CA's ID number, use the following command: **get ike ca**.

4. Route

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
```

5. Policies

```
set policy top name “To/From Paris” from trust to untrust Trust_LAN
    paris_office any tunnel vpn tokyo_paris
set policy top name “To/From Paris” from untrust to trust paris_office
    Trust_LAN any tunnel vpn tokyo_paris
save
```

CLI (Paris)

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.2.2.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 2.2.2.2/24
```

2. Addresses

```
set address trust Trust_LAN 10.2.2.0/24
set address untrust tokyo_office 10.1.1.0/24
```

3. VPN

Preshared Key

```
set ike gateway to_tokyo address 1.1.1.1 main outgoing-interface ethernet3
    preshare h1p8A24nG5 proposal pre-g2-3des-sha
set vpn paris_tokyo gateway to_tokyo sec-level compatible
```

(or)

Certificates

```
set ike gateway to_tokyo address 1.1.1.1 main outgoing-interface ethernet3
    proposal rsa-g2-3des-sha
set ike gateway to_tokyo cert peer-ca 1
set ike gateway to_tokyo cert peer-cert-type x509-sig
set vpn paris_tokyo gateway to_tokyo tunnel proposal nopfs-esp-3des-sha
```

4. Route

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 2.2.2.250
```

5. Policies

```
set policy top name “To/From Tokyo” from trust to untrust Trust_LAN
    tokyo_office any tunnel vpn paris_tokyo
set policy top name “To/From Tokyo” from untrust to trust tokyo_office
    Trust_LAN any tunnel vpn paris_tokyo
save
```

Example: Route-Based Site-to-Site VPN, Dynamic Peer

In this example, an AutoKey IKE VPN tunnel using either a preshared key or a pair of certificates (one at each end of the tunnel) provides a secure connection between NetScreen devices protecting the Tokyo and Paris offices. The Untrust zone interface for the NetScreen device at the Paris office has a static IP address. The ISP serving the Tokyo office assigns the IP address for the Untrust zone interface dynamically via DHCP. Because only the Paris NetScreen device has a fixed address for its Untrust zone, VPN traffic must originate from hosts in the Tokyo office. After a tunnel has been established, traffic through the tunnel can originate from either end. All security and tunnel zones are in the trust-vr.

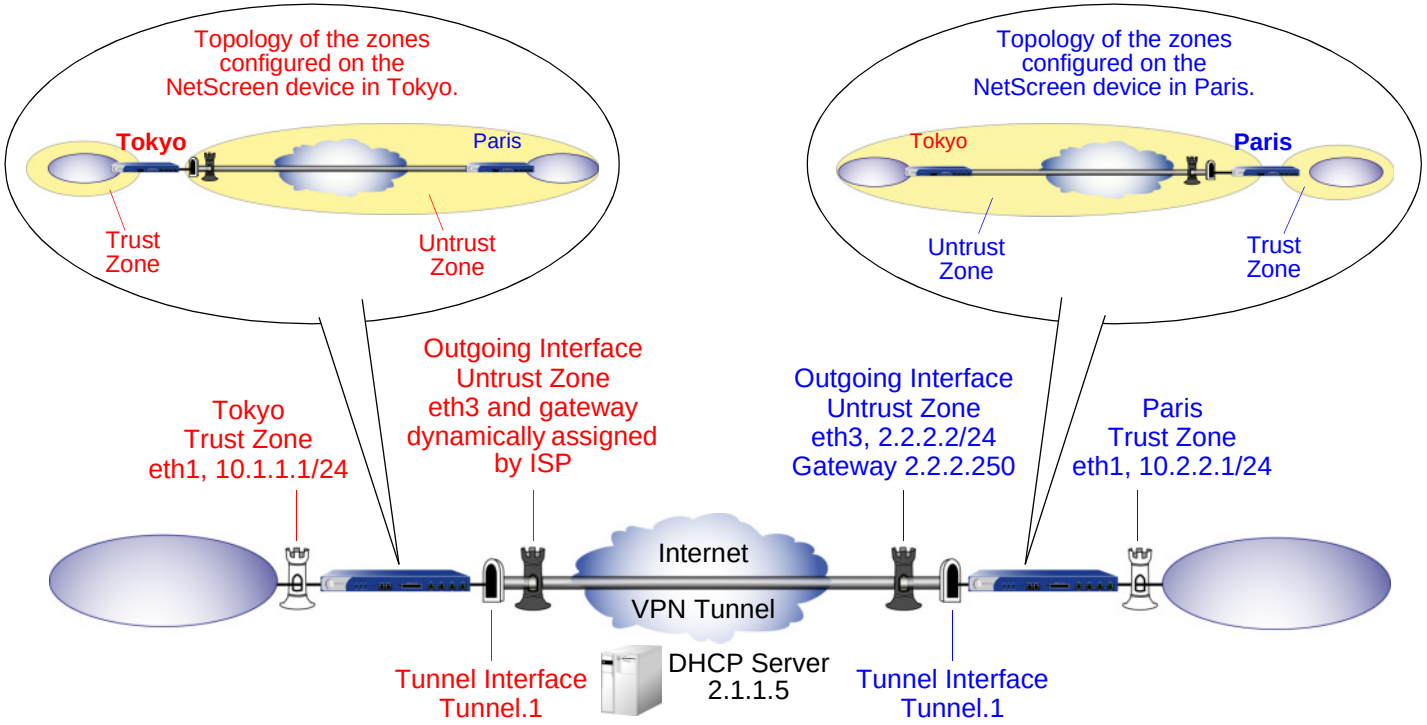

The preshared key is h1p8A24nG5. It is assumed that both participants already have RSA certificates from the certificate authority (CA) Verisign, and that the e-mail address *pmason@abc.com* appears in the local certificate on NetScreen-A. (For information about obtaining and loading certificates, see "Certificates and CRLs" on page 21.) For the Phase 1 and 2 security levels, you specify one Phase 1 proposal—either pre-g2-3des-sha for the preshared key method or rsa-g2-3des-sha for certificates—and select the "Compatible" set of proposals for Phase 2.

WebUI (Tokyo)

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **Apply**:

Zone Name: Untrust

Enter the following, and then click **OK**:

Obtain IP using DHCP: (select)[5]

Network > Interfaces > New Tunnel IF: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.1

Zone (VR): Untrust (trust-vr)

Unnumbered: (select)

Interface: ethernet3 (trust-vr)

5. You cannot specify the IP address of the DHCP server through the WebUI; however, you can do so through the CLI.

2. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Trust_LAN

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.0/24

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Paris_Office

IP Address/Domain Name:

IP/Netmask: (select), 10.2.2.0/24

Zone: Untrust

3. VPN

VPNs > AutoKey Advanced > Gateway > New: Enter the following, and then click **OK**:

Gateway Name: To_Paris

Security Level: Custom

Remote Gateway Type:

Static IP Address: (select), IP Address/Hostname: 2.2.2.2

Preshared Key

Preshared Key: h1p8A24nG5

Local ID: pmason@abc.com

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Gateway configuration page:

Security Level: Custom

Phase 1 Proposal (For Custom Security Level): pre-g2-3des-sha

Mode (Initiator): Aggressive

(or)

Certificates

Local ID: pmason@abc.com[6]

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Gateway configuration page:

Security Level: Custom

Phase 1 Proposal (For Custom Security Level): rsa-g2-3des-sha

Mode (Initiator): Aggressive

Preferred Certificate (optional):

Peer CA: Verisign

Peer Type: X509-SIG

6. The U-FQDN “pmason@abc.com” must appear in the SubjectAltName field in the certificate.

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: Tokyo_Paris

Security Level: Compatible

Remote Gateway:

Predefined: (select), To_Paris

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Bind to: Tunnel Interface: (select), tunnel.1

Proxy-ID: (select)

Local IP / Netmask: 10.1.1.0/24

Remote IP / Netmask: 10.2.2.0/24

Service: ANY

4. Routes

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 0.0.0.0[7]

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 10.2.2.0/24

Gateway: (select)

Interface: Tunnel.1

Gateway IP Address: 0.0.0.0

7. The ISP provides the gateway IP address dynamically through DHCP.

5. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Trust_LAN

Destination Address:

Address Book Entry: (select), Paris_Office

Service: Any

Action: Permit

Position at Top: (select)

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Paris_Office

Destination Address:

Address Book Entry: (select), Trust_LAN

Service: Any

Action: Permit

Position at Top: (select)

WebUI (Paris)

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.2.2.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address: 2.2.2.2/24

Network > Interfaces > New Tunnel IF: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.1

Zone (VR): Untrust (trust-vr)

Unnumbered: (select)

Interface: ethernet3 (trust-vr)

2. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Trust_LAN

IP Address/Domain Name:

IP/Netmask: (select), 10.2.2.0/24

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Tokyo_Office

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.0/24

Zone: Untrust

3. VPN

VPNs > AutoKey Advanced > Gateway > New: Enter the following, and then click **OK**:

Gateway Name: To_Tokyo

Security Level: Custom

Remote Gateway Type:

Dynamic IP Address: (select), Peer ID: pmason@abc.com

Preshared Key

Preshared Key: h1p8A24nG5

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Gateway configuration page:

Security Level: Custom

Phase 1 Proposal (For Custom Security Level):
pre-g2-3des-sha

Mode (Initiator): Aggressive

(or)

Certificates

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Gateway configuration page:

Security Level: Custom

Phase 1 Proposal (For Custom Security Level): rsa-g2-3des-sha

Mode (Initiator): Aggressive

Preferred Certificate (optional):

Peer CA: Verisign

Peer Type: X509-SIG

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: Paris_Tokyo

Security Level: Compatible

Remote Gateway:

Predefined: (select), To_Tokyo

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Bind to: Tunnel Interface: (select), tunnel.1

Proxy-ID: (select)

Local IP / Netmask: 10.2.2.0/24

Remote IP / Netmask: 10.1.1.0/24

Service: ANY

4. Routes

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: (select), 2.2.2.250

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 10.1.1.0/24

Gateway: (select)

Interface: Tunnel.1

Gateway IP Address: 0.0.0.0

5. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Trust_LAN

Destination Address:

Address Book Entry: (select), Tokyo_Office

Service: Any

Action: Permit

Position at Top: (select)

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Tokyo_Office

Destination Address:

Address Book Entry: (select), Trust_LAN

Service: Any

Action: Permit

Position at Top: (select)

CLI (Tokyo)

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 dhcp client
set interface ethernet3 dhcp client settings server 1.1.1.5

set interface tunnel.1 zone untrust
set interface tunnel.1 ip unnumbered interface ethernet3
```

2. Addresses

```
set address trust Trust_LAN 10.1.1.0/24
set address untrust Paris_Office 10.2.2.0/24
```

3. VPN

Preshared Key

```
set ike gateway To_Paris address 2.2.2.2 aggressive local-id pmason@abc.com
    outgoing-interface ethernet3 preshare h1p8A24nG5 proposal pre-g2-3des-sha
set vpn Tokyo_Paris gateway To_Paris tunnel proposal nopfs-esp-3des-sha
set vpn Tokyo_Paris bind interface tunnel.1
set vpn Tokyo_Paris proxy-id local-ip 10.1.1.0/24 remote-ip 10.2.2.0/24 any
```

(or)

Certificates

```
set ike gateway To_Paris address 2.2.2.2 aggressive local-id pmason@abc.com[8]
    outgoing-interface ethernet3 proposal rsa-g2-3des-sha
set ike gateway To_Paris cert peer-ca 1[9]
set ike gateway To_Paris cert peer-cert-type x509-sig
set vpn Tokyo_Paris gateway To_Paris tunnel proposal nopfs-esp-3des-sha
set vpn Tokyo_Paris bind interface tunnel.1
set vpn Tokyo_Paris proxy-id local-ip 10.1.1.0/24 remote-ip 10.2.2.0/24 any
```

4. Routes

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3[10]
set vrouter trust-vr route 10.2.2.0/24 interface tunnel.1
```

5. Policies

```
set policy top from trust to untrust Trust_LAN Paris_Office any permit
set policy top from untrust to trust Paris_Office Trust_LAN any permit
save
```

8. The U-FQDN “pmason@abc.com” must appear in the SubjectAltName field in the certificate.
9. The number 1 is the CA ID number. To discover the CA’s ID number, use the following command: **get ike ca**.
10. The ISP provides the gateway IP address dynamically through DHCP, so you cannot specify it here.

CLI (Paris)

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.2.2.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 2.2.2.2/24

set interface tunnel.1 zone untrust
set interface tunnel.1 ip unnumbered interface ethernet3
```

2. Addresses

```
set address trust Trust_LAN 10.2.2.0/24
set address untrust Tokyo_Office 10.1.1.0/24
```

3. VPN

Preshared Key

```
set ike gateway To_Tokyo dynamic pmason@abc.com aggressive outgoing-interface
    ethernet3 preshare h1p8A24nG5 proposal pre-g2-3des-sha
set vpn Paris_Tokyo gateway To_Tokyo tunnel proposal nopfs-esp-3des-sha
set vpn Paris_Tokyo bind interface tunnel.1
set vpn Paris_Tokyo proxy-id local-ip 10.2.2.0/24 remote-ip 10.1.1.0/24 any
```

(or)

Certificates

```
set ike gateway To_Tokyo dynamic pmason@abc.com aggressive outgoing-interface
    ethernet3 proposal rsa-g2-3des-sha
set ike gateway To_Tokyo cert peer-ca 1[11]
set ike gateway To_Tokyo cert peer-cert-type x509-sig
set vpn Paris_Tokyo gateway To_Tokyo tunnel proposal nopfs-esp-3des-sha
set vpn Paris_Tokyo bind interface tunnel.1
set vpn Paris_Tokyo proxy-id local-ip 10.2.2.0/24 remote-ip 10.1.1.0/24 any
```

4. Routes

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 2.2.2.250
set vrouter trust-vr route 10.1.1.0/24 interface tunnel.1
```

5. Policies

```
set policy top from trust to untrust Trust_LAN Tokyo_Office any permit
set policy top from untrust to trust Tokyo_Office Trust_LAN any permit
save
```

11. The number 1 is the CA ID number. To discover the CA's ID number, use the following command: **get ike ca**.

Example: Policy-Based Site-to-Site VPN, Dynamic Peer

In this example, a VPN tunnel securely connects the users in the Trust zone behind NetScreen-A to the mail server in the corporate DMZ zone, protected by NetScreen-B. The Untrust zone interface for NetScreen-B has a static IP address. The ISP serving NetScreen-A assigns the IP address for its Untrust zone interface dynamically via DHCP. Because only NetScreen-B has a fixed address for its Untrust zone, VPN traffic must originate from hosts behind NetScreen-A. After NetScreen-A has established the tunnel, traffic through the tunnel can originate from either end. All zones are in the trust-vr routing domain.

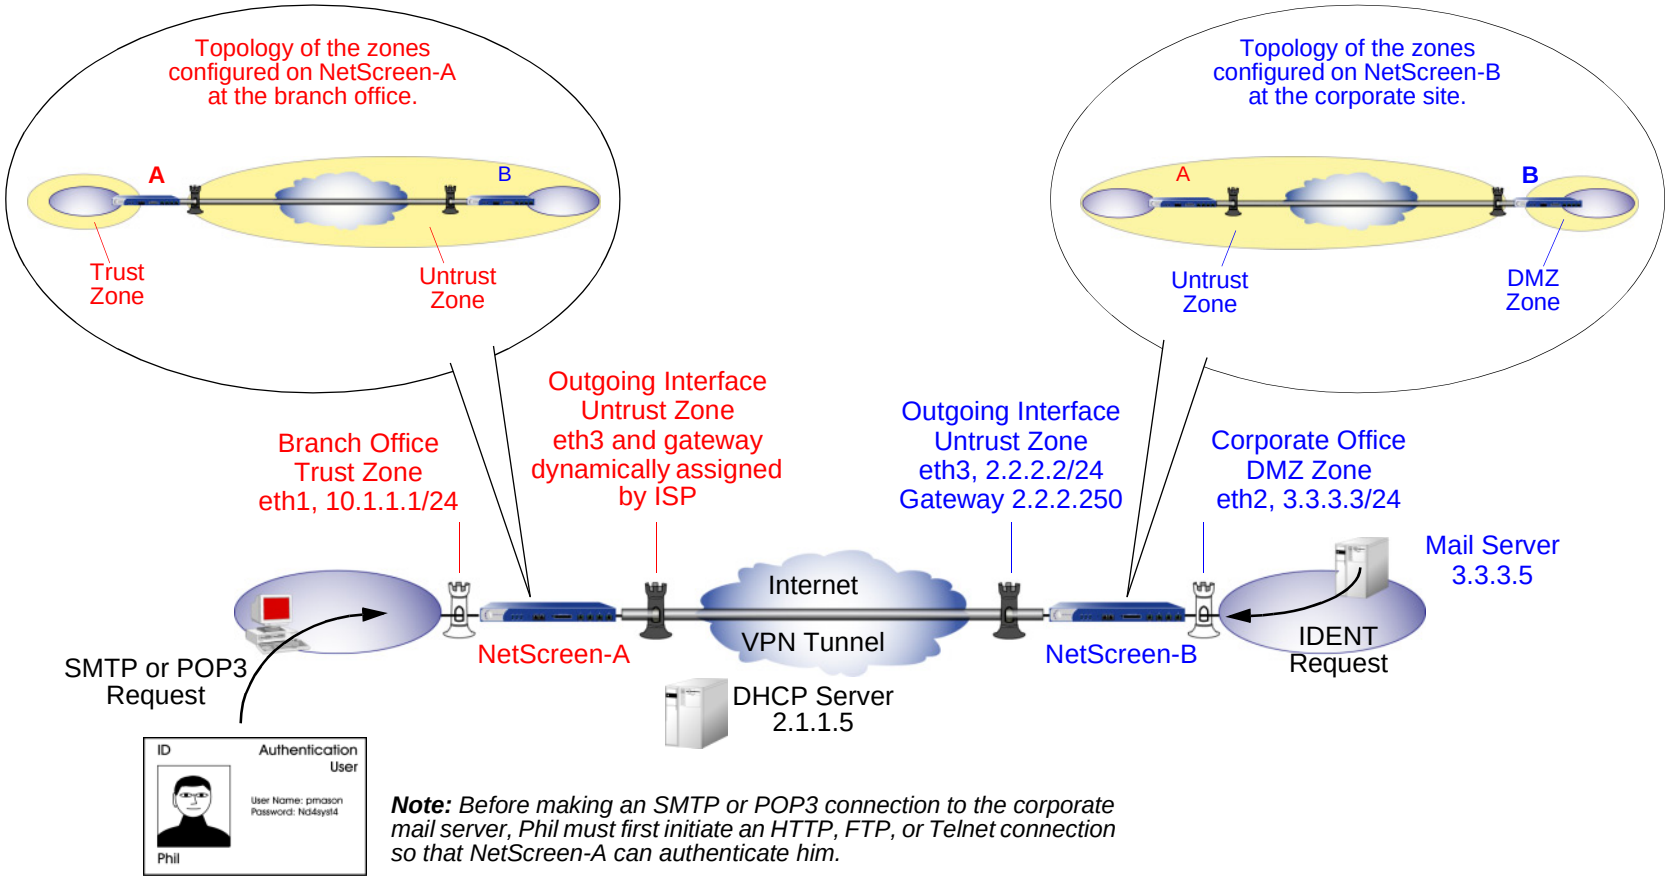

Note: *Before making an SMTP or POP3 connection to the corporate mail server, Phil must first initiate an HTTP, FTP, or Telnet connection so that NetScreen-A can authenticate him.*

In this example, the local auth user Phil (login name: pmason; password: Nd4syst4) wants to get his e-mail from the mail server at the corporate site. When he attempts to do so, he is authenticated twice: first, NetScreen-A authenticates him locally before allowing traffic from him through the tunnel[12]; second, the mail server program authenticates him, sending the IDENT request through the tunnel.

Note: *The mail server can send the IDENT request through the tunnel only if the NetScreen-A and B administrators add a custom service for it (TCP, port 113) and set up policies allowing that traffic through the tunnel to the 10.10.10.0/24 subnet.*

The preshared key is h1p8A24nG5. It is assumed that both participants already have RSA certificates from the certificate authority (CA) Verisign, and that the e-mail address *pmason@abc.com* appears in the local certificate on NetScreen-A. (For information about obtaining and loading certificates, see "Certificates and CRLs" on page 21.) For the Phase 1 and 2 security levels, you specify one Phase 1 proposal—either pre-g2-3des-sha for the preshared key method or rsa-g2-3des-sha for certificates—and select the predefined "Compatible" set of proposals for Phase 2.

WebUI (NetScreen-A)

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

12. Because Phil is an authentication user, before he can make an SMTP of POP3 request, he must first initiate an HTTP, FTP, or Telnet connection so that NetScreen-A can respond with a firewall user/login prompt to authenticate him. After NetScreen-A authenticates him, he has permission to contact the corporate mail server via the VPN tunnel.

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Obtain IP using DHCP: (select)[13]

2. User

Objects > Users > Local > New: Enter the following, and then click **OK**:

User Name: pmason

Status: Enable

Authentication User: (select)

User Password: Nd4syst4

Confirm Password: Nd4syst4

3. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Trusted network

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.0/24

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Mail Server

IP Address/Domain Name:

IP/Netmask: (select), 3.3.3.5/32

Zone: Untrust

13. You cannot specify the IP address of the DHCP server through the WebUI; however, you can do so through the CLI.

4. Services

Objects > Services > Custom > New: Enter the following, and then click **OK**:

Service Name: Ident

Service Timeout:

Use protocol default: (select)

Transport Protocol: TCP (select)

Source Port: Low 0, High 65535

Destination Port: Low 113, High 113

Objects > Services > Group > New: Enter the following, move the following services, and then click **OK**:

Group Name: Remote_Mail

Group Members << Available Members:

HTTP
FTP
Telnet
Ident
MAIL
POP3

5. VPN

VPNs > AutoKey Advanced > Gateway > New: Enter the following, and then click **OK**:

Gateway Name: To_Mail

Security Level: Custom

Remote Gateway Type:

Static IP Address: (select), IP Address/Hostname: 2.2.2.2

Preshared Key

Preshared Key: h1p8A24nG5

Local ID: pmason@abc.com

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Gateway configuration page:

Security Level: Custom

Phase 1 Proposal (For Custom Security Level): pre-g2-3des-sha

Mode (Initiator): Aggressive

(or)

Certificates

Local ID: pmason@abc.com

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Gateway configuration page:

Security Level: Custom

Phase 1 Proposal (For Custom Security Level): rsa-g2-3des-sha

Mode (Initiator): Aggressive

Preferred Certificate (optional):

Peer CA: Verisign

Peer Type: X509-SIG

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

Name: branch_corp

Security Level: Compatible

Remote Gateway Tunnel: To_Mail

6. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 0.0.0.0[14]

7. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Trusted network

Destination Address:

Address Book Entry: (select), Mail Server

Service: Remote_Mail

Action: Tunnel

VPN Tunnel: branch_corp

Modify matching bidirectional VPN policy: (select)

Position at Top: (select)

14. The ISP provides the gateway IP address dynamically through DHCP.

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Policy configuration page:

Authentication: (select)

Auth Server: Local

User: (select), Local Auth User - pmason

WebUI (NetScreen-B)

1. Interfaces

Network > Interfaces > Edit (for ethernet2): Enter the following, and then click **OK**:

Zone Name: DMZ

Static IP: (select this option when present)

IP Address/Netmask: 3.3.3.3/24

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 2.2.2.2/24

2. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Mail Server

IP Address/Domain Name:

IP/Netmask: (select), 3.3.3.5/32

Zone: DMZ

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: branch office

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.0/24

Zone: Untrust

3. Services

Objects > Services > Custom > New: Enter the following, and then click **OK**:

Service Name: Ident

Service Timeout:

Use protocol default: (select)

Transport Protocol: TCP (select)

Source Port: Low 0, High 65535

Destination Port: Low 113, High 113

Objects > Services > Group > New: Enter the following, move the following services, and then click **OK**:

Group Name: Remote_Mail

Group Members << Available Members:

Ident

MAIL

POP3

4. VPN

VPNs > AutoKey Advanced > Gateway > New: Enter the following, and then click **OK**:

Gateway Name: To_branch

Security Level: Custom

Remote Gateway Type:

Dynamic IP Address: (select), Peer ID: pmason@abc.com

Preshared Key

Preshared Key: h1p8A24nG5

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Gateway configuration page:

Security Level: Custom

Phase 1 Proposal (For Custom Security Level): pre-g2-3des-sha

Mode (Initiator): Aggressive

(or)

Certificates

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Gateway configuration page:

Security Level: Custom

Phase 1 Proposal (For Custom Security Level): rsa-g2-3des-sha

Mode (Initiator): Aggressive

Preferred Certificate (optional):

Peer CA: Verisign

Peer Type: X509-SIG

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: corp_branch

Security Level: Compatible

Remote Gateway:

Predefined: (select), To_branch

5. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 2.2.2.250

6. Policies

Policies > (From: DMZ, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Mail Server

Destination Address:

Address Book Entry: (select), branch office

Service: Remote_Mail

Action: Tunnel

VPN Tunnel: corp_branch

Modify matching bidirectional VPN policy: (select)

Position at Top: (select)

CLI (NetScreen-A)

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 dhcp client
set interface ethernet3 dhcp client settings server 1.1.1.5
```

2. User

```
set user pmason password Nd4syst4
```

3. Addresses

```
set address trust “trusted network” 10.1.1.0/24
set address untrust “mail server” 3.3.3.5/32
```

4. Services

```
set service ident protocol tcp src-port 0-65535 dst-port 113-113
set group service remote_mail
set group service remote_mail add http
set group service remote_mail add ftp
set group service remote_mail add telnet
set group service remote_mail add ident
set group service remote_mail add mail
set group service remote_mail add pop3
```

5. VPN

Preshared Key

```
set ike gateway to_mail address 2.2.2.2 aggressive local-id pmason@abc.com
    outgoing-interface ethernet3 preshare h1p8A24nG5 proposal pre-g2-3des-sha
set vpn branch_corp gateway to_mail sec-level compatible
```

(or)

Certificates

```
set ike gateway to_mail address 2.2.2.2 aggressive local-id pmason@abc.com[15]
    outgoing-interface ethernet3 proposal rsa-g2-3des-sha
set ike gateway to_mail cert peer-ca 1[16]
set ike gateway to_mail cert peer-cert-type x509-sig
set vpn branch_corp gateway to_mail sec-level compatible
```

6. Route

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3[17]
```

7. Policies

```
set policy top from trust to untrust “trusted network” “mail server”
    remote_mail tunnel vpn branch_corp auth server Local user pmason
set policy top from untrust to trust “mail server” “trusted network”
    remote_mail tunnel vpn branch_corp
save
```

15. The U-FQDN “pmason@abc.com” must appear in the SubjectAltName field in the certificate.
16. The number 1 is the CA ID number. To discover the CA’s ID number, use the following command: **get ike ca**.
17. The ISP provides the gateway IP address dynamically through DHCP.

CLI (NetScreen-B)

1. Interfaces

```
set interface ethernet2 zone dmz
set interface ethernet2 ip 3.3.3.3/24

set interface ethernet3 zone untrust
set interface ethernet3 ip 2.2.2.2/24
```

2. Addresses

```
set address dmz “mail server” 3.3.3.5/32
set address untrust “branch office” 10.1.1.0/24
```

3. Services

```
set service ident protocol tcp src-port 0-65535 dst-port 113-113
set group service remote_mail
set group service remote_mail add ident
set group service remote_mail add mail
set group service remote_mail add pop3
```

4. VPN

Preshared Key

```
set ike gateway to_branch dynamic pmason@abc.com aggressive outgoing-interface
    ethernet3 preshare h1p8A24nG5 proposal pre-g2-3des-sha
set vpn corp_branch gateway to_branch tunnel sec-level compatible
```

(or)

Certificates

```
set ike gateway to_branch dynamic pmason@abc.com aggressive outgoing-interface
    ethernet3 proposal rsa-g2-3des-sha
set ike gateway to_branch cert peer-ca 1[18]
set ike gateway to_branch cert peer-cert-type x509-sig
set vpn corp_branch gateway to_branch sec-level compatible
```

5. Route

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 2.2.2.250
```

6. Policies

```
set policy top from dmz to untrust “mail server” “branch office” remote_mail
    tunnel vpn corp_branch
set policy top from untrust to dmz “branch office” “mail server” remote_mail
    tunnel vpn corp_branch
save
```

18. The number 1 is the CA ID number. To discover the CA’s ID number, use the following command: **get ike ca**.

Example: Route-Based Site-to-Site VPN, Manual Key

In this example, a Manual Key tunnel provides a secure communication channel between offices in Tokyo and Paris. The Trust zones at each site are in NAT mode. The addresses are as follows:

- Tokyo:
  - Trust zone interface (ethernet1): 10.1.1.1/24
  - Untrust zone interface (ethernet3): 1.1.1.1/24
- Paris:
  - Trust zone interface (ethernet1): 10.2.2.1/24
  - Untrust zone interface (ethernet3): 2.2.2.2/24

The Trust and Untrust security zones and the Untrust-Tun tunnel zone are all in the trust-vr routing domain. The Untrust zone interface (ethernet3) serves as the outgoing interface for the VPN tunnel.

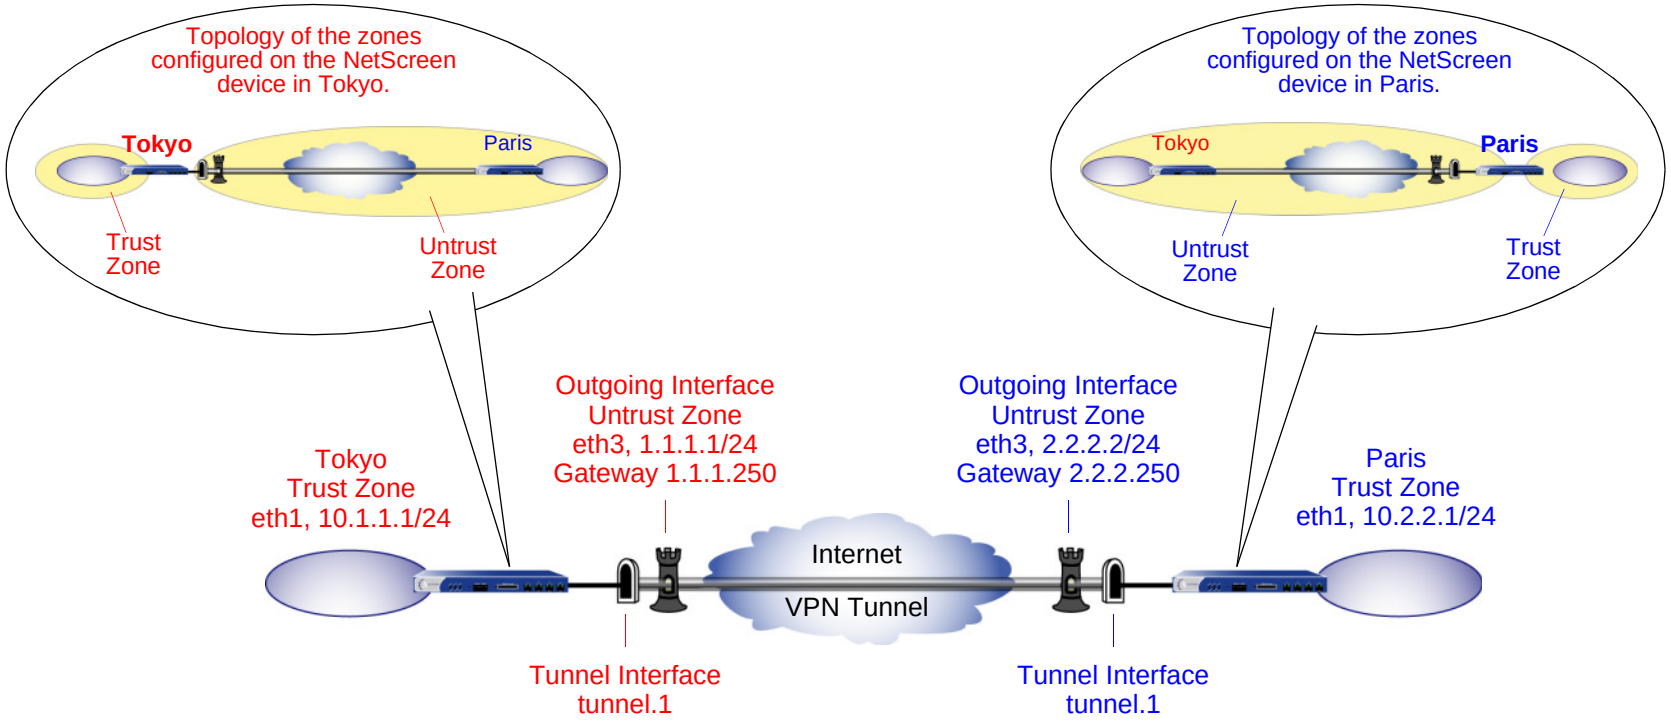

To set up the tunnel, perform the following steps on the NetScreen devices at both ends of the tunnel:

1. Assign IP addresses to the physical interfaces bound to the security zones and to the tunnel interface.
2. Configure the VPN tunnel, designate its outgoing interface in the Untrust zone, and bind it to the tunnel interface.
3. Enter the IP addresses for the local and remote endpoints in the address books for the Trust and Untrust zones.
4. Enter a default route to the external router in the trust-vr, and a route to the destination via the tunnel interface.
5. Set up policies for VPN traffic to pass between each site.

WebUI (Tokyo)

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

Network > Interfaces > New Tunnel IF: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.1

Zone (VR): Untrust (trust-vr)

Unnumbered: (select)

Interface: ethernet3 (trust-vr)

2. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Trust_LAN

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.0/24

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Paris_Office

IP Address/Domain Name:

IP/Netmask: (select), 10.2.2.0/24

Zone: Untrust

3. VPN

VPNs > Manual Key > New: Enter the following, and then click **OK**:

VPN Tunnel Name: Tokyo_Paris

Gateway IP: 2.2.2.2

Security Index: 3020 (Local), 3030 (Remote)

Outgoing Interface: ethernet3

ESP-CBC: (select)

Encryption Algorithm: 3DES-CBC

Generate Key by Password: asdlk24234

Authentication Algorithm: SHA-1

Generate Key by Password: PNas134a

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Manual Key tunnel configuration page:

Bind to: Tunnel Interface, tunnel.1

4. Routes

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 1.1.1.250

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 10.2.2.0/24

Gateway: (select)

Interface: tunnel.1

Gateway IP Address: 0.0.0.0

5. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Name: To Paris

Source Address:

Address Book Entry: (select), Trust_LAN

Destination Address:

Address Book Entry: (select), Paris_Office

Service: ANY

Action: Permit

Position at Top: (select)

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Name: From Paris

Source Address:

Address Book Entry: (select), Paris_Office

Destination Address:

Address Book Entry: (select), Trust_LAN

Service: ANY

Action: Permit

Position at Top: (select)

WebUI (Paris)

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.2.2.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 2.2.2.2/24

Network > Interfaces > New Tunnel IF: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.1

Zone (VR): Untrust (trust-vr)

Unnumbered: (select)

Interface: ethernet3 (trust-vr)

2. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Trust_LAN

IP Address/Domain Name:

IP/Netmask: (select), 10.2.2.0/24

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Tokyo_Office

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.0/24

Zone: Untrust

3. VPN

VPNs > Manual Key > New: Enter the following, and then click **OK**:

VPN Tunnel Name: Paris_Tokyo

Gateway IP: 1.1.1.1

Security Index: 3030 (Local), 3020 (Remote)

Outgoing Interface: ethernet3

ESP-CBC: (select)

Encryption Algorithm: 3DES-CBC

Generate Key by Password: asdlk24234

Authentication Algorithm: SHA-1

Generate Key by Password: PNas134a

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Manual Key tunnel configuration page:

Bind to: Tunnel Interface, tunnel.1

4. Routes

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 2.2.2.250

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 10.1.1.0/24

Gateway: (select)

Interface: tunnel.1

Gateway IP Address: 0.0.0.0

5. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Name: To Tokyo

Source Address:

Address Book Entry: (select), Trust_LAN

Destination Address:

Address Book Entry: (select), Tokyo_Office

Service: ANY

Action: Permit

Position at Top: (select)

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Name: From Tokyo

Source Address:

Address Book Entry: (select), Tokyo_Office

Destination Address:

Address Book Entry: (select), Trust_LAN

Service: ANY

Action: Permit

Position at Top: (select)

CLI (Tokyo)

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24

set interface tunnel.1 zone untrust
set interface tunnel.1 ip unnumbered interface ethernet3
```

2. Addresses

```
set address trust Trust_LAN 10.1.1.0/24
set address untrust Paris_Office 10.2.2.0/24
```

3. VPN

```
set vpn Tokyo_Paris manual 3020 3030 gateway 2.2.2.2 outgoing-interface
    ethernet3 esp 3des password asdlk24234 auth sha-1 password PNas134a
set vpn Tokyo_Paris bind interface tunnel.1
```

4. Routes

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
set vrouter trust-vr route 10.2.2.0/24 interface tunnel.1
```

5. Policies

```
set policy top name “To Paris” from trust to untrust Trust_LAN Paris_Office any
    permit
set policy top name “From Paris” from untrust to trust Paris_Office Trust_LAN
    any permit
save
```

CLI (Paris)

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.2.2.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 2.2.2.2/24

set interface tunnel.1 zone untrust
set interface tunnel.1 ip unnumbered interface ethernet3
```

2. Addresses

```
set address trust Trust_LAN 10.2.2.0/24
set address untrust Tokyo_Office 10.1.1.0/24
```

3. VPN

```
set vpn Paris_Tokyo manual 3030 3020 gateway 1.1.1.1 outgoing-interface
    ethernet3 esp 3des password asdlk24234 auth sha-1 password PNas134a
set vpn Paris_Tokyo bind interface tunnel.1
```

4. Routes

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 2.2.2.250
set vrouter trust-vr route 10.1.1.0/24 interface tunnel.1
```

5. Policies

```
set policy top name “To Tokyo” from trust to untrust Trust_LAN Tokyo_Office any
    permit
set policy top name “From Tokyo” from untrust to trust Tokyo_Office Trust_LAN
    any permit
save
```

Example: Policy-Based Site-to-Site VPN, Manual Key

In this example, a Manual Key tunnel provides a secure communication channel between offices in Tokyo and Paris, using ESP with 3DES encryption and SHA-1 authentication. The Trust zones at each site are in NAT mode. The addresses are as follows:

- Tokyo:
  - Trust interface (ethernet1): 10.1.1.1/24
  - Untrust interface (ethernet3): 1.1.1.1/24
- Paris:
  - Trust interface (ethernet1): 10.2.2.1/24
  - Untrust interface (ethernet3): 2.2.2.2/24

The Trust and Untrust security zones and the Untrust-Tun tunnel zone are in the trust-vr routing domain. The Untrust zone interface (ethernet3) serves as the outgoing interface for the VPN tunnel.

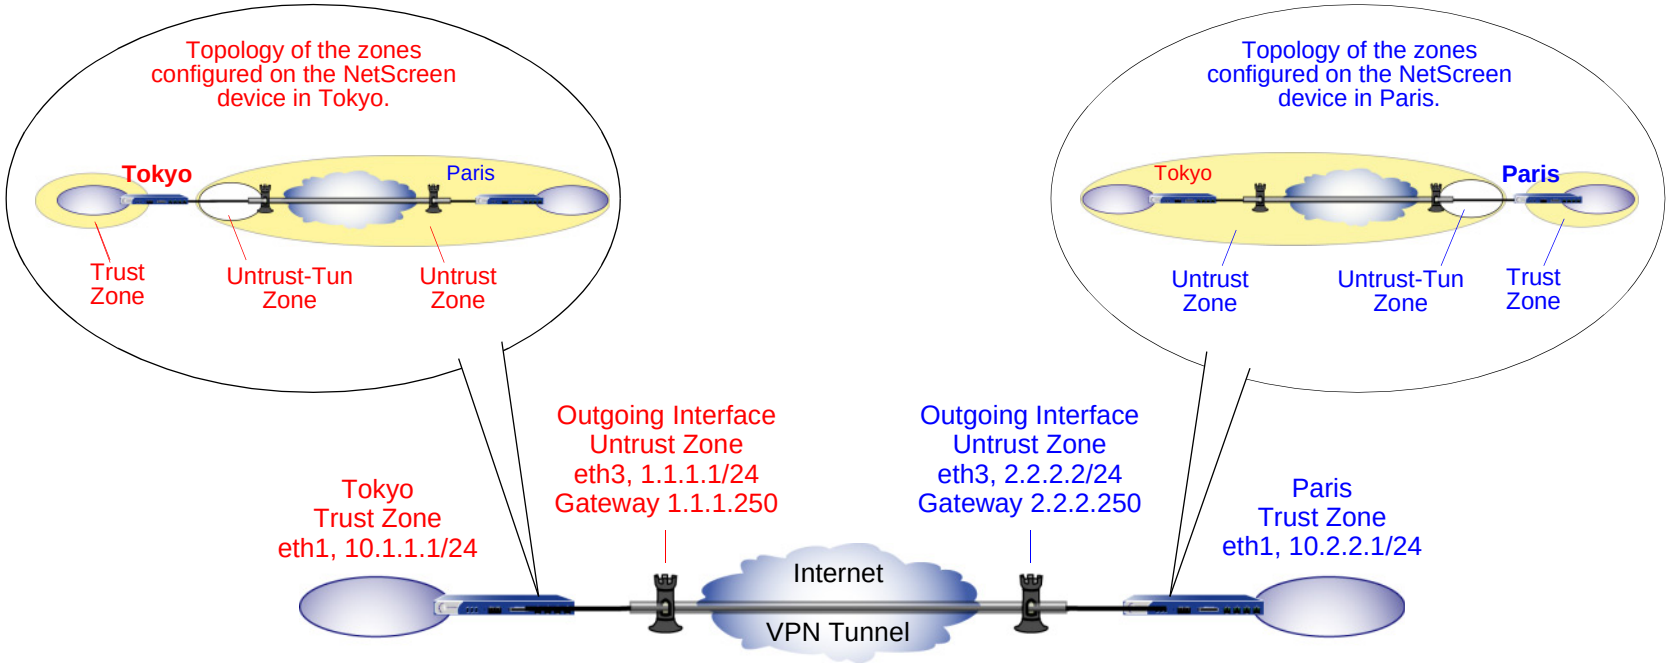

To set up the tunnel, perform the following five steps on the NetScreen devices at both ends of the tunnel:

1. Assign IP addresses to the physical interfaces bound to the security zones.
2. Configure the VPN tunnel, and designate its outgoing interface in the Untrust zone.
3. Enter the IP addresses for the local and remote endpoints in the Trust and Untrust address books.
4. Enter a default route to the external router.
5. Set up policies for VPN traffic to pass bidirectionally through the tunnel.

WebUI (Tokyo)

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

2. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Trust_LAN

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.0/24

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Paris_Office

IP Address/Domain Name:

IP/Netmask: (select), 10.2.2.0/24

Zone: Untrust

3. VPN

VPNs > Manual Key > New: Enter the following, and then click **OK**:

VPN Tunnel Name: Tokyo_Paris

Gateway IP: 2.2.2.2

Security Index: 3020 (Local), 3030 (Remote)

Outgoing Interface: ethernet3

ESP-CBC: (select)

Encryption Algorithm: 3DES-CBC

Generate Key by Password: asdlk24234

Authentication Algorithm: SHA-1

Generate Key by Password: PNas134a

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Manual Key tunnel configuration page:

Bind to: Tunnel Zone, Untrust-Tun

4. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 1.1.1.250

5. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Name: To/From Paris

Source Address:

Address Book Entry: (select), Trust_LAN

Destination Address:

Address Book Entry: (select), Paris_Office

Service: ANY

Action: Tunnel

Tunnel VPN: Tokyo_Paris

Modify matching bidirectional VPN policy: (select)

Position at Top: (select)

WebUI (Paris)

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.2.2.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 2.2.2.2/24

2. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Trust_LAN

IP Address/Domain Name:

IP/Netmask: (select), 10.2.2.0/24

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Tokyo_Office

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.0/24

Zone: Untrust

3. VPN

VPNs > Manual Key > New: Enter the following, and then click **OK**:

VPN Tunnel Name: Paris_Tokyo

Gateway IP: 1.1.1.1

Security Index (HEX Number): 3030 (Local), 3020 (Remote)

Outgoing Interface: ethernet3

ESP-CBC: (select)

Encryption Algorithm: 3DES-CBC

Generate Key by Password: asdlk24234

Authentication Algorithm: SHA-1

Generate Key by Password: PNas134a

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Manual Key tunnel configuration page:

Bind to: Tunnel Zone, Untrust-Tun

4. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 2.2.2.250

5. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Name: To/From Tokyo

Source Address:

Address Book Entry: (select), Trust_LAN

Destination Address:

Address Book Entry: (select), Tokyo_Office

Service: ANY

Action: Tunnel

Tunnel VPN: Paris_Tokyo

Modify matching bidirectional VPN policy: (select)

Position at Top: (select)

CLI (Tokyo)

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
```

2. Addresses

```
set address trust Trust_LAN 10.1.1.0/24
set address untrust paris_office 10.2.2.0/24
```

3. VPN

```
set vpn tokyo_paris manual 3020 3030 gateway 2.2.2.2 outgoing-interface
    ethernet3 esp 3des password asdlk24234 auth sha-1 password PNas134a
set vpn tokyo_paris bind zone untrust-tun
```

4. Route

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
```

5. Policies

```
set policy top name “To/From Paris” from trust to untrust Trust_LAN
    paris_office any tunnel vpn tokyo_paris
set policy top name “To/From Paris” from untrust to trust paris_office
    Trust_LAN any tunnel vpn tokyo_paris
save
```

CLI (Paris)

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.2.2.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 2.2.2.2/24
```

2. Addresses

```
set address trust Trust_LAN 10.2.2.0/24
set address untrust tokyo_office 10.1.1.0/24
```

3. VPN

```
set vpn paris_tokyo manual 3030 3020 gateway 1.1.1.1 outgoing-interface
    ethernet3 esp 3des password asdlk24234 auth sha-1 password PNas134a
set vpn paris_tokyo bind zone untrust-tun
```

4. Route

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 2.2.2.250
```

5. Policies

```
set policy top name “To/From Tokyo” from trust to untrust Trust_LAN
    tokyo_office any tunnel vpn paris_tokyo
set policy top name “To/From Tokyo” from untrust to trust tokyo_office
    Trust_LAN any tunnel vpn paris_tokyo
save
```

FQDN for Dynamic IKE Gateways

For an IKE peer that has a static fully qualified domain name (FQDN) but a dynamically assigned IP address, you can specify the FQDN in the local configuration for the remote gateway. For example, an Internet service provider (ISP) might assign IP addresses via DHCP to its customers. The ISP draws addresses from a pool of about 2000 addresses and assigns them when its customers come online. Although the IKE peer has a static FQDN, it has an unpredictably changing IP address. The IKE peer has three methods available for maintaining a Domain Name Service (DNS) mapping of its static FQDN to its dynamically assigned IP address (a process known as dynamic DNS).

- If the remote IKE peer is a NetScreen device, the admin can manually notify the DNS server to update its FQDN-to-IP address mapping each time the NetScreen device receives a new IP address from its ISP.
- If the remote IKE peer is another kind of VPN termination device that has dynamic DNS software running on it, that software can automatically notify the DNS server of its address changes so the server can update its FQDN-to-IP address mapping table.
- If the remote IKE peer is a NetScreen device or any other kind of VPN termination device, a host behind it can run an FQDN-to-IP address automatic update program that alerts the DNS server of address changes.

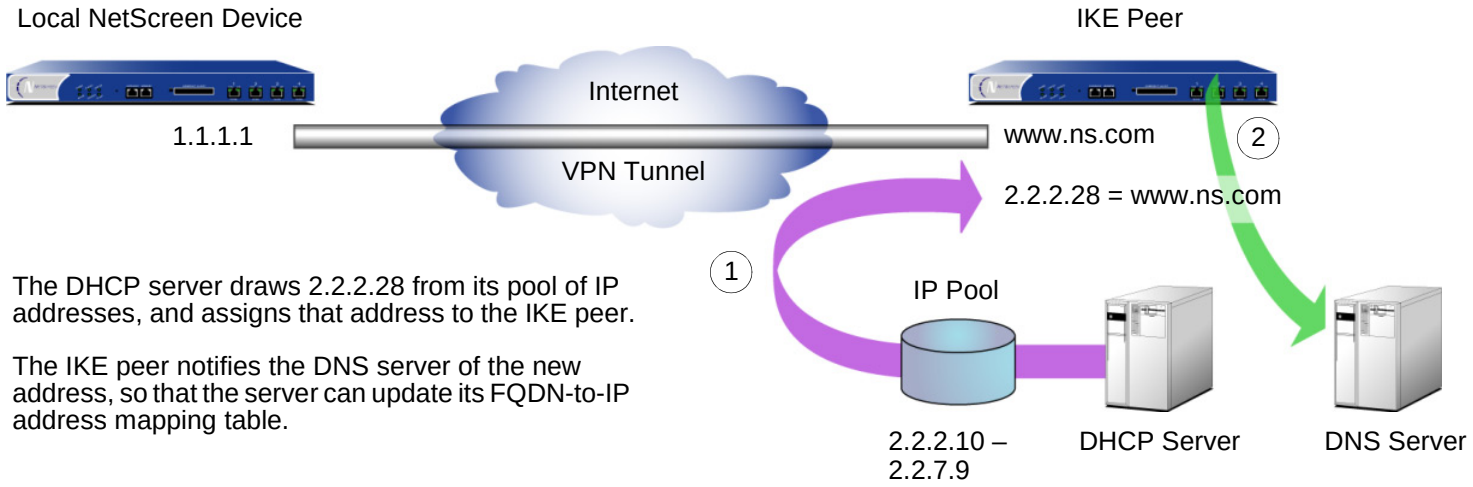

1. The DHCP server draws 2.2.2.28 from its pool of IP addresses, and assigns that address to the IKE peer.
2. The IKE peer notifies the DNS server of the new address, so that the server can update its FQDN-to-IP address mapping table.

Without needing to know the current IP address of a remote IKE peer, you can now configure an AutoKey IKE VPN tunnel to that peer using its FQDN instead of an IP address.

Aliases

You can also use an alias for the FQDN of the remote IKE peer if the DNS server that the local NetScreen device queries returns only one IP address. If the DNS server returns several IP addresses, the local device uses the first one it receives. Because there is no guarantee for the order of the addresses in the response from the DNS server, the local NetScreen device might use the wrong IP address and IKE negotiations might fail.

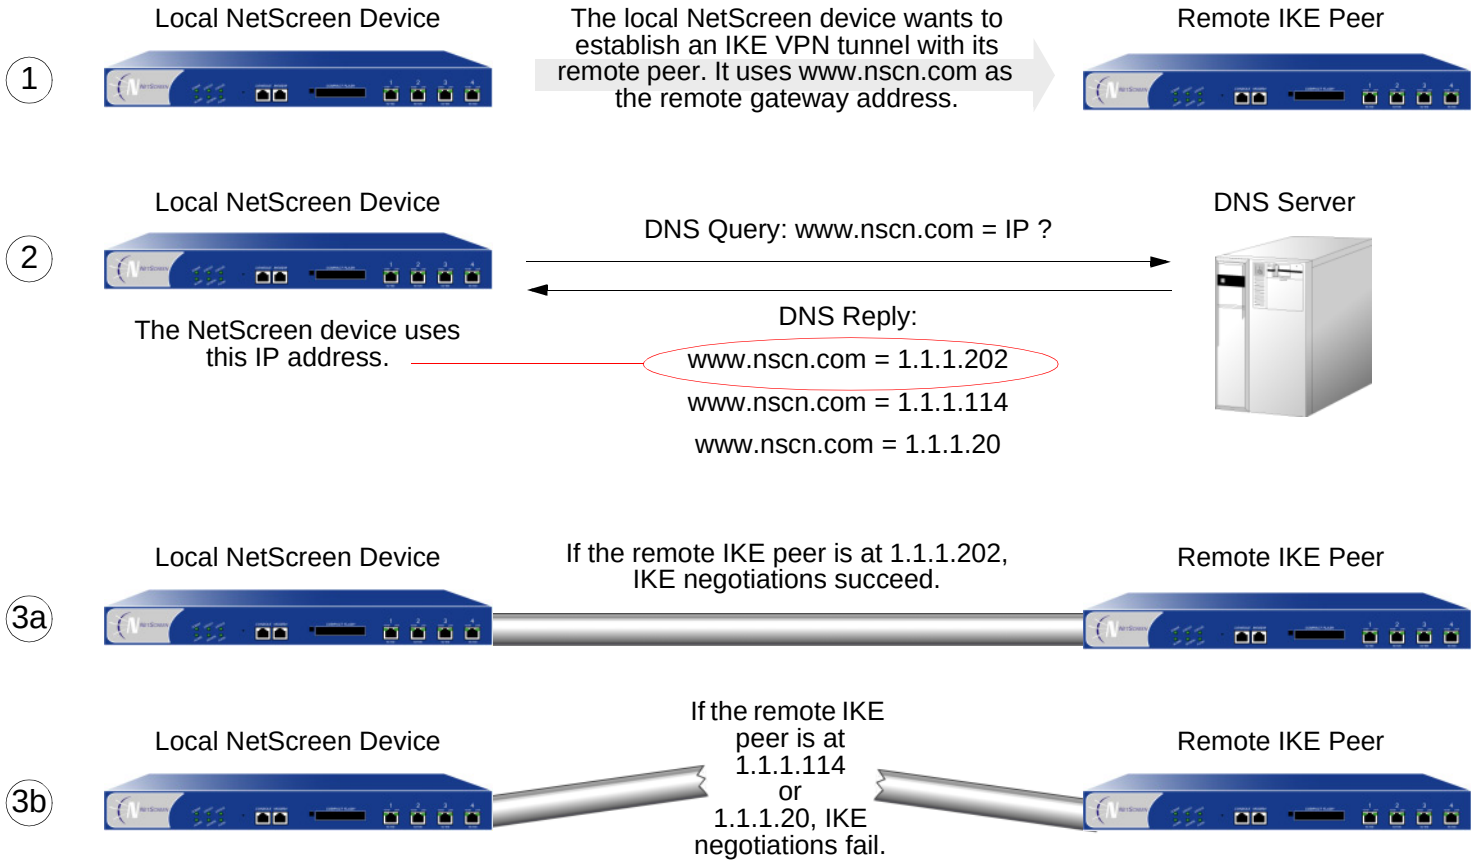

Example: AutoKey IKE Peer with FQDN

In this example, an AutoKey IKE VPN tunnel using either a preshared secret or a pair of certificates (one at each end of the tunnel) provides a secure connection between two offices in Tokyo and Paris. The Paris office has a dynamically assigned IP address, so the Tokyo office uses the remote peer's FQDN (www.nspar.com) as the address of the remote gateway in its VPN tunnel configuration.

The following configuration is for a route-based VPN tunnel. For the Phase 1 and 2 security levels, you specify one Phase 1 proposal—either pre-g2-3des-sha for the preshared key method or rsa-g2-3des-sha for certificates—and select the predefined "Compatible" set of proposals for Phase 2. All zones are in the trust-vr.

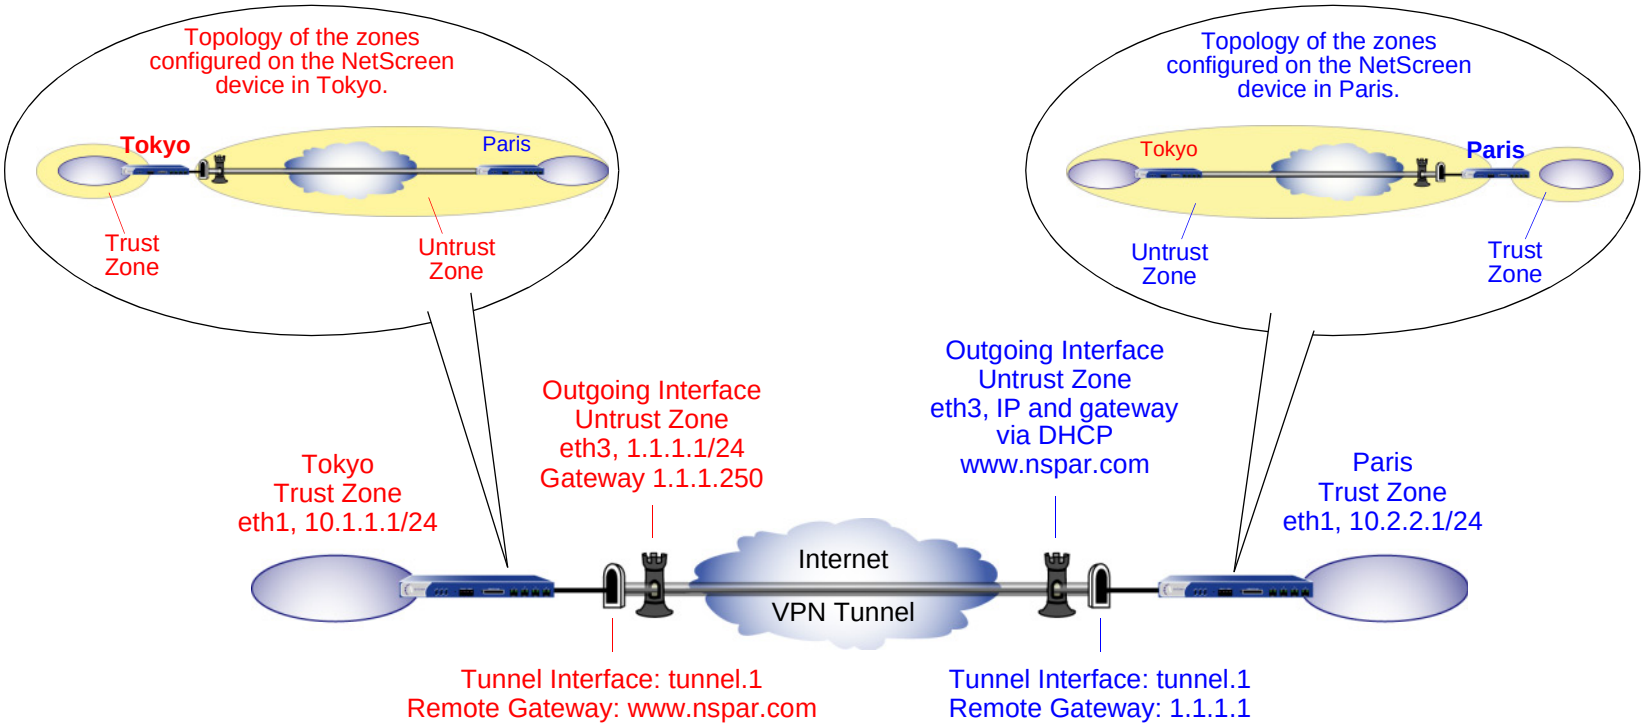

Setting up a route-based AutoKey IKE tunnel using either a preshared secret or certificates involves the following steps:

1. Assign IP addresses to the physical interfaces bound to the security zones and to the tunnel interface.
2. Define the remote gateway and key exchange mode, and specify either a preshared secret or a certificate
3. Configure the VPN tunnel, designate its outgoing interface in the Untrust zone, bind it to the tunnel interface, and configure its proxy-ID.
4. Enter the IP addresses for the local and remote endpoints in the Trust and Untrust address books.
5. Enter a default route to the external router in the trust-vr, and a route to the destination via the tunnel interface.
6. Set up policies for traffic to pass between each site.

In the following examples, the preshared key is h1p8A24nG5. It is assumed that both participants already have RSA certificates and are using Entrust as the certificate authority (CA). (For information about obtaining and loading certificates, see *NetScreen Concepts & Examples ScreenOS Reference Guide, Volume 4, VPNs*.)

WebUI (Tokyo)

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1

Network > Interfaces > New Tunnel IF: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.1

Zone (VR): Untrust (trust-vr)

Unnumbered: (select)

Interface: ethernet3 (trust-vr)

2. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Trust_LAN

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.0/24

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Paris_Office

IP Address/Domain Name:

IP/Netmask: (select), 10.2.2.0/24

Zone: Untrust

3. VPN

VPNs > AutoKey Advanced > Gateway > New: Enter the following, and then click **OK**:

Gateway Name: To_Paris

Security Level: Custom

Remote Gateway Type:

Static IP Address: (select), IP Address/Hostname: www.nspar.com

Preshared Key

Preshared Key: h1p8A24nG5

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Gateway configuration page:

Security Level: Custom

Phase 1 Proposal (for Custom Security Level): pre-g2-3des-sha

Mode (Initiator): Main (ID Protection)

(or)

Certificates

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Gateway configuration page:

Security Level: Custom

Phase 1 Proposal (for Custom Security Level): rsa-g2-3des-sha

Preferred certificate (optional)

Peer CA: Entrust

Peer Type: X509-SIG

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: Tokyo_Paris

Security Level: Compatible

Remote Gateway:

Predefined: (select), To_Paris

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Security Level: Compatible

Bind to: Tunnel Interface, tunnel.1

Proxy-ID: (select)

Local IP / Netmask: 10.1.1.0/24

Remote IP / Netmask: 10.2.2.0/24

Service: ANY

4. Routes

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 0.0.0.0[19]

19. The ISP provides the gateway IP address dynamically through DHCP.

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 10.2.2.0/24

Gateway: (select)

Interface: tunnel.1

Gateway IP Address: 0.0.0.0

5. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Name: To Paris

Source Address: Trust_LAN

Destination Address: Paris_Office

Service: ANY

Action: Permit

Position at Top: (select)

Policies > Policy (From: Untrust, To: Trust) > New Policy: Enter the following, and then click **OK**:

Name: From Paris

Source Address: Paris_Office

Destination Address: Trust_LAN

Service: ANY

Action: Permit

Position at Top: (select)

WebUI (Paris)

1. Host Name and Domain Name

Network > DNS: Enter the following, and then click **Apply**:

Host Name: www

Domain Name: nspar.com

2. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.2.2.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Obtain IP using DHCP: (select)

Network > Interfaces > New Tunnel IF: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.1

Zone (VR): Untrust (trust-vr)

Unnumbered: (select)

Interface: ethernet3 (trust-vr)

3. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Trust_LAN

IP Address/Domain Name:

IP/Netmask: (select), 10.2.2.0/24

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Tokyo_Office

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.0/24

Zone: Untrust

4. VPN

VPNs > AutoKey Advanced > Gateway > New: Enter the following, and then click **OK**:

Gateway Name: To_Tokyo

Security Level: Custom

Remote Gateway Type:

Static IP Address: (select), IP Address/Hostname: 1.1.1.1

Preshared Key

Preshared Key: h1p8A24nG5

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Gateway configuration page:

Security Level: Custom

Phase 1 Proposal (For Custom Security Level):
pre-g2-3des-sha

Mode (Initiator): Main (ID Protection)

(or)

Certificates

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Gateway configuration page:

Security Level: Custom

Phase 1 Proposal (for Custom Security Level): rsa-g2-3des-sha

Preferred certificate (optional)

Peer CA: Entrust

Peer Type: X509-SIG

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

Name: Paris_Tokyo

Security Level: Custom

Remote Gateway:

Predefined: (select), To_Tokyo

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Security Level: Compatible

Bind to: Tunnel Interface, tunnel.1

Proxy-ID: (select)
Local IP / Netmask: 10.2.2.0/24
Remote IP / Netmask: 10.1.1.0/24
Service: ANY

5. Routes

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0
Gateway: (select)
Interface: ethernet3
Gateway IP Address: 2.2.2.250

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 10.1.1.0/24
Gateway: (select)
Interface: tunnel.1
Gateway IP Address: 0.0.0.0

6. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Name: To Tokyo
Source Address: Trust_LAN
Destination Address: Tokyo_Office
Service: ANY
Action: Permit
Position at Top: (select)

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Name: From Tokyo

Source Address: Tokyo_Office

Destination Address: Trust_LAN

Service: ANY

Action: Permit

Position at Top: (select)

CLI (Tokyo)

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24

set interface tunnel.1 zone untrust
set interface tunnel.1 ip unnumbered interface ethernet3
```

2. Addresses

```
set address trust Trust_LAN 10.1.1.0/24
set address untrust paris_office 10.2.2.0/24
```

3. VPN

Preshared Key

```
set ike gateway to_paris address www.nspar.com main outgoing-interface
    ethernet3 preshare h1p8A24nG5 proposal pre-g2-3des-sha
set vpn tokyo_paris gateway to_paris sec-level compatible
set vpn tokyo_paris bind interface tunnel.1
set vpn tokyo_paris proxy-id local-ip 10.1.1.0/24 remote-ip 10.2.2.0/24 any
```

(or)

Certificate

```
set ike gateway to_paris address www.nspar.com main outgoing-interface
    ethernet3 proposal rsa-g2-3des-sha
set ike gateway to_paris cert peer-ca 1[20]
set ike gateway to_paris cert peer-cert-type x509-sig
set vpn tokyo_paris gateway to_paris sec-level compatible
set vpn tokyo_paris bind interface tunnel.1
set vpn tokyo_paris proxy-id local-ip 10.1.1.0/24 remote-ip 10.2.2.0/24 any
```

4. Routes

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
set vrouter trust-vr route 10.2.2.0/24 interface tunnel.1
```

5. Policies

```
set policy top name “To Paris” from trust to untrust Trust_LAN paris_office any
    permit
set policy top name “From Paris” from untrust to trust paris_office Trust_LAN
    any permit
save
```

20. The number 1 is the CA ID number. To discover the CA’s ID number, use the following command: **get ike ca**.

CLI (Paris)

1. Host Name and Domain Name

```
set hostname www
set domain nspar.com
```

2. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.2.2.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip dhcp-client enable

set interface tunnel.1 zone untrust
set interface tunnel.1 ip unnumbered interface ethernet3
```

3. Addresses

```
set address trust Trust_LAN 10.2.2.0/24
set address untrust tokyo_office 10.1.1.0/24
```

4. VPN

Preshared Key

```
set ike gateway to_tokyo address 1.1.1.1 main outgoing-interface ethernet3
    preshare h1p8A24nG5 proposal pre-g2-3des-sha
set vpn paris_tokyo gateway to_tokyo sec-level compatible
set vpn paris_tokyo bind interface tunnel.1
set vpn paris_tokyo proxy-id local-ip 10.2.2.0/24 remote-ip 10.1.1.0/24 any
```

(or)

Certificate

```
set ike gateway to_tokyo address 1.1.1.1 main outgoing-interface ethernet3
    proposal rsa-g2-3des-sha
set ike gateway to_tokyo cert peer-ca 1[3]
set ike gateway to_tokyo cert peer-cert-type x509-sig
set vpn paris_tokyo gateway to_tokyo sec-level compatible
set vpn paris_tokyo bind interface tunnel.1
set vpn paris_tokyo proxy-id local-ip 10.2.2.0/24 remote-ip 10.1.1.0/24 any
```

5. Routes

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 2.2.2.250
set vrouter trust-vr route 10.1.1.0/24 interface tunnel.1
```

6. Policies

```
set policy top name “To Tokyo” from trust to untrust Trust_LAN tokyo_office any
    permit
set policy top name “From Tokyo” from untrust to trust tokyo_office Trust_LAN
    any permit
save
```

VPN Sites with Overlapping Addresses

Because the range of private IP addresses is relatively small, there is a good chance that the addresses of protected networks of two VPN peers overlap[21]. For bidirectional VPN traffic between two end entities with overlapping addresses, the NetScreen devices at both ends of the tunnel must apply source and destination network address translation (NAT-src and NAT-dst) to the VPN traffic passing between them.

For NAT-src, the interfaces at both ends of the tunnel must have IP addresses in mutually unique subnets, with a dynamic IP (DIP) pool in each of those subnets [22]. The policies regulating outbound VPN traffic can then apply NAT-src using DIP pool addresses to translate original source addresses to those in a neutral address space.

To provide NAT-dst on inbound VPN traffic, there are two options:

- Policy-based NAT-dst: A policy can apply NAT-dst to translate inbound VPN traffic to an address that is either in the same subnet as the tunnel interface—but not in the same range as the local DIP pool used for outbound VPN traffic—or to an address in another subnet to which the NetScreen device has an entry in its route table. (For information about routing considerations when configuring NAT-dst, see “Routing for Destination Translation” on page **2**-282.)
- Mapped IP (MIP): A policy can reference a MIP as the destination address. The MIP uses an address in the same subnet as the tunnel interface—but not in the same range as the local DIP pool used for outbound VPN traffic. (For information about MIPs, see “Mapped IP Addresses” on page **2**-331.)

VPN traffic between sites with overlapping addresses requires address translation in both directions. Because the source address on outbound traffic cannot be the same as the destination address on inbound traffic—the NAT-dst address or MIP cannot be in the DIP pool—the addresses referenced in the inbound and outbound policies cannot be symmetrical.

21. An overlapping address space is when the IP address range in two networks are partially or completely the same.

22. The range of addresses in a DIP pool must be in the same subnet as the tunnel interface, but the pool must not include the interface IP address or any MIP or VIP addresses that might also be in that subnet. For security zone interfaces, you can also define an extended IP address and an accompanying DIP pool in a different subnet from that of the interface IP address. For more information, see “Extended Interface and DIP” on page **2**-175.

When you want the NetScreen device to perform source and destination address translation on bidirectional VPN traffic through the same tunnel, you have two choices:

- You can define a proxy ID[23] for a policy-based VPN configuration. When you specifically reference a VPN tunnel in a policy, the NetScreen device derives a proxy ID from the components in the policy that references that tunnel. The NetScreen device derives the proxy ID when you first create the policy, and each time the device reboots thereafter. However, if you manually define a proxy ID for a VPN tunnel that is referenced in a policy, the NetScreen device applies the user-defined proxy ID, not the proxy ID derived from the policy.
- You can use a route-based VPN tunnel configuration, which must have a user-defined proxy ID. With a route-based VPN tunnel configuration, you do not specifically reference a VPN tunnel in a policy. Instead, the policy controls access (permit or deny) to a particular destination. The route to that destination points to a tunnel interface that in turn is bound to a VPN tunnel. Because the VPN tunnel is not directly associated with a policy from which it can derive a proxy ID from the source address, destination address, and service, you must manually define a proxy ID for it. (Note that a route-based VPN configuration also allows you to create multiple policies that make use of a single VPN tunnel; that is, a single Phase 2 SA.)

Consider the addresses in following illustration of a VPN tunnel between two sites with overlapping address spaces:

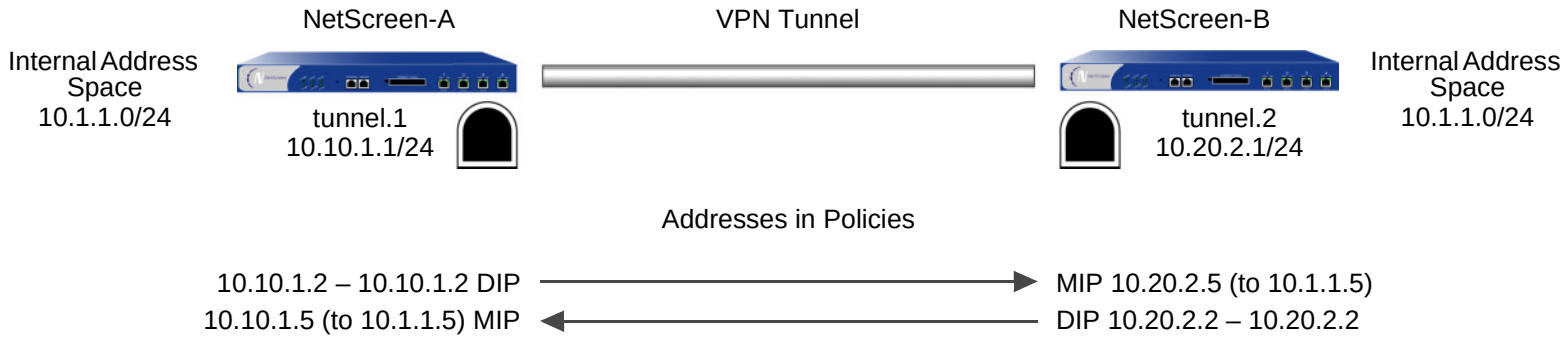

23. A proxy ID is a kind of agreement between IKE peers to permit traffic through a tunnel if the traffic matches a specified tuple of local address, remote address, and service.

If the NetScreen devices in the previous illustration derive proxy IDs from the policies, as they do in policy-based VPN configurations, then the inbound and outbound policies produce the following proxy IDs:

NetScreen-A			
	Local	Remote	Service
Outbound	10.10.1.2/32	10.20.2.5/32	Any
Inbound	10.10.1.5/32	10.20.2.2/32	Any

NetScreen-B			
	Local	Remote	Service
Inbound	10.20.2.5/32	10.10.1.2/32	Any
Outbound	10.20.2.2/32	10.10.1.5/32	Any

As you can see, there are two proxy IDs: one for outbound VPN traffic and another for inbound. When NetScreen-A first sends traffic from 10.10.1.2/32 to 10.20.2.5/32, the two peers perform IKE negotiations and produce Phase 1 and Phase 2 security associations (SAs). The Phase 2 SA results in the above outbound proxy ID for NetScreen-A, and the inbound proxy ID for NetScreen-B.

If NetScreen-B then sends traffic to NetScreen-A, the policy lookup for traffic from 10.20.2.2/32 to 10.10.1.5/32 indicates that there is no active Phase 2 SA for such a proxy ID. Therefore, the two peers use the existing Phase 1 SA (assuming that its lifetime has not yet expired) to negotiate a different Phase 2 SA. The resulting proxy IDs are shown above as the inbound proxy ID for NetScreen-A and the outbound proxy ID for NetScreen-B. There are two Phase 2 SAs—two VPN tunnels—because the addresses are asymmetrical and require different proxy IDs.

To create just one tunnel for bidirectional VPN traffic, you can define the following proxy IDs with addresses whose scope includes both the translated source and destination addresses at each end of the tunnel:

NetScreen-A				NetScreen-B		
Local	Remote	Service		Local	Remote	Service
10.10.1.0/24	10.20.2.0/24	Any		10.20.2.0/24	10.10.1.0/24	Any
			or			
0.0.0.0/0	0.0.0.0/0	Any		0.0.0.0/0	0.0.0.0/0	Any

The above proxy IDs encompass addresses appearing in both inbound and outbound VPN traffic between the two sites. The address 10.10.1.0/24 includes both the DIP pool 10.10.1.2 – 10.10.1.2 and the MIP 10.10.1.5. Likewise, the address 10.20.2.0/24 includes both the DIP pool 10.20.2.2 – 10.20.2.2 and the MIP 10.20.2.5[24]. The above

24. The address 0.0.0.0/0 includes all IP addresses, and thus the addresses of the DIP pool and MIP.

proxy IDs are symmetrical; that is, the local address for NetScreen-A is the remote address for NetScreen-B, and vice versa. If NetScreen-A sends traffic to NetScreen-B, the Phase 2 SA and proxy ID also apply to traffic sent from NetScreen-B to NetScreen-A. Thus, a single Phase 2 SA—that is, a single VPN tunnel—is all that is required for bidirectional traffic between the two sites.

To create one VPN tunnel for bidirectional traffic between sites with overlapping address spaces when the addresses for NAT-src and NAT-dst configured on the same device are in different subnets from each other, the proxy ID for the tunnel must be (local IP) 0.0.0.0/0 – (remote IP) 0.0.0.0/0 – *service type*. If you want to use more restrictive addresses in the proxy ID, then the addresses for NAT-src and NAT-dst must be in the same subnet.

Example: Tunnel Interface with NAT-Src and NAT-Dst

In this example, you configure a VPN tunnel between “NetScreen-A” at a corporate site and “NetScreen-B” at a branch office. The address space for the VPN end entities overlaps; they both use addresses in the 10.1.1.0/24 subnet. To overcome this conflict, you use NAT-src to translate the source address on outbound VPN traffic and NAT-dst to translate the destination address on inbound VPN traffic. The policies permit all addresses in the corporate LAN to reach an FTP server at the branch site, and for all addresses at the branch office site to reach an FTP server at the corporate site.

Note: *For more information about source and destination network address translation (NAT-src and NAT-dst), see Chapter 8, “Address Translation”.*

The tunnel configurations at both ends of the tunnel use the following parameters: AutoKey IKE, preshared key (“netscreen1”), and the security level predefined as “Compatible” for both Phase 1 and Phase 2 proposals. (For details about these proposals, see “Tunnel Negotiation” on page 11.)

The outgoing interface on NetScreen-A at the corporate site is ethernet3, which has IP address 1.1.1.1/24 and is bound to the Untrust zone. NetScreen-B at the branch office uses this address as its remote IKE gateway.

The outgoing interface on NetScreen-B at the branch office is ethernet3 , which has IP address 2.2.2.2/24 and is bound to the Untrust zone. NetScreen-A at the corporate site uses this address as its remote IKE gateway.

The Trust zone interface on both NetScreen devices is ethernet1 and has IP address 10.1.1.1/24. All zones on both NetScreen devices are in the trust-vr routing domain.

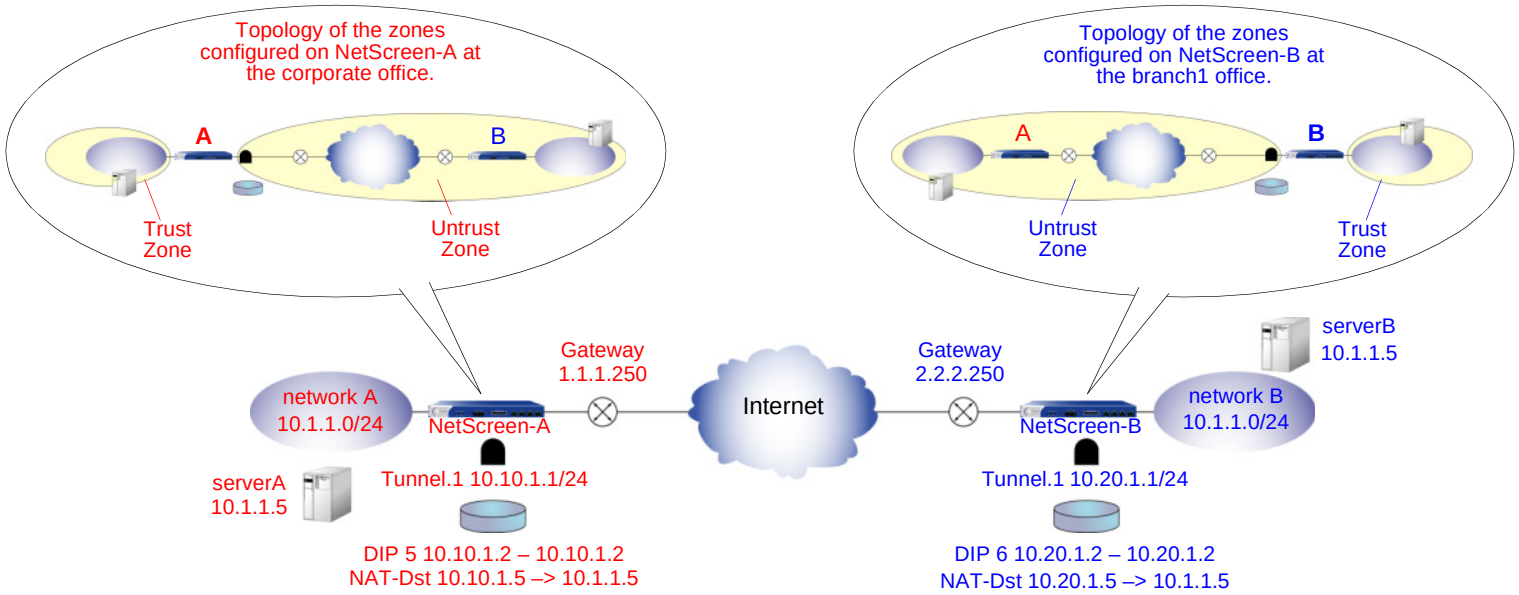

Users at network A can access server B. Users at network B can access server A.

All traffic flows through the VPN tunnel between the two sites.

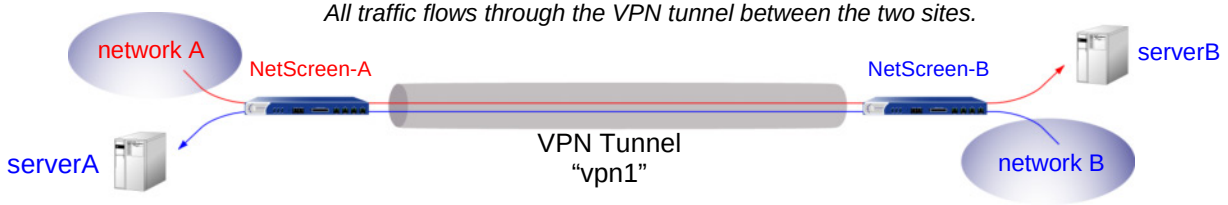

WebUI (NetScreen-A)

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

Network > Interfaces > New Tunnel IF: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.1

Zone (VR): Untrust (trust-vr)

Fixed IP: (select)

IP Address / Netmask: 10.10.1.1/24

2. DIP

Network > Interfaces > Edit (for tunnel.1) > DIP > New: Enter the following, and then click **OK**:

ID: 5

IP Address Range: (select), 10.10.1.2 ~ 10.10.1.2

Port Translation: (select)

In the same subnet as the interface IP or its secondary IPs: (select)

3. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: corp

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.0/24

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: virtualA

IP Address/Domain Name:

IP/Netmask: (select), 10.10.1.5/32

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: branch1

IP Address/Domain Name:

IP/Netmask: (select), 10.20.1.2/32

Zone: Untrust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: serverB

IP Address/Domain Name:

IP/Netmask: (select), 10.20.1.5/32

Zone: Untrust

4. VPN

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: vpn1

Security Level: Compatible

Remote Gateway: Create a Simple Gateway: (select)

Gateway Name: branch1

Type: Static IP: (select), Address/Hostname: 2.2.2.2

Preshared Key: netscreen1

Security Level: Compatible

Outgoing Interface: ethernet3[25]

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Bind to: Tunnel Interface, tunnel.1

Proxy-ID: (select)

Local IP / Netmask: 10.10.1.0/24

Remote IP / Netmask: 10.20.1.0/24

Service: ANY

25. The outgoing interface does not have to be in the same zone to which the tunnel interface is bound, although in this case they are in the same zone.

5. Routes

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 10.20.1.0/24

Gateway: (select)

Interface: tunnel.1

Gateway IP Address: 0.0.0.0

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet1/2(untrust-vr)

Gateway IP Address: 1.1.1.250

6. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), corp

Destination Address:

Address Book Entry: (select), serverB

Service: FTP

Action: Permit

Position at Top: (select)

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Policy configuration page:

NAT:

Source Translation: (select)

DIP On: 5 (10.10.1.2–10.10.1.2)/X-late

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), branch1

Destination Address:

Address Book Entry: (select), virtualA

Service: FTP

Action: Permit

Position at Top: (select)

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Policy configuration page:

NAT:

Destination Translation: (select)

Translate to IP: (select), 10.1.1.5

Map to Port: (clear)

WebUI (NetScreen-B)

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 2.2.2.2/24

Network > Interfaces > New Tunnel IF: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.1

Zone (VR): Untrust (trust-vr)

Fixed IP: (select)

IP Address / Netmask: 10.20.1.1/24

2. DIP

Network > Interfaces > Edit (for tunnel.1) > DIP > New: Enter the following, and then click **OK**:

ID: 6

IP Address Range: (select), 10.20.1.2 ~ 10.20.1.2

Port Translation: (select)

In the same subnet as the interface IP or its secondary IPs: (select)

3. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: branch1

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.0/24

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: virtualB

IP Address/Domain Name:

IP/Netmask: (select), 10.20.1.5/32

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: corp

IP Address/Domain Name:

IP/Netmask: (select), 10.10.1.2/32

Zone: Untrust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: serverA

IP Address/Domain Name:

IP/Netmask: (select), 10.10.1.5/32

Zone: Untrust

4. VPN

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: vpn1

Security Level: Compatible

Remote Gateway: Create a Simple Gateway: (select)

Gateway Name: corp

Type: Static IP: (select), Address/Hostname: 1.1.1.1

Preshared Key: netscreen1

Security Level: Compatible

Outgoing Interface: ethernet3[26]

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Bind to: Tunnel Interface, tunnel.1

Proxy-ID: (select)

Local IP / Netmask: 10.20.1.0/24

Remote IP / Netmask: 10.10.1.0/24

Service: ANY

26. The outgoing interface does not have to be in the same zone to which the tunnel interface is bound, although in this case they are in the same zone.

5. Routes

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 10.10.1.0/24

Gateway: (select)

Interface: tunnel.1

Gateway IP Address: 0.0.0.0

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet1/2(untrust-vr)

Gateway IP Address: 2.2.2.250

6. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), corp

Destination Address:

Address Book Entry: (select), serverA

Service: FTP

Action: Permit

Position at Top: (select)

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Policy configuration page:

NAT:

Source Translation: (select)

DIP on: 6 (10.20.1.2–10.20.1.2)/X-late

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), corp

Destination Address:

Address Book Entry: (select), virtualB

Service: FTP

Action: Permit

Position at Top: (select)

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Policy configuration page:

NAT:

Destination Translation: (select)

Translate to IP: 10.1.1.5

Map to Port: (clear)

CLI (NetScreen-A)

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24

set interface tunnel.1 zone untrust
set interface tunnel.1 ip 10.10.1.1/24
```

2. DIP

```
set interface tunnel.1 dip 5 10.10.1.2 10.10.1.2
```

3. Addresses

```
set address trust corp 10.1.1.0/24
set address trust virtualA 10.10.1.5/32
set address untrust branch1 10.20.1.2/32
set address untrust serverB 10.20.1.5/32
```

4. VPN

```
set ike gateway branch1 address 2.2.2.2 outgoing-interface ethernet3[27] preshare
    netscreen1 sec-level compatible
set vpn vpn1 gateway branch1 sec-level compatible
set vpn vpn1 bind interface tunnel.1
set vpn vpn1 proxy-id local-ip 10.10.1.0/24 remote-ip 10.20.1.0/24 any
```

27. The outgoing interface does not have to be in the same zone to which the tunnel interface is bound, although in this case they are in the same zone.

5. Routes

```
set vrouter trust-vr route 10.20.1.0/24 interface tunnel.1
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
```

6. Policies

```
set policy top from trust to untrust corp serverB ftp nat src dip-id 5 permit
set policy top from untrust to trust branch1 virtualA ftp nat dst ip 10.1.1.5
    permit
save
```

CLI (NetScreen-B)

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 2.2.2.2/24

set interface tunnel.1 zone untrust
set interface tunnel.1 ip 10.20.1.1/24
```

2. DIP

```
set interface tunnel.1 dip 6 10.20.1.2 10.20.1.2
```

3. Addresses

```
set address trust branch1 10.1.1.0/24
set address trust virtualB 10.20.1.5/32
set address untrust corp 10.10.1.2/32
set address untrust serverA 10.10.1.5/32
```

4. VPN

```
set ike gateway corp address 1.1.1.1 outgoing-interface ethernet3[28] preshare
    netscreen1 sec-level compatible
set vpn vpn1 gateway corp sec-level compatible
set vpn vpn1 bind interface tunnel.1
set vpn vpn1 proxy-id local-ip 10.20.1.0/24 remote-ip 10.10.1.0/24 any
```

5. Routes

```
set vrouter trust-vr route 10.10.1.0/24 interface tunnel.1
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 2.2.2.250
```

6. Policies

```
set policy top from trust to untrust branch1 serverA ftp nat src dip-id 6
    permit
set policy top from untrust to trust corp virtualB ftp nat dst ip 10.1.1.5
    permit
save
```

28. The outgoing interface does not have to be in the same zone to which the tunnel interface is bound, although in this case they are in the same zone.

TRANSPARENT MODE VPN

When the NetScreen device interfaces are in Transparent mode (that is, they have no IP addresses and are operating at Layer 2 in the OSI model[29]), you can use the VLAN1 IP address as a VPN termination point. In place of an outgoing interface, as used when the interfaces are in Route or NAT mode (that is, they have IP addresses and are operating at Layer 3), a VPN tunnel references an outgoing zone. By default, a tunnel uses the V1-Untrust zone as its outgoing zone. If you have multiple interfaces bound to the same outgoing zone, the VPN tunnel can use any one of them.

Note: *At the time of this release, a NetScreen device whose interfaces are in Transparent mode supports only policy-based VPNs. For more information about Transparent mode, see "Transparent Mode" on page* ***2*** *-92.*

29. The OSI model is a networking industry standard model of network protocol architecture. The OSI model consists of seven layers, in which layer 2 is the data link layer and layer 3 is the network layer.

Example: Transparent Mode, Policy-Based AutoKey IKE VPN

In this example, you set up a policy-based AutoKey IKE VPN tunnel between two NetScreen devices with interfaces operating in Transparent mode.

Note: *It is not necessary that the interfaces of both NetScreen devices be in Transparent mode. The interfaces of the device at one end of the tunnel can be in Transparent mode and those of the other device can be in Route or NAT mode.*

The key elements of the configuration for the NetScreen devices at both ends of the tunnel are as follows:

Configuration Elements	NetScreen-A	NetScreen-B
V1-Trust Zone	Interface: ethernet1, 0.0.0.0/0 (enable management for the local admin)	Interface: ethernet1, 0.0.0.0/0 (enable management for the local admin)
V1-Untrust Zone	Interface: ethernet3, 0.0.0.0/0	Interface: ethernet3, 0.0.0.0/0
VLAN1 Interface	IP Address: 1.1.1.1/24 Manage IP: 1.1.1.2*	IP Address: 2.2.2.2/24 Manage IP: 2.2.2.3
Addresses	local_lan: 1.1.1.0/24 in V1-Trust peer_lan: 2.2.2.0/24 in V1-Untrust	local_lan: 2.2.2.0/24 in V1-Trust peer_lan: 1.1.1.0/24 in V1-Untrust
IKE gateway	gw1, 2.2.2.2, preshared key h1p8A24nG5, security: compatible	gw1, 1.1.1.1, preshared key h1p8A24nG5, security: compatible
VPN tunnel	security: compatible	security: compatible
Policies	local_lan -> peer_lan, any service, vpn1 peer_lan -> local_lan, any service, vpn1	local_lan -> peer_lan, any service, vpn1 peer_lan -> local_lan, any service, vpn1
External Router	IP Address: 1.1.1.250	IP Address: 2.2.2.250
Route	0.0.0.0/0, use VLAN1 interface to gateway 1.1.1.250	0.0.0.0/0, use VLAN1 interface to gateway 2.2.2.250

* You can separate administrative from VPN traffic by using the manage IP address to receive administrative traffic and the VLAN1 address to terminate VPN traffic.

Configuring a policy-based AutoKey IKE tunnel for a NetScreen device whose interfaces are in Transparent mode involves the following steps:

1. Remove any IP addresses from the physical interfaces, and bind them to the layer-2 security zones.
2. Assign an IP address and manage IP address to the VLAN1 interface.
3. Enter the IP addresses for the local and remote endpoints in the address books for the V1-Trust and V1-Untrust zones.
4. Configure the VPN tunnel and designate its outgoing zone as the V1-Untrust zone.
5. Enter a default route to the external router in the trust-vr.
6. Set up policies for VPN traffic to pass between each site.

WebUI (NetScreen-A)

1. Interfaces

Note: *Moving the VLAN1 IP address to a different subnet causes the NetScreen device to delete any routes involving the previous VLAN1 interface. When configuring a NetScreen device through the WebUI, your workstation must reach the first VLAN1 address and then be in the same subnet as the new address. After changing the VLAN1 address, you must then change the IP address of your workstation so that it is in the same subnet as the new VLAN1 address. You might also have to relocate your workstation to a subnet physically adjacent to the NetScreen device.*

Network > Interfaces > Edit (for the VLAN1 interface): Enter the following, and then click **OK**:

IP Address/Netmask: 1.1.1.1/24

Manage IP: 1.1.1.2

Management Services: WebUI, Telnet, Ping[30]

30. You enable the management options for WebUI, Telnet, and Ping on both the V1-Trust zone and the VLAN1 interface so that a local admin in the V1-Trust zone can reach the VLAN1 manage IP address. If management via the WebUI is not already enabled on VLAN1 and the V1-Trust zone interfaces, you cannot reach the NetScreen device through the WebUI to make these settings. Instead, you must first set WebUI manageability on these interfaces through a console connection.

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Management Services: WebUI, Telnet

Other Services: Ping

Select the following, and then click **OK**:

Zone Name: V1-Trust

IP Address/Netmask: 0.0.0.0/0

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: V1-Untrust

IP Address/Netmask: 0.0.0.0/0

2. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: local_lan

IP Address/Domain Name:

IP/Netmask: (select), 1.1.1.0/24

Zone: V1-Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: peer_lan

IP Address/Domain Name:

IP/Netmask: (select), 2.2.2.0/24

Zone: V1-Untrust

3. VPN

VPNs > AutoKey Advanced > Gateway > New: Enter the following, and then click **OK**:

Gateway Name: gw1

Security Level: Compatible

Remote Gateway Type:

Static IP Address: (select), IP Address/Hostname: 2.2.2.2

Preshared Key: h1p8A24nG5

Outgoing Zone: V1-Untrust

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: vpn1

Security Level: Compatible

Remote Gateway:

Predefined: (select), gw1

4. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: VLAN1 (VLAN)

Gateway IP Address: 1.1.1.250

5. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), local_lan

Destination Address:

Address Book Entry: (select), peer_lan

Service: ANY

Action: Tunnel

Tunnel VPN: vpn1

Modify matching bidirectional VPN policy: (select)

Position at Top: (select)

WebUI (NetScreen-B)

1. Interfaces

Note: *Moving the VLAN1 IP address to a different subnet causes the NetScreen device to delete any routes involving the previous VLAN1 interface. When configuring a NetScreen device through the WebUI, your workstation must reach the first VLAN1 address and then be in the same subnet as the new address. After changing the VLAN1 address, you must then change the IP address of your workstation so that it is in the same subnet as the new VLAN1 address. You might also have to relocate your workstation to a subnet physically adjacent to the NetScreen device.*

Network > Interfaces > Edit (for the VLAN1 interface): Enter the following, and then click **OK**:

IP Address/Netmask: 2.2.2.2/24

Manage IP: 2.2.2.3

Management Services: WebUI[31], Telnet, Ping

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Management Services: WebUI, Telnet

Other Services: Ping

Select the following, and then click **OK**:

Zone Name: V1-Trust

IP Address/Netmask: 0.0.0.0/0

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: V1-Untrust

IP Address/Netmask: 0.0.0.0/0

31. If management via the WebUI is not already enabled on VLAN1 and the V1-Trust zone interfaces, you cannot reach the NetScreen device through the WebUI to make these settings. Instead, you must first set WebUI manageability on these interfaces through a console connection.

2. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: local_lan

IP Address/Domain Name:

IP/Netmask: (select), 2.2.2.0/24

Zone: V1-Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: peer_lan

IP Address/Domain Name:

IP/Netmask: (select), 1.1.1.0/24

Zone: V1-Untrust

3. VPN

VPNs > AutoKey Advanced > Gateway > New: Enter the following, and then click **OK**:

Gateway Name: gw1

Security Level: Compatible

Remote Gateway Type:

Static IP Address: (select), IP Address/Hostname: 1.1.1.1

Preshared Key: h1p8A24nG5

Outgoing Zone: V1-Untrust

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: vpn1

Security Level: Compatible

Remote Gateway:

Predefined: (select), gw1

4. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: VLAN1 (VLAN)

Gateway IP Address: 2.2.2.250

5. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), local_lan

Destination Address:

Address Book Entry: (select), peer_lan

Service: ANY

Action: Tunnel

Tunnel VPN: vpn1

Modify matching bidirectional VPN policy: (select)

Position at Top: (select)

CLI (NetScreen-A)

1. Interfaces and Zones

```
unset interface ethernet1 ip
unset interface ethernet1 zone
set interface ethernet1 zone v1-trust
set zone v1-trust manage web
set zone v1-trust manage telnet
set zone v1-trust manage ping[32]

unset interface ethernet3 ip
unset interface ethernet3 zone
set interface ethernet3 zone v1-untrust

set interface vlan1 ip 1.1.1.1/24
set interface vlan1 manage-ip 1.1.1.2
set interface vlan1 manage web
set interface vlan1 manage telnet
set interface vlan1 manage ping
```

2. Addresses

```
set address v1-trust local_lan 1.1.1.0/24
set address v1-untrust peer_lan 2.2.2.0/24
```

3. VPN

```
set ike gateway gw1 address 2.2.2.2 main outgoing-interface v1-untrust preshare
    h1p8A24nG5 sec-level compatible
set vpn vpn1 gateway gw1 sec-level compatible
```

32. You enable the management options for WebUI, Telnet, and Ping on both the V1-Trust zone and the VLAN1 interface so that a local admin in the V1-Trust zone can reach the VLAN1 manage IP address.

4. Routes

```
set vrouter trust-vr route 0.0.0.0/0 interface vlan1 gateway 1.1.1.250
```

5. Policies

```
set policy top from v1-trust to v1-untrust local_lan peer_lan any tunnel vpn
    vpn1
set policy top from v1-untrust to v1-trust peer_lan local_lan any tunnel vpn
    vpn1
save
```

CLI (NetScreen-B)

1. Interfaces and Zones

```
unset interface ethernet1 ip
unset interface ethernet1 zone
set interface ethernet1 zone v1-trust
set zone v1-trust manage

unset interface ethernet3 ip
unset interface ethernet3 zone
set interface ethernet3 zone v1-untrust

set interface vlan1 ip 2.2.2.2/24
set interface vlan1 manage-ip 2.2.2.3
set interface vlan1 manage
```

2. Addresses

```
set address v1-trust local_lan 2.2.2.0/24
set address v1-untrust peer_lan 1.1.1.0/24
```

3. VPN

```
set ike gateway gw1 address 1.1.1.1 main outgoing-interface v1-untrust preshare
    h1p8A24nG5 sec-level compatible
set vpn vpn1 gateway gw1 sec-level compatible
```

4. Routes

```
set vrouter trust-vr route 0.0.0.0/0 interface vlan1 gateway 2.2.2.250
```

5. Policies

```
set policy top from v1-trust to v1-untrust local_lan peer_lan any tunnel vpn
    vpn1
set policy top from v1-untrust to v1-trust peer_lan local_lan any tunnel vpn
    vpn1
save
```

5

Dialup VPNs

NetScreen devices can support dialup VPN connections. You can configure a NetScreen device that has a static IP address to secure an IPSec tunnel with a NetScreen-Remote client or with another NetScreen device with a dynamic IP address.

This chapter offers examples of the following dialup VPN concepts:

- "Dialup VPNs" on page 200
  - "Example: Policy-Based Dialup VPN, AutoKey IKE" on page 201
  - "Example: Route-Based Dialup VPN, Dynamic Peer" on page 209
  - "Example: Policy-Based Dialup VPN, Dynamic Peer" on page 220
  - "Example: Bidirectional Dialup VPN Policies" on page 230
- "Group IKE ID" on page 237
  - "Example: Group IKE ID (Certificates)" on page 243
  - "Example: Group IKE ID (Preshared Keys)" on page 252
- "Shared IKE IDs" on page 259
  - "Example: Shared IKE ID (Preshared Keys)" on page 260

DIALUP VPNS

You can configure tunnels for VPN dialup users on a per-user basis or form users into a VPN dialup group for which you need only configure one tunnel. You can also create a group IKE ID user, which allows you to define one user whose IKE ID is used as part of the IKE IDs of dialup IKE users. This approach is particularly timesaving when there are large groups of dialup users because you do not have to configure each IKE user individually.

Note: *For more information on creating IKE user groups, see "IKE Users and User Groups" on page* ***2*** *-431. For more information about the Group IKE ID feature, see "Group IKE ID" on page 237.*

If the dialup client can support a virtual internal IP address, which the NetScreen-Remote does, you can also create a dynamic peer dialup VPN, AutoKey IKE tunnel (with a preshared key or certificates). You can configure a NetScreen security gateway with a static IP address to secure an IPSec tunnel with a NetScreen-Remote client or with another NetScreen device with a dynamic IP address.

Note: *For background information about the available VPN options, see Chapter 1, "IPSec". For guidance when choosing among the various options, see Chapter 3, "VPN Guidelines".*

You can configure policy-based VPN tunnels for VPN dialup users. For a dialup dynamic peer client[1], you can configure either a policy-based or route-based VPN. Because a dialup dynamic peer client can support a virtual internal IP address, which the NetScreen-Remote does, you can configure a routing table entry to that virtual internal address via a designated tunnel interface. Doing so allows you to configure a route-based VPN tunnel between the NetScreen device and that peer.

Note: *The dialup dynamic peer is nearly identical to the Site-to-Site dynamic peer except that the internal IP address for the dialup client is a virtual address.*

1. A dialup dynamic peer client is a dialup client that supports a virtual internal IP address.

Example: Policy-Based Dialup VPN, AutoKey IKE

In this example, an AutoKey IKE tunnel using either a preshared key or a pair of certificates (one at each end of the tunnel[2]) provides the secure communication channel between the IKE user Wendy and the UNIX server. The tunnel again uses ESP with 3DES encryption and SHA-1 authentication.

Setting up the AutoKey IKE tunnel using AutoKey IKE with either a preshared key or certificates requires you to do the following at the corporate site:

1. Configure interfaces for the Trust and Untrust zones, both of which are in the trust-vr routing domain.
2. Enter the address of the UNIX server in the Trust zone address book.
3. Define Wendy as an IKE user.
4. Configure the remote gateway and AutoKey IKE VPN.
5. Set up a default route.
6. Create a policy from the Untrust zone to the Trust zone permitting access to the UNIX from the dialup user.

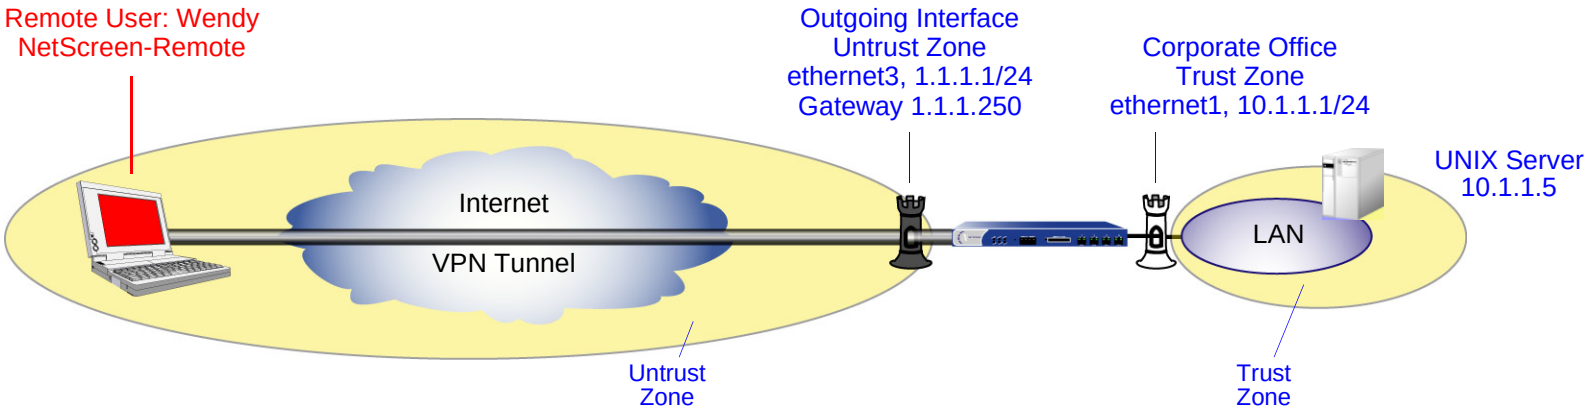

2. The preshared key is h1p8A24nG5. It is assumed that both participants already have certificates. For more information about certificates, see "Certificates and CRLs" on page 21.

The preshared key is h1p8A24nG5. This example assumes that both participants already have RSA certificates issued by Verisign, and that the local certificate on the NetScreen-Remote contains the U-FQDN wparker@email.com. (For information about obtaining and loading certificates, see “Certificates and CRLs” on page 21.) For the Phase 1 and 2 security levels, you specify one Phase 1 proposal—either pre-g2-3des-sha for the preshared key method or rsa-g2-3des-sha for certificates—and select the predefined “Compatible” set of proposals for Phase 2.

WebUI

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

2. Address

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: UNIX

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.5/32

Zone: Trust

3. User

Objects > Users > Local > New: Enter the following, and then click **OK**:

User Name: Wendy

Status: Enable (select)

IKE User: (select)

Simple Identity: (select)

IKE Identity: wparker@email.com

4. VPN

VPNs > AutoKey Advanced > Gateway > New: Enter the following, and then click **OK**:

Gateway Name: Wendy_NSR

Security Level: Custom

Remote Gateway Type:

Dialup User: (select), User: Wendy

Preshared Key

Preshared Key: h1p8A24nG5

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Gateway configuration page:

Security Level: Custom

Phase 1 Proposal (For Custom Security Level):
pre-g2-3des-sha

Mode (Initiator): Aggressive

(or)

Certificates

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Gateway configuration page:

Security Level: Custom

Phase 1 Proposal (For Custom Security Level): rsa-g2-3des-sha

Mode (Initiator): Aggressive

Preferred Certificate (optional):

Peer CA: Verisign

Peer Type: X509-SIG

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: Wendy_UNIX

Security Level: Compatible

Remote Gateway:

Predefined: (select), Wendy_NSR

5. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 1.1.1.250

6. Policy

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Dial-Up VPN

Destination Address:

Address Book Entry: (select), UNIX

Service: ANY

Action: Tunnel

Tunnel VPN: Wendy_UNIX

Modify matching bidirectional VPN policy: (clear)

Position at Top: (select)

CLI

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
```

2. Address

```
set address trust unix 10.1.1.5/32
```

3. User

```
set user wendy ike-id u-fqdn wparker@email.com
```

4. VPN

Preshared Key

```
set ike gateway wendy_nsr dialup wendy aggressive outgoing-interface ethernet3
    preshare h1p8A24nG5 proposal pre-g2-3des-sha
set vpn wendy_unix gateway wendy_nsr sec-level compatible
```

(or)

Certificates

```
set ike gateway wendy_nsr dialup wendy aggressive outgoing-interface ethernet3
    proposal rsa-g2-3des-sha
set ike gateway wendy_nsr cert peer-ca 1[3]
set ike gateway wendy_nsr cert peer-cert-type x509-sig
set vpn wendy_unix gateway wendy_nsr sec-level compatible
```

5. Route

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
```

6. Policy

```
set policy top from untrust to trust “Dial-Up VPN” unix any tunnel vpn
    wendy_unix
save
```

3. The number 1 is the CA ID number. To discover the CA’s ID number, use the following command: **get pki x509 list ca-cert**.

NetScreen-Remote Security Policy Editor

1. Click **Options > Secure > Specified Connections**.
2. Click **Add a new connection**, and type **UNIX** next to the new connection icon that appears.
3. Configure the connection options:

   Connection Security: Secure

   Remote Party Identity and Addressing:

   ID Type: IP Address, 10.1.1.5

   Protocol: All

   Connect using Secure Gateway Tunnel: (select)

   ID Type: IP Address, 1.1.1.1
4. Click the **PLUS** symbol, located to the left of the UNIX icon, to expand the connection policy.
5. Click **My Identity**: Do either of the following:

   Click **Pre-shared Key** > **Enter Key**: Type **h1p8A24nG5**, and then click **OK**.

   ID Type: (select **E-mail Address**), and type **wparker@email.com**.

   (or)

   Select a certificate from the Select Certificate drop-down list.

   ID Type: (select **E-mail Address**)[4]
6. Click the **Security Policy** icon, and select **Aggressive Mode**.
7. Click the **PLUS** symbol, located to the left of the Security Policy icon, and then the **PLUS** symbol to the left of Authentication (Phase 1) and Key Exchange (Phase 2) to expand the policy further.

4. The e-mail address from the certificate appears in the identifier field automatically.

8. Click **Authentication (Phase 1)** > **Proposal 1**: Select the following Encryption and Data Integrity Algorithms:

   Encrypt Alg: Triple DES

   Hash Alg: SHA-1

   Key Group: Diffie-Hellman Group 2

9. Click **Key Exchange (Phase 2)** > **Proposal 1**: Select the following IPSec Protocols:

   Encapsulation Protocol (ESP): (select)

   Encrypt Alg: Triple DES

   Hash Alg: SHA-1

   Encapsulation: Tunnel

10. Click **Key Exchange (Phase 2)** > **Create New Proposal**: Select the following IPSec Protocols:

    Encapsulation Protocol (ESP): (select)

    Encrypt Alg: Triple DES

    Hash Alg: MD5

    Encapsulation: Tunnel

11. Click **Key Exchange (Phase 2)** > **Create New Proposal**: Select the following IPSec Protocols:

    Encapsulation Protocol (ESP): (select)

    Encrypt Alg: DES

    Hash Alg: SHA-1

    Encapsulation: Tunnel

12. Click **Key Exchange (Phase 2)** > **Create New Proposal**: Select the following IPSec Protocols:

    Encapsulation Protocol (ESP): (select)

    Encrypt Alg: DES

    Hash Alg: MD5

    Encapsulation: Tunnel

13. Click **Save**.

Example: Route-Based Dialup VPN, Dynamic Peer

In this example, a VPN tunnel securely connects the user behind the NetScreen-Remote to the Untrust zone interface of the NetScreen device protecting the mail server in the DMZ zone. The Untrust zone interface has a static IP address. The NetScreen-Remote client has a dynamically assigned external IP address and a static (virtual) internal IP address. The administrator of the NetScreen device must know the peer's internal IP address so that he can add it to the Untrust address book for use in policies to tunnel traffic from that source. After the NetScreen-Remote client establishes the tunnel, traffic through the tunnel can originate from either end.

All zones on the NetScreen device are in the trust-vr routing domain.

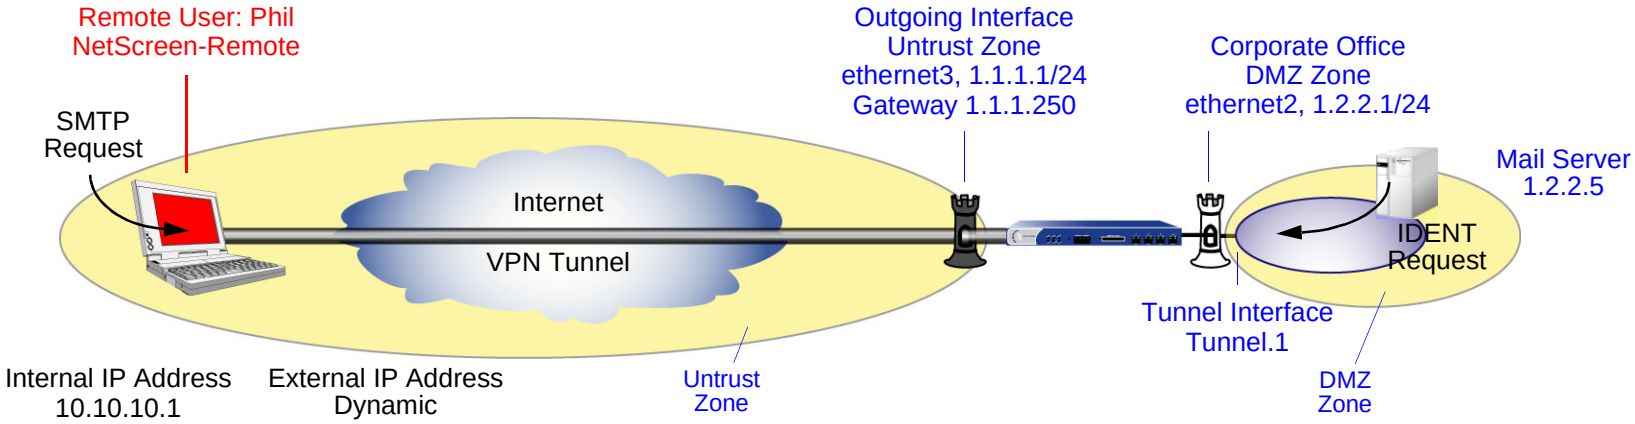

In this example, Phil wants to get his e-mail from the mail server at the company site. When he attempts to do so, he is authenticated by the mail server program, which sends him an IDENT request through the tunnel.

Note: *The mail server can send the IDENT request through the tunnel only if the NetScreen administrator adds a custom service for it (TCP, port 113) and sets up an outgoing policy allowing that traffic through the tunnel to 10.10.10.1.*

The preshared key is h1p8A24nG5. It is assumed that both participants already have RSA certificates issued by Verisign, and that the local certificate on the NetScreen-Remote contains the U-FQDN *pm@netscreen.com.* (For information about obtaining and loading certificates, see "Certificates and CRLs" on page 21.) For the Phase 1 and 2 security levels, you specify one Phase 1 proposal—either pre-g2-3des-sha for the preshared key method or rsa-g2-3des-sha for certificates—and select the predefined "Compatible" set of proposals for Phase 2.

WebUI

1. Interfaces

Network > Interfaces > Edit (for ethernet2): Enter the following, and then click **OK**:

Zone Name: DMZ

Static IP: (select this option when present)

IP Address/Netmask: 1.2.2.1/24

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

Network > Interfaces > New Tunnel IF: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.1

Zone (VR): Untrust (trust-vr)

Unnumbered: (select)

Interface: ethernet3 (trust-vr)

2. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Mail Server

IP Address/Domain Name:

IP/Netmask: (select), 1.2.2.5/32

Zone: DMZ

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Phil

IP Address/Domain Name:

IP/Netmask: (select), 10.10.10.1/32

Zone: Untrust

3. Services

Objects > Services > Custom > New: Enter the following, and then click **OK**:

Service Name: Ident

Service Timeout:

Use protocol default: (select)

Transport Protocol: TCP (select)

Source Port: Low 1, High 65535

Destination Port: Low 113, High 113

Objects > Services > Group > New: Enter the following, move the following services, and then click **OK**:

Group Name: Remote_Mail

Group Members << Available Members:

Ident

MAIL

POP3

4. VPN

VPNs > AutoKey Advanced > Gateway > New: Enter the following, and then click **OK**:

Gateway Name: To_Phil

Security Level: Custom

Remote Gateway Type:

Dynamic IP Address: (select), Peer ID: pm@netscreen.com

Preshared Key

Preshared Key: h1p8A24nG5

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Gateway configuration page:

Security Level: Custom

Phase 1 Proposal (For Custom Security Level): pre-g2-3des-sha

Mode (Initiator): Aggressive

(or)

Certificates

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Gateway configuration page:

Security Level: Custom

Phase 1 Proposal (For Custom Security Level): rsa-g2-3des-sha

Mode (Initiator): Aggressive

Preferred Certificate (optional):

Peer CA: Verisign

Peer Type: X509-SIG

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: corp_Phil

Security Level: Compatible

Remote Gateway:

Predefined: (select), To_Phil

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Bind to: Tunnel Interface: (select), tunnel.1

Proxy-ID: (select)

Local IP / Netmask: 1.2.2.5/32

Remote IP / Netmask: 10.10.10.1/32

Service: Any

5. Routes

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 1.1.1.250

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 10.10.10.1/32

Gateway: (select)

Interface: tunnel.1

Gateway IP Address: 0.0.0.0

6. Policies

Policies > (From: Untrust, To: DMZ) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Phil

Destination Address:

Address Book Entry: (select), Mail Server

Service: Remote_Mail

Action: Permit

Position at Top: (select)

Policies > (From: DMZ, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Mail Server

Destination Address:

Address Book Entry: (select), Phil

Service: Remote_Mail

Action: Permit

Position at Top: (select)

CLI

1. Interfaces

```
set interface ethernet2 zone dmz
set interface ethernet2 ip 1.2.2.1/24

set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24

set interface tunnel.1 zone untrust
set interface tunnel.1 ip unnumbered interface ethernet3
```

2. Addresses

```
set address dmz “Mail Server” 1.2.2.5/32
set address untrust phil 10.10.10.1/32
```

3. Services

```
set service ident protocol tcp src-port 1-65535 dst-port 113-113
set group service remote_mail
set group service remote_mail add ident
set group service remote_mail add mail
set group service remote_mail add pop3
```

4. VPN

Preshared Key

```
set ike gateway to_phil dynamic pm@netscreen.com aggressive outgoing-interface
    ethernet3 preshare h1p8A24nG5 proposal pre-g2-3des-sha
set vpn corp_phil gateway to_phil sec-level compatible
set vpn corp_phil bind interface tunnel.1
set vpn corp_phil proxy-id local-ip 1.2.2.5/32 remote-ip 10.10.10.1/32 any
```

(or)

Certificates

```
set ike gateway to_phil dynamic pm@netscreen.com aggressive outgoing-interface
    ethernet3 proposal rsa-g2-3des-sha
set ike gateway to_phil cert peer-ca 1[5]
set ike gateway to_phil cert peer-cert-type x509-sig
set vpn corp_phil gateway to_phil sec-level compatible
set vpn corp_phil bind interface tunnel.1
set vpn corp_phil proxy-id local-ip 1.2.2.5/32 remote-ip 10.10.10.1/32 any
```

5. Routes

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
set vrouter trust-vr route 10.10.10.1/32 interface tunnel.1
```

6. Policies

```
set policy top from dmz to untrust “Mail Server” phil remote_mail permit
set policy top from untrust to dmz phil “Mail Server” remote_mail permit
save
```

5. The number 1 is the CA ID number. To discover the CA’s ID number, use the following command: **get pki x509 list ca-cert**.

NetScreen-Remote

1. Click **Options** > **Global Policy Settings**, and select the **Allow to Specify Internal Network Address** check box.
2. **Options > Secure > Specified Connections**.
3. Click the **Add a new connection** button, and type **Mail** next to the new connection icon that appears.
4. Configure the connection options:

   Connection Security: Secure

   Remote Party Identity and Addressing:

   ID Type: IP Address, 1.2.2.5

   Protocol: All

   Connect using Secure Gateway Tunnel: (select)

   ID Type: IP Address, 1.1.1.1
5. Click the **PLUS** symbol, located to the left of the unix icon, to expand the connection policy.
6. Click the **Security Policy** icon, and select **Aggressive Mode**.
7. Click **My Identity** and do either of the following:

   Click **Pre-shared Key** > **Enter Key**: Type **h1p8A24nG5**, and then click **OK**.

   ID Type: E-mail Address; pm@netscreen.com

   Internal Network IP Address: 10.10.10.1

   or

   Select the certificate that contains the e-mail address “pm@netscreen.com” from the Select Certificate drop-down list.

   ID Type: E-mail Address; pm@netscreen.com

   Internal Network IP Address: 10.10.10.1

8. Click the **PLUS** symbol, located to the left of the Security Policy icon, and then the **PLUS** symbol to the left of Authentication (Phase 1) and Key Exchange (Phase 2) to expand the policy further.
9. Click **Authentication (Phase 1)** > **Proposal 1**: Select the following Encryption and Data Integrity Algorithms:

   Encrypt Alg: Triple DES

   Hash Alg: SHA-1

   Key Group: Diffie-Hellman Group 2
10. Click **Key Exchange (Phase 2)** > **Proposal 1**: Select the following IPSec Protocols:

    Encapsulation Protocol (ESP): (select)

    Encrypt Alg: Triple DES

    Hash Alg: SHA-1

    Encapsulation: Tunnel
11. Click **Key Exchange (Phase 2)** > **Create New Proposal**: Select the following IPSec Protocols:

    Encapsulation Protocol (ESP): (select)

    Encrypt Alg: Triple DES

    Hash Alg: MD5

    Encapsulation: Tunnel
12. Click **Key Exchange (Phase 2)** > **Create New Proposal**: Select the following IPSec Protocols:

    Encapsulation Protocol (ESP): (select)

    Encrypt Alg: DES

    Hash Alg: SHA-1

    Encapsulation: Tunnel

13. Click **Key Exchange (Phase 2)** > **Create New Proposal**: Select the following IPSec Protocols:

    Encapsulation Protocol (ESP): (select)

    Encrypt Alg: DES

    Hash Alg: MD5

    Encapsulation: Tunnel

14. Click the **Save** button.

Example: Policy-Based Dialup VPN, Dynamic Peer

In this example, a VPN tunnel securely connects the user behind the NetScreen-Remote to the Untrust zone interface of the NetScreen device protecting the mail server in the DMZ zone. The Untrust zone interface has a static IP address. The NetScreen-Remote client has a dynamically assigned external IP address and a static (virtual) internal IP address. The administrator of the NetScreen device must know the client's internal IP address so that he can add it to the Untrust address book for use in policies to tunnel traffic from that source. After the NetScreen-Remote client establishes the tunnel, traffic through the tunnel can originate from either end.

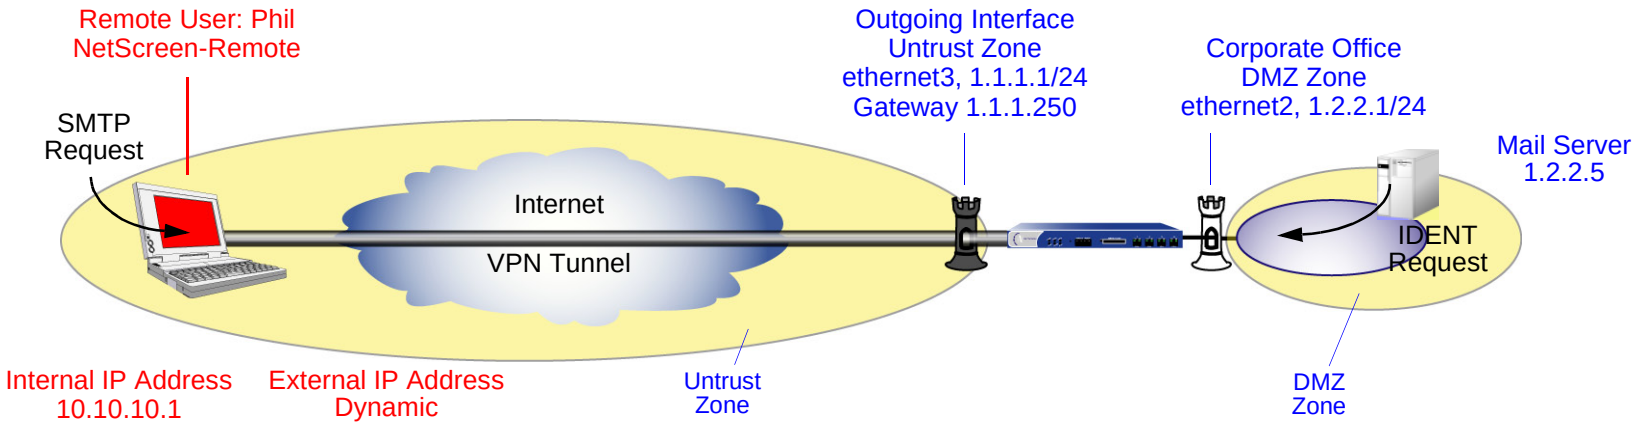

In this example, Phil wants to get his e-mail from the mail server at the company site. When he attempts to do so, he is authenticated by the mail server program, which sends him an IDENT request through the tunnel.

Note: *The mail server can send the IDENT request through the tunnel only if the NetScreen administrator adds a custom service for it (TCP, port 113) and sets up an outgoing policy allowing that traffic through the tunnel to 10.10.10.1.*

The preshared key is h1p8A24nG5. This example assumes that both participants have RSA certificates issued by Verisign, and that the local certificate on the NetScreen-Remote contains the U-FQDN *pm@netscreen.com.* (For more information about obtaining and loading certificates, see "Certificates and CRLs" on page 21.) For the Phase 1 and 2 security levels, you specify one Phase 1 proposal—either pre-g2-3des-sha for the preshared key method or rsa-g2-3des-sha for certificates—and select the predefined "Compatible" set of proposals for Phase 2.

WebUI

1. Interfaces

Network > Interfaces > Edit (for ethernet2): Enter the following, and then click **OK**:

Zone Name: DMZ

Static IP: (select this option when present)

IP Address/Netmask: 1.2.2.1/24

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

2. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Mail Server

IP Address/Domain Name:

IP/Netmask: (select), 1.2.2.5/32

Zone: DMZ

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Phil

IP Address/Domain Name:

IP/Netmask: (select), 10.10.10.1/32

Zone: Untrust

3. Services

Objects > Services > Custom > New: Enter the following, and then click **OK**:

Service Name: Ident

Service Timeout:

Use protocol default: (select)

Transport Protocol: TCP (select)

Source Port: Low 1, High 65535

Destination Port: Low 113, High 113

Objects > Services > Group > New: Enter the following, move the following services, and then click **OK**:

Group Name: Remote_Mail

Group Members << Available Members:

Ident
MAIL
POP3

4. VPN

VPNs > AutoKey Advanced > Gateway > New: Enter the following, and then click **OK**:

Gateway Name: To_Phil

Security Level: Custom

Remote Gateway Type:

Dynamic IP Address: (select), Peer ID: pm@netscreen.com

Preshared Key

Preshared Key: h1p8A24nG5

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Gateway configuration page:

Security Level: Custom

Phase 1 Proposal (For Custom Security Level): pre-g2-3des-sha

Mode (Initiator): Aggressive

(Or)

Certificates

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Gateway configuration page:

Security Level: Custom

Phase 1 Proposal (For Custom Security Level): rsa-g2-3des-sha

Mode (Initiator): Aggressive

Preferred Certificate (optional):
Peer CA: Verisign
Peer Type: X509-SIG

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: corp_Phil
Security Level: Compatible
Remote Gateway:
Predefined: (select), To_Phil

5. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0
Gateway: (select)
Interface: ethernet3
Gateway IP Address: 1.1.1.250

6. Policies

Policies > (From: Untrust, To: DMZ) New: Enter the following, and then click **OK**:

Source Address:
Address Book Entry: (select), Phil
Destination Address:
Address Book Entry: (select), Mail Server
Service: Remote_Mail
Action: Tunnel
VPN Tunnel: corp_Phil
Modify matching bidirectional VPN policy: (select)
Position at Top: (select)

CLI

1. Interfaces

```
set interface ethernet2 zone dmz
set interface ethernet2 ip 1.2.2.1/24

set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
```

2. Addresses

```
set address dmz “mail server” 1.2.2.5/32
set address untrust phil 10.10.10.1/32
```

3. Services

```
set service ident protocol tcp src-port 1-65535 dst-port 113-113
set group service remote_mail
set group service remote_mail add ident
set group service remote_mail add mail
set group service remote_mail add pop3
```

4. VPN

Preshared Key

```
set ike gateway to_phil dynamic pm@netscreen.com aggressive outgoing-interface
    ethernet3 preshare h1p8A24nG5 proposal pre-g2-3des-sha
set vpn corp_phil gateway to_phil sec-level compatible
```

(or)

Certificates

```
set ike gateway to_phil dynamic pm@netscreen.com aggressive outgoing-interface
    ethernet3 proposal rsa-g2-3des-sha
set ike gateway to_phil cert peer-ca 1[6]
set ike gateway to_phil cert peer-cert-type x509-sig
set vpn corp_phil gateway to_phil sec-level compatible
```

5. Route

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
```

6. Policies

```
set policy top from untrust to dmz phil “mail server” remote_mail tunnel vpn
    corp_phil
set policy top from dmz to untrust “mail server” phil remote_mail tunnel vpn
    corp_phil
save
```

6. The number 1 is the CA ID number. To discover the CA’s ID number, use the following command: **get pki x509 list ca-cert**.

NetScreen-Remote

1. Click **Options** > **Global Policy Settings**, and select **Allow to Specify Internal Network Address**.
2. **Options > Secure > Specified Connections**.
3. Click **Add a new connection**, and type **Mail** next to the new connection icon that appears.
4. Configure the connection options:

   Connection Security: Secure

   Remote Party Identity and Addressing:

   ID Type: IP Address, 1.2.2.5

   Protocol: All

   Connect using Secure Gateway Tunnel: (select)

   ID Type: IP Address, 1.1.1.1
5. Click the **PLUS** symbol, located to the left of the unix icon, to expand the connection policy.
6. Click the **Security Policy** icon, and select **Aggressive Mode**.
7. Click **My Identity** and do either of the following:

   Click **Pre-shared Key** > **Enter Key**: Type **h1p8A24nG5**, and then click **OK**.

   Internal Network IP Address: 10.10.10.1

   ID Type: E-mail Address; pm@netscreen.com

   or

   Select the certificate that contains the e-mail address “pmason@email.com” from the Select Certificate drop-down list.

   Internal Network IP Address: 10.10.10.1

   ID Type: E-mail Address; pm@netscreen.com
8. Click the **PLUS** symbol, located to the left of the Security Policy icon, and then the **PLUS** symbol to the left of Authentication (Phase 1) and Key Exchange (Phase 2) to expand the policy further.

9. Click **Authentication (Phase 1)** > **Proposal 1**: Select the following Encryption and Data Integrity Algorithms:

   Encrypt Alg: Triple DES
   Hash Alg: SHA-1
   Key Group: Diffie-Hellman Group 2

10. Click **Key Exchange (Phase 2)** > **Proposal 1**: Select the following IPSec Protocols:

    Encapsulation Protocol (ESP): (select)
    Encrypt Alg: Triple DES
    Hash Alg: SHA-1
    Encapsulation: Tunnel

11. Click **Key Exchange (Phase 2)** > **Create New Proposal**: Select the following IPSec Protocols:

    Encapsulation Protocol (ESP): (select)
    Encrypt Alg: Triple DES
    Hash Alg: MD5
    Encapsulation: Tunnel

12. Click **Key Exchange (Phase 2)** > **Create New Proposal**: Select the following IPSec Protocols:

    Encapsulation Protocol (ESP): (select)
    Encrypt Alg: DES
    Hash Alg: SHA-1
    Encapsulation: Tunnel

13. Click **Key Exchange (Phase 2)** > **Create New Proposal**: Select the following IPSec Protocols:

    Encapsulation Protocol (ESP): (select)
    Encrypt Alg: DES
    Hash Alg: MD5
    Encapsulation: Tunnel

14. Click **Save**.

Bidirectional Policies for Dialup VPN Users

You can create bidirectional policies for dialup VPNs. This configuration provides similar functionality as a dynamic peer VPN configuration. However, with a dynamic peer VPN configuration, the NetScreen device admin must know the internal IP address space of the dialup user, so that the admin can use it as the destination address when configuring an outgoing policy (see “Example: Policy-Based Dialup VPN, Dynamic Peer” on page 220). With a dialup VPN user configuration, the admin at the LAN site does not need to know the internal address space of the dialup user. The NetScreen device protecting the LAN uses the predefined address “Dial-Up VPN” as the source address in the incoming policy and the destination in the outgoing policy.

The ability to create bidirectional policies for a dialup VPN tunnel allows traffic to originate from the LAN end of the VPN connection after the connection has been established. Note that unlike a dialup dynamic peer VPN tunnel, this feature requires that the services on the incoming and outgoing policies be identical.

Note: *NetScreen does not support service groups and address groups in bidirectional policies referencing a dialup VPN configuration.*

Be mindful that the internal address space of two or more concurrently connected dialup VPN users might overlap. For example, dialup users A and B might both have an internal IP address space of 10.2.2.0/24. If that happens, the NetScreen device sends all outbound VPN traffic to both user A and user B through the VPN referenced in the first policy it finds in the policy list. For example, if the outbound policy referencing the VPN to user A appears first in the policy list, then the NetScreen device sends all outbound VPN traffic intended for users A and B to user A.

Similarly, the internal address of a dialup user might happen to overlap an address in any other policy—whether or not that other policy references a VPN tunnel. If that occurs, the NetScreen device applies the first policy that matches the basic traffic attributes of source address, destination address, source port number, destination port number, service. To avoid a bidirectional dialup VPN policy with a dynamically derived address superseding another policy with a static address, NetScreen recommends positioning the bidirectional dialup VPN policy lower in the policy list.

Example: Bidirectional Dialup VPN Policies

In this example, you configure bidirectional policies for a dialup AutoKey IKE VPN tunnel named *VPN_dial* for IKE user *dialup-j* with IKE ID *jf@ns.com*. For Phase 1 negotiations, you use the proposal *pre-g2-3des-sha*, with the preshared key *Jf11d7uU*. You select the predefined "Compatible" set of proposals for Phase 2 negotiations.

The IKE user initiates a VPN connection to the NetScreen device from the Untrust zone to reach corporate servers in the Trust zone. After the IKE user establishes the VPN connection, traffic can initiate from either end of the tunnel.

The Trust zone interface is ethernet1, has IP address 10.1.1.1/24, and is in NAT mode. The Untrust zone interface is ethernet3 and has IP address 1.1.1.1/24. The default route points to the external router at 1.1.1.250.

WebUI

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

2. Objects

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: trust_net

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.0/24

Zone: Trust

Objects > Users > Local > New: Enter the following, and then click **OK**:

User Name: dialup-j

Status: Enable

IKE User: (select)

Simple Identity: (select); jf@ns.com

3. VPN

VPNs > AutoKey Advanced > Gateway > New: Enter the following, and then click **OK**:

Gateway Name: dialup1

Security Level: Custom

Remote Gateway Type:

Dialup User: (select); dialup-j

Preshared Key: Jf11d7uU

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Gateway configuration page:

Security Level: Custom

Phase 1 Proposal (For Custom Security Level):
pre-g2-3des-sha

Mode (Initiator): Aggressive

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: VPN_dial

Security Level: Compatible

Remote Gateway:

Create a Simple Gateway: (select)

Gateway Name: dialup1

Type:

Dialup User: (select); dialup-j

Preshared Key: Jf11d7uU

Security Level: Compatible

Outgoing Interface: ethernet3

4. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet1

Gateway IP Address: 1.1.1.250

5. Policies

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Dial-Up VPN

Destination Address:

Address Book Entry: (select), trust_net

Service: ANY

Action: Tunnel

VPN Tunnel: VPN_dial

Modify matching bidirectional VPN policy: (select)

CLI

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
```

2. Objects

```
set address trust trust_net 10.1.1.0/24
set user dialup-j ike-id u-fqdn jf@ns.com
```

3. VPN

```
set ike gateway dialup1 dialup dialup-j aggressive outgoing-interface ethernet3
    preshare Jf11d7uU proposal pre-g2-3des-sha
set vpn VPN_dial gateway dialup1 sec-level compatible
```

4. Route

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
```

5. Policies

```
set policy from untrust to trust “Dial-Up VPN” trust_net any tunnel vpn
    VPN_dial
set policy from trust to untrust trust_net “Dial-Up VPN” any tunnel vpn
    VPN_dial
save
```

NetScreen-Remote Security Policy Editor

1. Click **Options > Secure > Specified Connections**.
2. Click **Add a new connection**, and type **Corp** next to the new connection icon that appears.
3. Configure the connection options:

   Connection Security: Secure

   Remote Party Identity and Addressing

   ID Type: IP Subnet

   Subnet: 10.1.1.0

   Mask: 255.255.255.0

   Protocol: All

   Connect using Secure Gateway Tunnel: (select)

   ID Type: IP Address, 1.1.1.1
4. Click the **PLUS** symbol, located to the left of the UNIX icon, to expand the connection policy.
5. Click **My Identity**: Do either of the following:

   Click **Pre-shared Key** > **Enter Key**: Type **Jf11d7uU**, and then click **OK**.

   ID Type: (select **E-mail Address**), and type **jf@ns.com**.
6. Click the **Security Policy** icon, and select **Aggressive Mode**.
7. Click the **PLUS** symbol, located to the left of the Security Policy icon, and then the **PLUS** symbol to the left of Authentication (Phase 1) and Key Exchange (Phase 2) to expand the policy further.

8. Click **Authentication (Phase 1)** > **Proposal 1**: Select the following Encryption and Data Integrity Algorithms:

   Encrypt Alg: Triple DES
   Hash Alg: SHA-1
   Key Group: Diffie-Hellman Group 2

9. Click **Key Exchange (Phase 2)** > **Proposal 1**: Select the following IPSec Protocols:

   Encapsulation Protocol (ESP): (select)
   Encrypt Alg: Triple DES
   Hash Alg: SHA-1
   Encapsulation: Tunnel

10. Click **Key Exchange (Phase 2)** > **Create New Proposal**: Select the following IPSec Protocols:

    Encapsulation Protocol (ESP): (select)
    Encrypt Alg: Triple DES
    Hash Alg: MD5
    Encapsulation: Tunnel

11. Click **Key Exchange (Phase 2)** > **Create New Proposal**: Select the following IPSec Protocols:

    Encapsulation Protocol (ESP): (select)
    Encrypt Alg: DES
    Hash Alg: SHA-1
    Encapsulation: Tunnel

12. Click **Key Exchange (Phase 2)** > **Create New Proposal**: Select the following IPSec Protocols:

    Encapsulation Protocol (ESP): (select)
    Encrypt Alg: DES
    Hash Alg: MD5
    Encapsulation: Tunnel

13. Click **Save**.

GROUP IKE ID

Some organizations have many dialup VPN users. For example, a sales department might have hundreds of users, many of whom require secure dialup communication when off site. With so many users, it is impractical to create a separate user definition, dialup VPN configuration, and policy for each one.

To avoid this difficulty, the Group IKE ID method makes one user definition available for multiple users. The group IKE ID user definition applies to all users having certificates with specified values in the distinguished name (dn) or to all users whose full IKE ID and preshared key on their VPN client match a partial IKE ID and preshared key on the NetScreen device.

Note: *When a dialup IKE user connects to the NetScreen device, the NetScreen device first extracts and uses the full IKE ID to search its peer gateway records in case the user does not belong to a group IKE ID user group. If the full IKE ID search produces no matching entry, the NetScreen device then checks for a partial IKE ID match between the incoming embedded IKE ID and a configured group IKE ID user.*

You add a single group IKE ID user to an IKE dialup VPN user group and specify the maximum number of concurrent connections that that group supports. The maximum number of concurrent sessions cannot exceed the maximum number of allowed Phase 1 SAs or the maximum number of VPN tunnels allowed on the NetScreen platform.

Group IKE ID with Certificates

Group IKE ID with certificates is a technique for performing IKE authentication for a group of dialup IKE users without configuring a separate user profile for each one. Instead, the NetScreen device uses a single group IKE ID user profile that contains a partial IKE ID. A dialup IKE user can successfully build a VPN tunnel to a NetScreen device if the VPN configuration on his VPN client specifies a certificate that contains distinguished name elements that match those configured as the partial IKE ID definition in the group IKE ID user profile on the NetScreen device.

Group IKE ID with Certificates

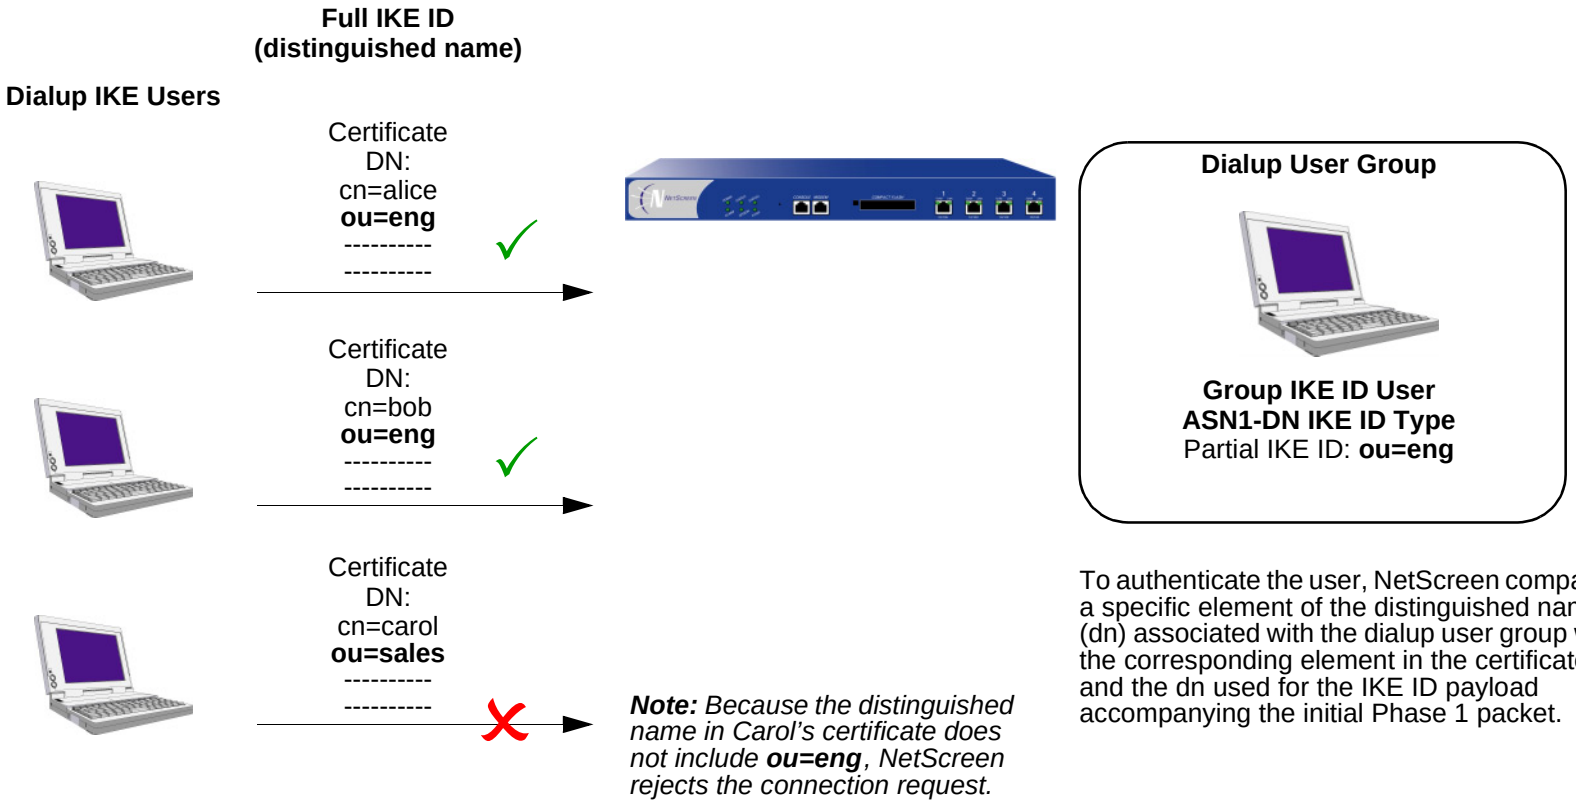

You can set up group IKE ID with certificates as follows:

On the NetScreen Device:

1. Create a new group IKE ID user with a partial IKE identity (such as *ou=sales,o=netscreen*), and specify how many dialup users can use the group IKE ID profile to log on.
2. Assign the new group IKE ID user to a dialup user group[7], and name the group.
3. In the dialup AutoKey IKE VPN configuration, specify the name of the dialup user group, that the Phase 1 negotiations be in Aggressive mode, and that certificates (RSA or DSA, depending on the type of certificate loaded on the dialup VPN clients) be used for authentication.
4. Create a policy permitting inbound traffic via the specified dialup VPN.

On the VPN Client:

1. Obtain and load a certificate whose distinguished name contains the same information as defined in the partial IKE ID on the NetScreen device.
2. Configure a VPN tunnel to the NetScreen device using Aggressive mode for Phase 1 negotiations, specify the certificate that you have previously loaded, and select *Distinguished Name* for the local IKE ID type.

Thereafter, each individual dialup IKE user with a certificate with distinguished name elements that match the partial IKE ID defined in the group IKE ID user profile can successfully build a VPN tunnel to the NetScreen device. For example, if the group IKE ID user has IKE ID *OU=sales,O=netscreen,* the NetScreen device accepts Phase 1 negotiations from any user with a certificate containing those elements in its distinguished name. The maximum number of such dialup IKE users that can connect to the NetScreen device depends upon the maximum number of concurrent sessions that you specify in the group IKE ID user profile.

7. You can put only one group IKE ID user in an IKE user group.

Wildcard and Container ASN1-DN IKE ID Types

When you define the IKE ID for a group IKE user, you must use the Abstract Syntax Notation, version 1, distinguished name (ASN1-DN) as the IKE ID type of identity configuration. This notation is a string of values, which is frequently, though not always, ordered from general to specific. For example:

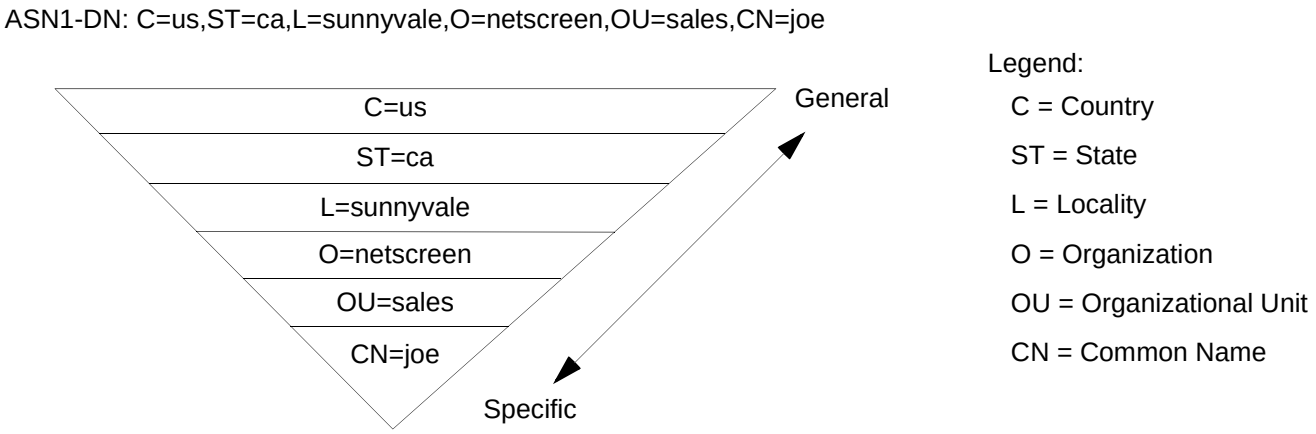

When configuring the group IKE ID user, you must specify the peer's ASN1-DN ID as one of two types:

- **Wildcard:** NetScreen authenticates a dialup IKE user's ID if the values in the dialup IKE user's ASN1-DN identity fields match those in the group IKE user's ASN1-DN identity fields. The wildcard ID type supports only one value per identity field (for example, "ou=eng" or "ou=sw", but not "ou=eng,ou=sw"). The ordering of the identity fields in the two ASN1-DN strings is inconsequential.
- **Container:** NetScreen authenticates a dialup IKE user's ID if the values in the dialup IKE user's ASN1-DN identity fields exactly match the values in the group IKE user's ASN1-DN identity fields. The container ID type supports multiple entries for each identity field (for example, "ou=eng,ou=sw,ou=screenos"). The ordering of the values in the identity fields of the two ASN1-DN strings must be identical.

When configuring an ASN1-DN ID for a remote IKE user, specify the type as either "wildcard" or "container" and define the ASN1-DN ID that you expect to receive in the peer's certificate (for example, "c=us,st=ca,cn=jrogers"). When configuring an ASN1-DN ID for a local IKE ID, use the following keyword: [DistinguishedName]. Include the brackets and spell it exactly as shown.

Wildcard ASN1-DN IKE ID

A wildcard ASN1-DN requires values in the remote peer's distinguished name IKE ID to match values in the group IKE user's partial ASN1-DN IKE ID. The sequencing of these values in the ASN1-DN string is inconsequential. For example, if the dialup IKE user's ID and the group IKE user's ID are as follows

- Dialup IKE user's full ASN1-DN IKE ID: CN=christine,OU=finance,**O=netscreen**,ST=ca,**C=us**
- Group IKE user's partial ASN1-DN IKE ID: **C=us**,**O=netscreen**

then a wildcard ASN1-DN IKE ID successfully matches the two IKE IDs, even though the order of values in the two IDs is different.

The dialup IKE user's ASN1-DN contains the values specified in the group IKE user's ASN1-DN. The order of the values does not matter.

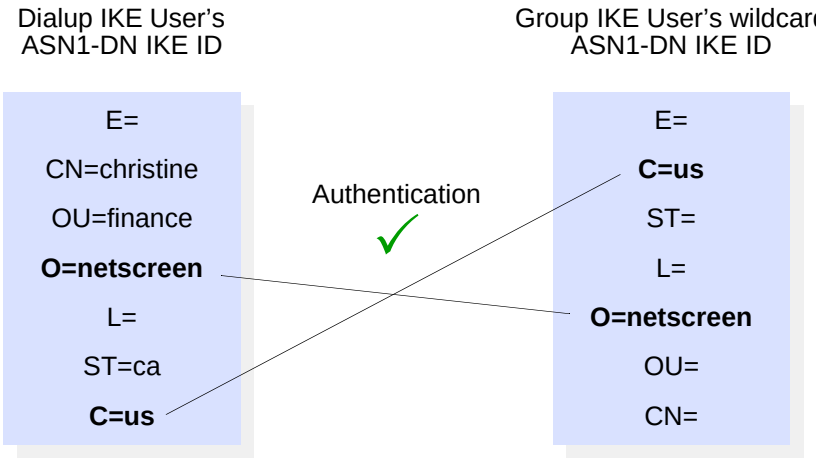

Container ASN1-DN IKE ID

A container ASN1-DN ID allows the group IKE user's ID to have multiple entries in each identity field. NetScreen authenticates a dialup IKE user if the dialup user's ID contains values that exactly match the values in the group IKE user's ID. Unlike the wildcard type, the order of the ASN1-DN fields must be identical in both the dialup IKE user's and group IKE user's IDs and the order of multiple values in those fields must be identical.

The first dialup IKE user's ASN1-DN contains exact matches of the group IKE user's ASN1-DN. The order of the multiple entries in the OU ID field is also identical.

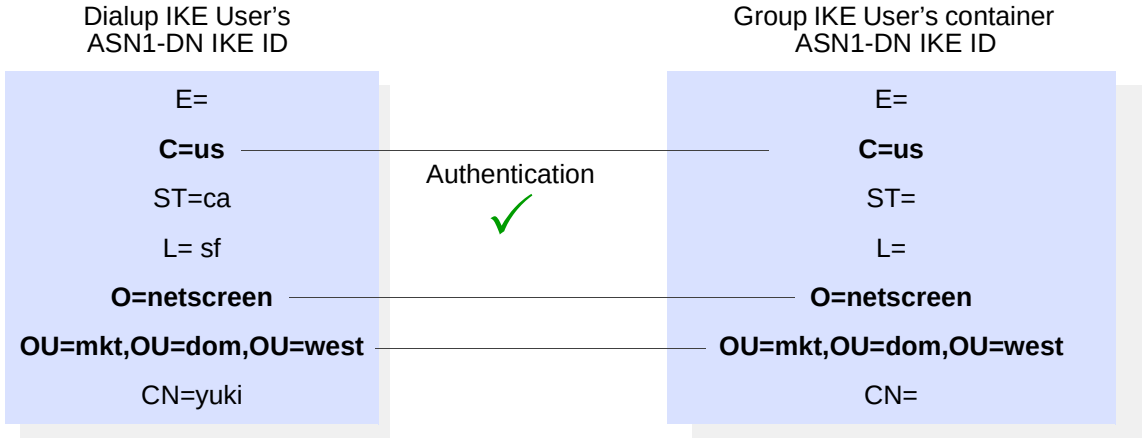

The second dialup IKE user's ASN1-DN contains exact matches of the group IKE user's ASN1-DN. However, the order of the multiple entries in the OU ID field is not identical.

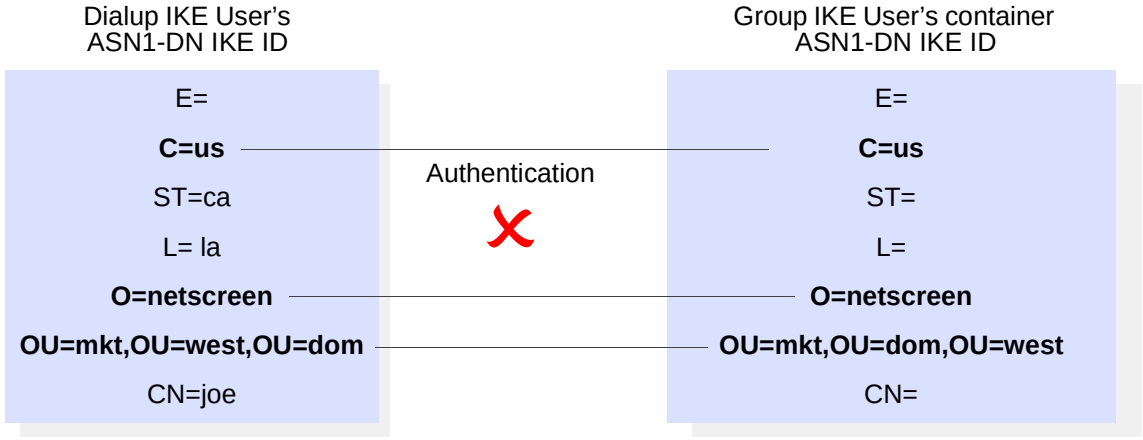

Example: Group IKE ID (Certificates)

In this example, you create a new group IKE ID user definition named *User1*. You configure it to accept up to 10 Phase 1 negotiations concurrently from VPN clients with RSA certificates containing *O=netscreen* and *OU=marketing*. The certificate authority (CA) is Verisign. You name the dialup IKE user group *office_1*.

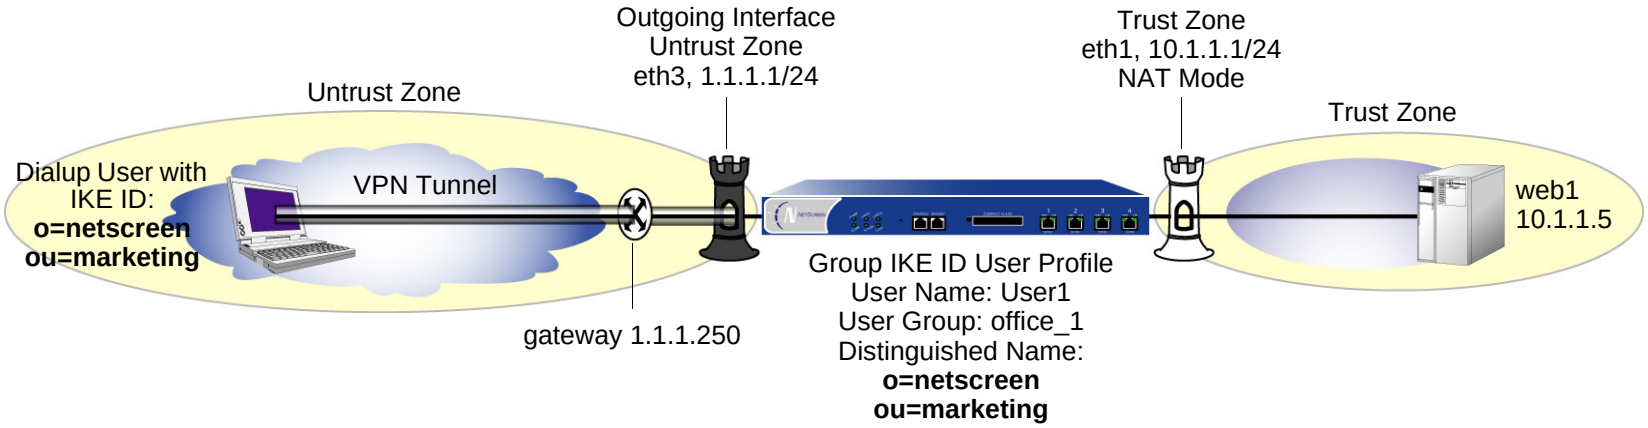

The dialup IKE users send a distinguished name as their IKE ID. The distinguished name (dn) in a certificate for a dialup IKE user in this group might appear as the following concatenated string:

> C=us,ST=ca,L=sunnyvale,**O=netscreen**,**OU=marketing**,CN=michael zhang,CN=a2010002,CN=ns500, CN=4085557800,CN=rsa-key,CN=10.10.5.44

Because the values *O=netscreen* and *OU=marketing* appear in the peer's certificate and the user uses the distinguished name as its IKE ID type, the NetScreen device authenticates the user.

For the Phase 1 and 2 security levels, you specify one Phase 1 proposal—rsa-g2-3des-sha for certificates—and select the predefined "Compatible" set of proposals for Phase 2.

You configure a dialup VPN and a policy permitting HTTP traffic via the VPN tunnel to reach the Web server *Web1*. The configuration of the remote VPN client (using NetScreen-Remote) is also included.

WebUI

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **OK**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

2. Address

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: web1

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.5/32

Zone: Trust

3. Users

Objects > Users > Local > New: Enter the following, then click **OK**:

User Name: User1

Status Enable: (select)

IKE User: (select)

Number of Multiple Logins with same ID: 10

Use Distinguished Name For ID: (select)

OU: marketing

Organization: netscreen

Objects > User Groups > Local > New: Type **office_1** in the Group Name field, do the following, and then click **OK**:

Select **User1** and use the **<<** button to move her from the Available Members column to the Group Members column.

4. VPN

VPNs > AutoKey Advanced > Gateway > New: Enter the following, and then click **OK**:

Gateway Name: Corp_GW

Security Level: Custom

Remote Gateway Type: Dialup User Group: (select), Group: office_1

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Gateway configuration page:

Security Level: Custom

Phase 1 Proposal (For Custom Security Level):
rsa-g2-3des-sha

Mode (Initiator): Aggressive

Preferred Certificate (optional):

Peer CA: Verisign

Peer Type: X509-SIG

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: Corp_VPN

Security Level: Compatible

Remote Gateway: Predefined: (select), Corp_GW

5. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 1.1.1.250

6. Policy

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Dial-Up VPN

Destination Address:

Address Book Entry: (select), web1

Service: HTTP

Action: Tunnel

Tunnel VPN: Corp_VPN

Modify matching bidirectional VPN policy: (clear)

Position at Top: (select)

CLI

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
```

2. Address

```
set address trust web1 10.1.1.5/32
```

3. Users

```
set user User1 ike-id asn1-dn wildcard o=netscreen,ou=marketing share-limit 10
set user-group office_1 user User1
```

4. VPN

```
set ike gateway Corp_GW dialup office_1 aggressive outgoing-interface ethernet3
    proposal rsa-g2-3des-sha
set ike gateway Corp_GW cert peer-ca 1[8]
set ike gateway Corp_GW cert peer-cert-type x509-sig
set vpn Corp_VPN gateway Corp_GW sec-level compatible
```

5. Route

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
```

6. Policy

```
set policy top from untrust to trust “Dial-Up VPN” web1 http tunnel vpn Corp_VPN
save
```

8. The number 1 is the CA ID number. To discover the CA’s ID number, use the following command: **get pki x509 list ca-cert**.

NetScreen-Remote Security Policy Editor

1. Click **Options > Secure > Specified Connections**.
2. Click **Add a new connection**, and type **web1** next to the new connection icon that appears.
3. Configure the connection options:

   Connection Security: Secure

   Remote Party Identity and Addressing

   ID Type: IP Address, 10.1.1.5

   Protocol: Highlight **All**, type **HTTP**, press the **Tab** key, and type **80**.

   Connect using Secure Gateway Tunnel: (select)

   ID Type: IP Address, 1.1.1.1
4. Click the **PLUS** symbol, located to the left of the web1 icon, to expand the connection policy.
5. Click **My Identity**: Select the certificate that has *o=netscreen,ou=marketing* as elements in its distinguished name from the Select Certificate drop-down list[9].

   ID Type: Select **Distinguished Name** from the drop-down list.
6. Click the **Security Policy** icon, and select **Aggressive Mode**.
7. Click the **PLUS** symbol, located to the left of the Security Policy icon, and then the **PLUS** symbol to the left of Authentication (Phase 1) and Key Exchange (Phase 2) to expand the policy further.
8. Click **Authentication (Phase 1)** > **Proposal 1**: Select the following Encryption and Data Integrity Algorithms:

   Authentication Method: RSA Signatures

   Encrypt Alg: Triple DES

   Hash Alg: SHA-1

   Key Group: Diffie-Hellman Group 2

9. This example assumes that you have already loaded a suitable certificate on the NetScreen-Remote client. For information on loading certificates on the NetScreen-Remote, refer to NetScreen-Remote documentation.

9. Click **Key Exchange (Phase 2)** > **Proposal 1**: Select the following IPSec Protocols:
   Encapsulation Protocol (ESP): (select)
   Encrypt Alg: Triple DES
   Hash Alg: SHA-1
   Encapsulation: Tunnel
10. Click **Key Exchange (Phase 2)** > **Create New Proposal**: Select the following IPSec Protocols:
   Encapsulation Protocol (ESP): (select)
   Encrypt Alg: Triple DES
   Hash Alg: MD5
   Encapsulation: Tunnel
11. Click **Key Exchange (Phase 2)** > **Create New Proposal**: Select the following IPSec Protocols:
   Encapsulation Protocol (ESP): (select)
   Encrypt Alg: DES
   Hash Alg: SHA-1
   Encapsulation: Tunnel
12. Click **Key Exchange (Phase 2)** > **Create New Proposal**: Select the following IPSec Protocols:
   Encapsulation Protocol (ESP): (select)
   Encrypt Alg: DES
   Hash Alg: MD5
   Encapsulation: Tunnel
13. Click **Save**.

Group IKE ID with Preshared Keys

Group IKE ID with preshared keys is a technique for performing IKE authentication for a group of dialup IKE users without configuring a separate user profile for each one. Instead, the NetScreen device uses a single group IKE ID user profile, which contains a partial IKE ID. A dialup IKE user can successfully build a VPN tunnel to a NetScreen device if the VPN configuration on his VPN client has the correct preshared key and if the rightmost part of the user's full IKE ID matches the group IKE ID user profile's partial IKE ID.

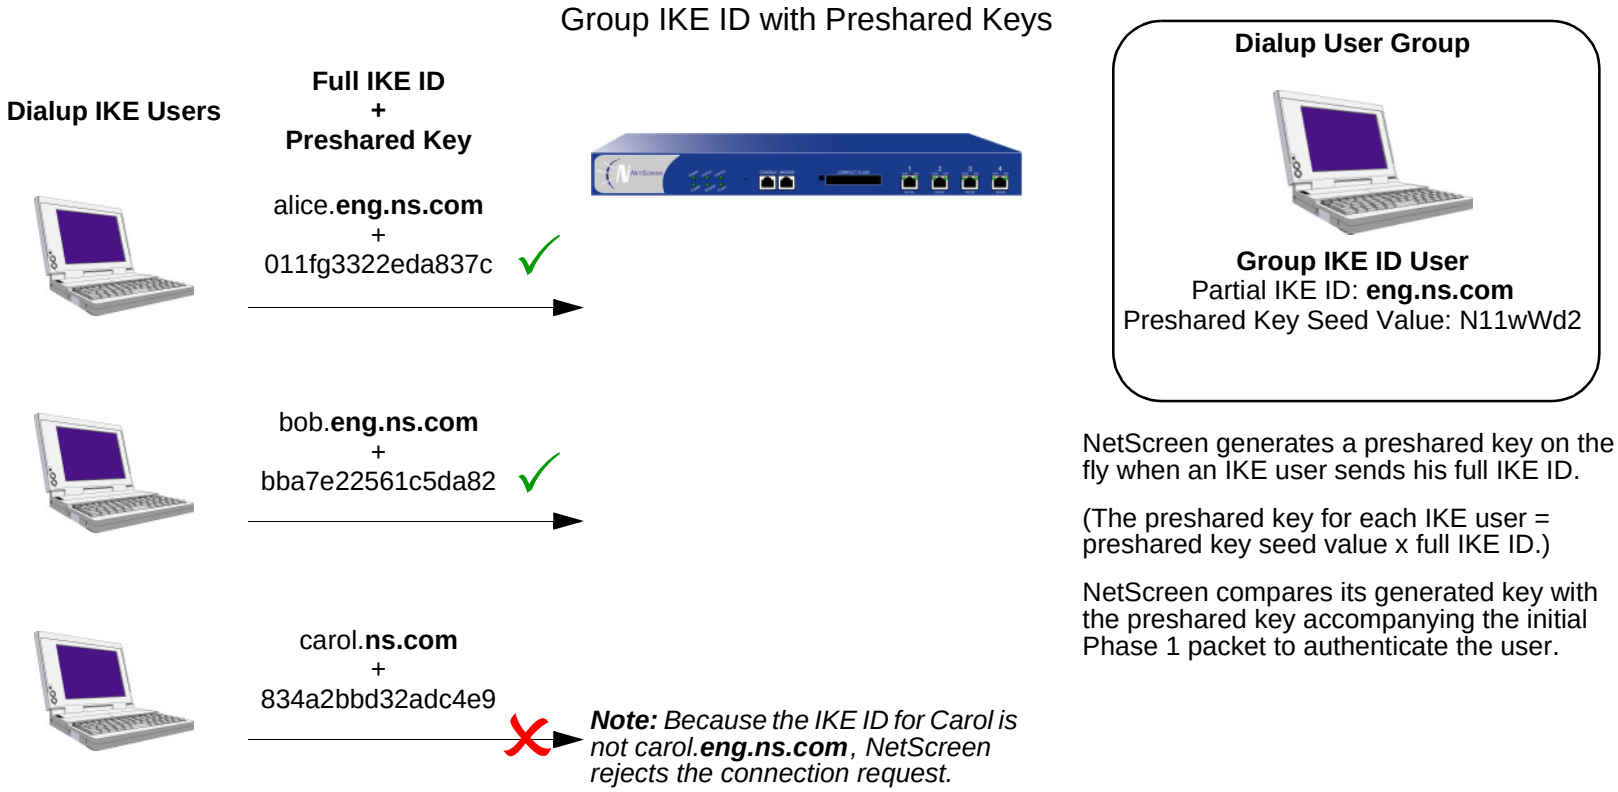

The IKE ID type that you can use for the Group IKE ID with Preshared Key feature can be either an e-mail address or a fully qualified domain name (FQDN).

You can set up group IKE ID with preshared keys as follows:

On the NetScreen Device:

1. Create a new group IKE ID user with a partial IKE identity (such as **netscreen.com**), and specify how many dialup users can use the group IKE ID profile to log on.
2. Assign the new group IKE ID user to a dialup user group.
3. In the dialup AutoKey IKE VPN configuration, assign a name for the remote gateway (such as **road1**), specify the dialup user group, and enter a preshared key seed value.
4. Use the following CLI command to generate an individual dialup user's preshared key using the preshared key seed value and the full user IKE ID (such as **joe@netscreen.com**)

   exec ike preshare-gen *name_str usr_name_str*

   (for example) **exec ike preshare-gen road1 joe@netscreen.com**
5. Record the preshared key for use when configuring the remote VPN client.

On the VPN Client:

Configure a VPN tunnel to the NetScreen device using Aggressive mode for Phase 1 negotiations and enter the preshared key that you previously generated on the NetScreen device.

Thereafter, the NetScreen device can successfully authenticate each individual user whose full IKE ID contains a section that matches the partial group IKE ID user profile. For example, if the group IKE ID user has IKE identity **netscreen.com**, any user with that domain name in his IKE ID can initiate Phase 1 IKE negotiations in Aggressive mode with the NetScreen device. For example: **alice@netscreen.com**, **bob@netscreen.com** and **carol@netscreen.com**. How many such users can log on depends upon a maximum number of concurrent sessions specified in the group IKE ID user profile.

Example: Group IKE ID (Preshared Keys)

In this example, you create a new group IKE ID user named *User2*. You configure it to accept up to 10 Phase 1 negotiations concurrently from VPN clients with preshared keys containing an IKE ID ending with the string *netscreen.com*. The seed value for the preshared key is *jk930k*. You name the dialup IKE user group *office_2*.

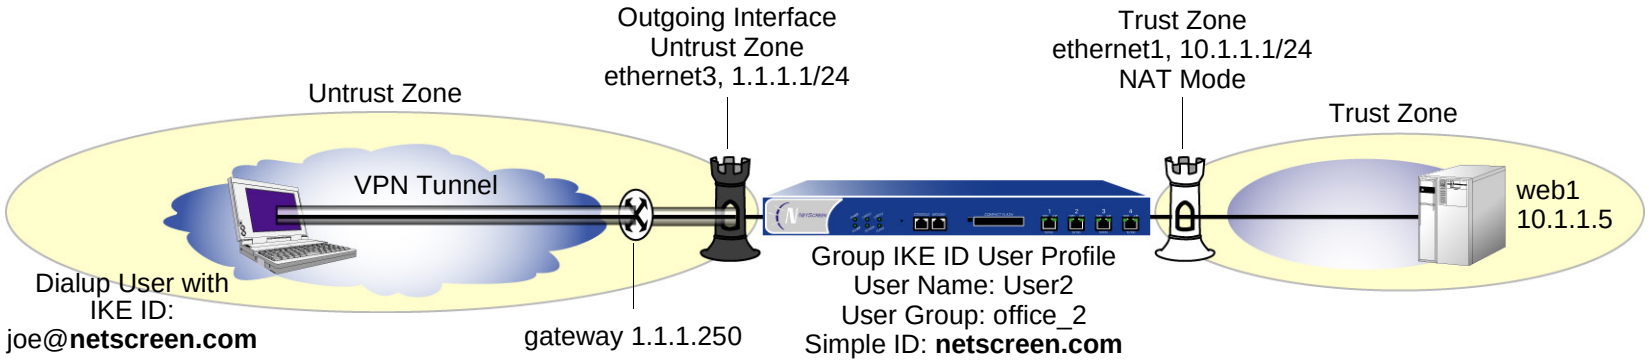

For both the Phase 1 and 2 negotiations, you select the security level predefined as “Compatible”. All the security zones are in the trust-vr routing domain.

WebUI

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **OK**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

2. Address

Objects > Addresses > List > New : Enter the following, and then click **OK**:

Address Name: web1

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.5/32

Zone: Trust

3. Users

Objects > Users > Local > New: Enter the following, then click **OK**:

User Name: User2

Status: Enable

IKE User: (select)

Number of Multiple Logins with same ID: 10

Simple Identity: (select)

IKE Identity: netscreen.com

Objects > User Groups > Local > New: Type **office_2** in the Group Name field, do the following, and then click **OK**:

Select **User2** and use the **<<** button to move him from the Available Members column to the Group Members column.

4. VPN

Note: *The WebUI allows you to enter only a value for a preshared key, not a seed value from which the NetScreen device derives a preshared key. To enter a preshared key seed value when configuring an IKE gateway, you must use the CLI.*

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: Corp_VPN
Security Level: Compatible
Remote Gateway: Predefined: (select), Corp_GW

5. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0
Gateway: (select)
Interface: ethernet3
Gateway IP Address: 1.1.1.250

6. Policy

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:
Address Book Entry: (select), Dial-Up VPN
Destination Address:
Address Book Entry: (select), web1
Service: HTTP
Action: Tunnel
Tunnel VPN: Corp_VPN
Modify matching bidirectional VPN policy: (clear)
Position at Top: (select)

CLI

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
```

2. Address

```
set address trust web1 10.1.1.5/32
```

3. Users

```
set user User2 ike-id u-fqdn netscreen.com share-limit 10
set user-group office_2 user User2
```

4. VPN

```
set ike gateway Corp_GW dialup office_2 aggressive seed-preshare jk930k
    sec-level compatible
set vpn Corp_VPN gateway Corp_GW sec-level compatible
```

5. Route

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
```

6. Policy

```
set policy top from untrust to trust “Dial-Up VPN” web1 http tunnel vpn
    Corp_VPN
save
```

Obtaining the Preshared Key

You can only obtain the preshared key by using the following CLI command:

exec ike preshare-gen *name_str usr_name_str*

The preshared key, based on the preshared key seed value *jk930k* (as specified in the configuration for the remote gateway named *Corp_GW*) and the full identity of individual user *joe@netscreen.com* is *11ccce1d396f8f29ffa93d11257f691af96916f2*.

NetScreen-Remote Security Policy Editor

1. Click **Options > Secure > Specified Connections**.
2. Click **Add a new connection**, and type **web1** next to the new connection icon that appears.
3. Configure the connection options:

   Connection Security: Secure

   Remote Party Identity and Addressing

   ID Type: IP Address, 10.1.1.5

   Protocol: Highlight **All**, type **HTTP**, press the **Tab** key, and type **80**.

   Connect using Secure Gateway Tunnel: (select)

   ID Type: IP Address, 1.1.1.1
4. Click the **PLUS** symbol, located to the left of the web1 icon, to expand the connection policy.
5. Click the **Security Policy** icon, and select **Aggressive Mode**.
6. Click **My Identity**: Click **Pre-shared Key** > **Enter Key**: Type **11ccce1d396f8f29ffa93d11257f691af96916f2**, and then click **OK**.

   ID Type: (select **E-mail Address**), and type **joe@netscreen.com**.
7. Click the **PLUS** symbol, located to the left of the Security Policy icon, and then click the **PLUS** symbol to the left of Authentication (Phase 1) and Key Exchange (Phase 2) to expand the policy further.

8. Click **Authentication (Phase 1)** > **Proposal 1**: Select the following Encryption and Data Integrity Algorithms:

   Authentication Method: Pre-Shared Key

   Encrypt Alg: Triple DES

   Hash Alg: SHA-1

   Key Group: Diffie-Hellman Group 2

9. Click **Authentication (Phase 1)** > **Create New Proposal**: Select the following IPSec Protocols:

   Encapsulation Protocol (ESP): (select)

   Encrypt Alg: Triple DES

   Hash Alg: MD5

   Key Group: Diffie-Hellman Group 2

10. Click **Authentication (Phase 1)** > **Create New Proposal**: Select the following IPSec Protocols:

    Encapsulation Protocol (ESP): (select)

    Encrypt Alg: DES

    Hash Alg: SHA-1

    Key Group: Diffie-Hellman Group 2

11. Click **Authentication (Phase 1)** > **Create New Proposal**: Select the following IPSec Protocols:

    Encapsulation Protocol (ESP): (select)

    Encrypt Alg: DES

    Hash Alg: MD5

    Key Group: Diffie-Hellman Group 2

12. Click **Key Exchange (Phase 2)** > **Proposal 1**: Select the following IPSec Protocols:

    Encapsulation Protocol (ESP): (select)

    Encrypt Alg: Triple DES

Hash Alg: SHA-1

Encapsulation: Tunnel

13. Click **Key Exchange (Phase 2)** > **Create New Proposal**: Select the following IPSec Protocols:

    Encapsulation Protocol (ESP): (select)

    Encrypt Alg: Triple DES

    Hash Alg: MD5

    Encapsulation: Tunnel

14. Click **Key Exchange (Phase 2)** > **Create New Proposal**: Select the following IPSec Protocols:

    Encapsulation Protocol (ESP): (select)

    Encrypt Alg: DES

    Hash Alg: SHA-1

    Encapsulation: Tunnel

15. Click **Key Exchange (Phase 2)** > **Create New Proposal**: Select the following IPSec Protocols:

    Encapsulation Protocol (ESP): (select)

    Encrypt Alg: DES

    Hash Alg: MD5

    Encapsulation: Tunnel

16. Click **Save**.

SHARED IKE IDS

The shared IKE ID feature facilitates the deployment of a large number of dialup users. With this feature, the NetScreen device authenticates multiple dialup VPN users using a single group IKE ID and preshared key. Thus, it provides IPSec protection for large remote user groups through a common VPN configuration.

This feature is similar to the Group IKE ID with pre-shared keys feature, with the following differences:

- With the group IKE ID feature, the IKE ID can be an e-mail address or an FQDN (fully-qualified domain name). For this feature, the IKE ID must be an e-mail address.
- Instead of using the preshared key seed value and the full user IKE ID to generate a preshared key for each user, you specify a single preshared key for all users in the group.
- You must use XAuth to authenticate the individual users.

To set up a shared IKE ID and preshared key on the NetScreen device:

1. Create a new group IKE ID user, and specify how many dialup users can use the group IKE ID to log on. For this feature, use an e-mail address as the IKE ID.
2. Assign the new group IKE ID to a dialup user group.
3. In the dialup-to-LAN autokey IKE VPN configuration, create a shared IKE ID gateway.
4. Define the XAuth users and enable XAuth on the remote IKE gateway.

On the VPN Client:

Configure a VPN tunnel to the NetScreen device using Aggressive mode for Phase 1 negotiations and enter the preshared key that you previously defined on the NetScreen device.Thereafter, the NetScreen device authenticates each remote user as follows:

During Phase 1 negotiations, the NetScreen device first authenticates the VPN client by matching the IKE ID and preshared key that the client sends with the IKE ID and preshared key on the NetScreen device. If there is a match, then the NetScreen device uses XAuth to authenticate the individual user. It sends a login prompt to the user at the remote site between Phase 1 and Phase 2 IKE negotiations. If the remote user successfully logs on with the correct user name and password, Phase 2 negotiations begin.

Example: Shared IKE ID (Preshared Keys)

In this example, you create a new group IKE ID user named Remote_Sales. It accepts up to 250 Phase 1 negotiations concurrently from VPN clients with the same preshared key (abcd1234). You name the dialup IKE user group *R_S*. In addition, you configure two XAuth users, Joe and MIke.

For both the Phase 1 and 2 negotiations, you select the security level predefined as “Compatible”. All the security zones are in the trust-vr routing domain.

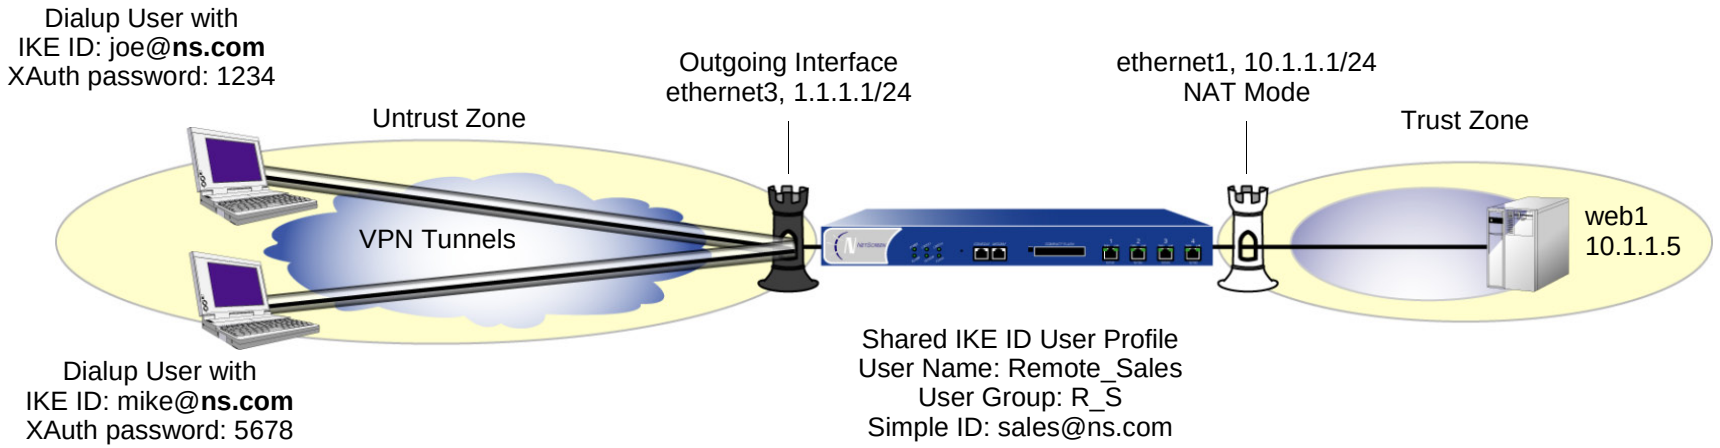

WebUI

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

2. Address

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: web1

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.5/32

Zone: Trust

3. Users

Objects > Users > Local > New: Enter the following, then click **OK**:

User Name: Remote_Sales

Status: Enable

IKE User: (select)

Number of Multiple Logins with same ID: 250

Simple Identity: (select)

IKE Identity: sales@ns.com

Objects > User Groups > Local > New: Type **R_S** in the Group Name field, do the following, and then click **OK**:

Select **Remote_sales** and use the **<<** button to move him from the Available Members column to the Group Members column.

Objects > Users > Local > New: Enter the following, then click **OK**:

User Name: Joe

Status: Enable

XAuth User: (select)

Password: 1234

Confirm Password: 1234

Objects > Users > Local > New: Enter the following, then click **OK**:

User Name: Mike

Status: Enable

XAuth User: (select)

Password: 5678

Confirm Password: 5678

4. VPN

VPNs > AutoKey Advanced > Gateway > New: Enter the following, and then click **OK**:

Gateway Name: sales_gateway

Security Level: Compatible (select)

Remote Gateway Type: Dialup Group (select), R_S

Preshared Key: abcd1234

Outgoing Interface: ethernet3

> Advanced: Enter the following, and then click **Return** to return to the base Gateway configuration page:

Enable XAuth: (select)

Local Authentication: (select)

Allow Any: (select)

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: Sales_VPN

Security Level: Compatible

Remote Gateway: Predefined: (select) sales_gateway

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Bind to: Tunnel Zone, Untrust-Tun

5. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 1.1.1.250

6. Policy

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Dial-Up VPN

Destination Address:

Address Book Entry: (select), web1

Service: HTTP

Action: Tunnel

Tunnel VPN: Sales_VPN

Modify matching bidirectional VPN policy: (clear)

Position at Top: (select)

CLI

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
```

2. Address

```
set address trust web1 10.1.1.5/32
```

3. Users

```
set user Remote_Sales ike-id sales@ns.com share-limit 250
set user-group R_S user Remote_Sales
set user Joe password 1234
set user Joe type xauth
set user Mike password 5678
set user Mike type xauth
```

4. VPN

```
set ike gateway sales_gateway dialup R_S aggressive outgoing-interface
    ethernet3 preshare abcd1234 sec-level compatible
set ike gateway sales_gateway xauth
set vpn sales_vpn gateway sales_gateway sec-level compatible
set vpn sales_vpn bind zone untrust-tun
```

5. Route

```
set route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
```

6. Policy

```
set policy top from untrust to trust “Dial-Up VPN” web1 http tunnel vpn sales_vpn
save
```

NetScreen-Remote Security Policy Editor

This example shows the configuration for the user named Joe.

1. Click **Options > Secure > Specified Connections**.
2. Click **Add a new connection**, and type **web1** next to the new connection icon that appears.
3. Configure the connection options:

   Connection Security: Secure

   Remote Party ID Type: IP Address

   IP Address: 10.1.1.5

   Connect using Secure Gateway Tunnel: (select)

   ID Type: IP Address; 1.1.1.1
4. Click the **PLUS** symbol, located to the left of the web1 icon, to expand the connection policy.
5. Click the **Security Policy** icon, and select **Aggressive Mode**.
6. Click **My Identity**: Click **Pre-shared Key** > **Enter Key**: Type **abcd1234**, and then click **OK**.

   ID Type: (select **E-mail Address**), and type **sales@ns.com**.
7. Click the **PLUS** symbol, located to the left of the Security Policy icon, and then click the **PLUS** symbol to the left of Authentication (Phase 1) and Key Exchange (Phase 2) to expand the policy further.
8. Click **Authentication (Phase 1)** > **Proposal 1**: Select the following Encryption and Data Integrity Algorithms:

   Authentication Method: Pre-Shared Key; Extended Authentication

   Encrypt Alg: Triple DES

   Hash Alg: SHA-1

   Key Group: Diffie-Hellman Group 2

9. Click **Authentication (Phase 1)** > **Create New Proposal**: Select the following IPSec Protocols:
   Encapsulation Protocol (ESP): (select)
   Encrypt Alg: Triple DES
   Hash Alg: MD5
   Key Group: Diffie-Hellman Group 2
10. Click **Authentication (Phase 1)** > **Create New Proposal**: Select the following IPSec Protocols:
   Encapsulation Protocol (ESP): (select)
   Encrypt Alg: DES
   Hash Alg: SHA-1
   Key Group: Diffie-Hellman Group 2
11. Click **Authentication (Phase 1)** > **Create New Proposal**: Select the following IPSec Protocols:
   Encapsulation Protocol (ESP): (select)
   Encrypt Alg: DES
   Hash Alg: MD5
   Key Group: Diffie-Hellman Group 2
12. Click **Key Exchange (Phase 2)** > **Proposal 1**: Select the following IPSec Protocols:
   Encapsulation Protocol (ESP): (select)
   Encrypt Alg: Triple DES
   Hash Alg: SHA-1
   Encapsulation: Tunnel
13. Click **Key Exchange (Phase 2)** > **Create New Proposal**: Select the following IPSec Protocols:
   Encapsulation Protocol (ESP): (select)
   Encrypt Alg: Triple DES
   Hash Alg: MD5
   Encapsulation: Tunnel

14. Click **Key Exchange (Phase 2)** > **Create New Proposal**: Select the following IPSec Protocols:
    Encapsulation Protocol (ESP): (select)
    Encrypt Alg: DES
    Hash Alg: SHA-1
    Encapsulation: Tunnel
15. Click **Key Exchange (Phase 2)** > **Create New Proposal**: Select the following IPSec Protocols:
    Encapsulation Protocol (ESP): (select)
    Encrypt Alg: DES
    Hash Alg: MD5
    Encapsulation: Tunnel
16. Click **Save**.

6

L2TP

This chapter provides an introduction to Layer 2 Tunneling Protocol (L2TP), its use alone and with IPSec support, and then some configuration examples for L2TP and L2TP-over-IPSec:

- "Introduction to L2TP" on page 270
- "Packet Encapsulation and Decapsulation" on page 274
- "L2TP Parameters" on page 276
  - "Example: Configuring an IP Pool and L2TP Default Settings" on page 277
- "L2TP and L2TP-over-IPSec" on page 279
  - "Example: Configuring L2TP" on page 280
  - "Example: Configuring L2TP-over-IPSec" on page 286

INTRODUCTION TO L2TP

Layer 2 Tunneling Protocol (L2TP) provides a way for a dial-up user to make a virtual Point-to-Point Protocol (PPP) connection to an L2TP network server (LNS), which can be a NetScreen device. L2TP sends PPP frames through a tunnel between an L2TP access concentrator (LAC) and the LNS.

Originally, L2TP was designed so that a LAC residing at an ISP site tunneled to an LNS at either another ISP or corporate site. The L2TP tunnel did not extend completely to the dial-up user's computer, but only to the LAC at the dial-up user's local ISP. (This is sometimes referred to as a compulsory L2TP configuration.)

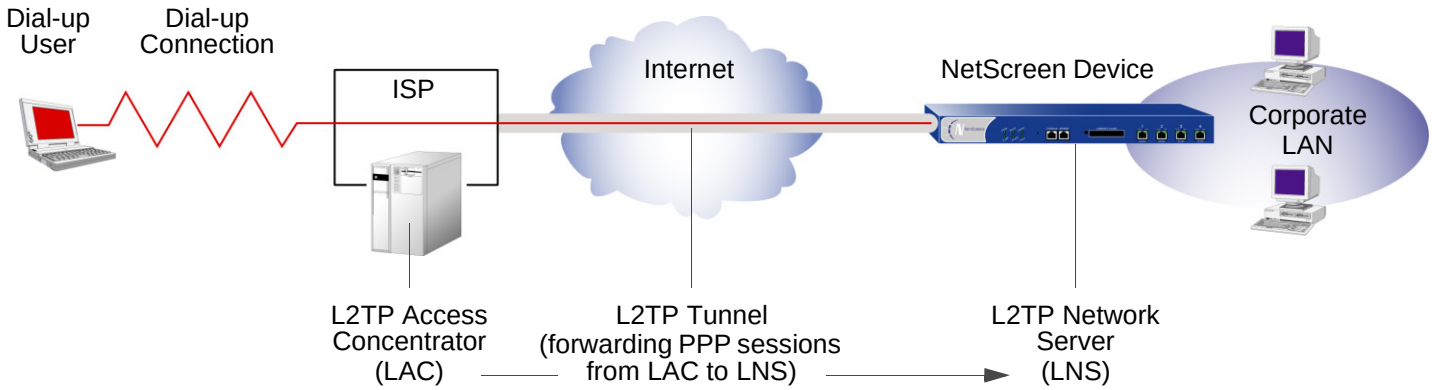

With the capability of a NetScreen-Remote client on Windows 2000 or Windows NT, or a Windows 2000 client by itself, to act as a LAC, the L2TP tunnel can extend directly to the dial-up user's computer, thus providing end-to-end tunneling. (This approach is sometimes referred to as a voluntary L2TP configuration.)

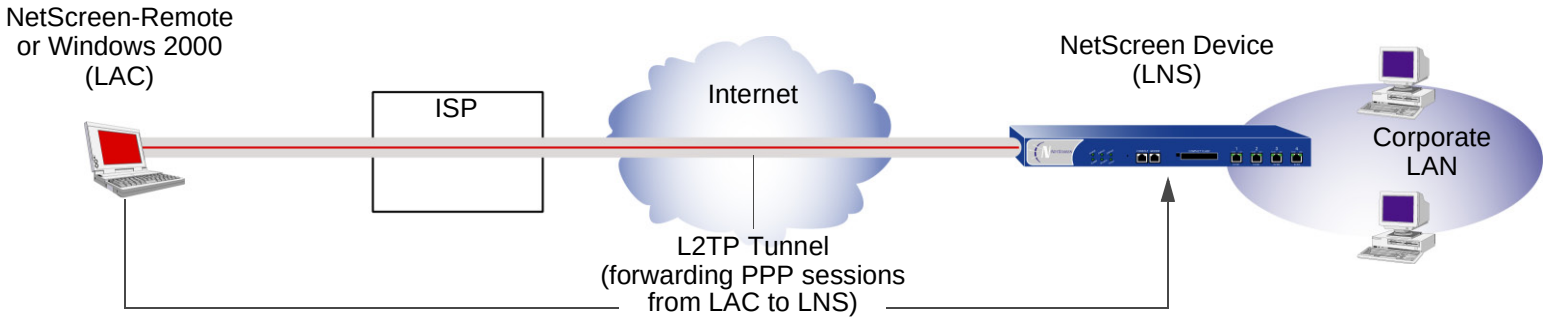

Because the PPP link extends from the dial-up user across the Internet to the NetScreen device (LNS), it is the NetScreen device, not the ISP, that assigns the client its IP address, DNS and WINS servers addresses, and authenticates the user, either from the local database or from an external auth server (RADIUS, SecurID, or LDAP).

In fact, the client receives two IP addresses—one for its physical connection to the ISP, and a logical one from the LNS. When the client contacts its ISP, perhaps using PPP, the ISP makes IP and DNS assignments, and authenticates the client. This allows users to connect to the Internet with a public IP address, which becomes the outer IP address of the L2TP tunnel.

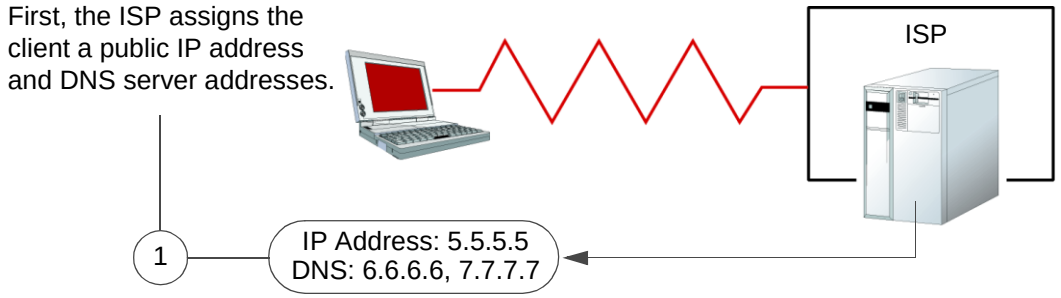

Then, when the L2TP tunnel forwards the encapsulated PPP frames to the NetScreen device, the NetScreen device assigns the client an IP address, and DNS and WINS settings. The IP address can be from the set of private addresses not used on the Internet. This address becomes the inner IP address of the L2TP tunnel.

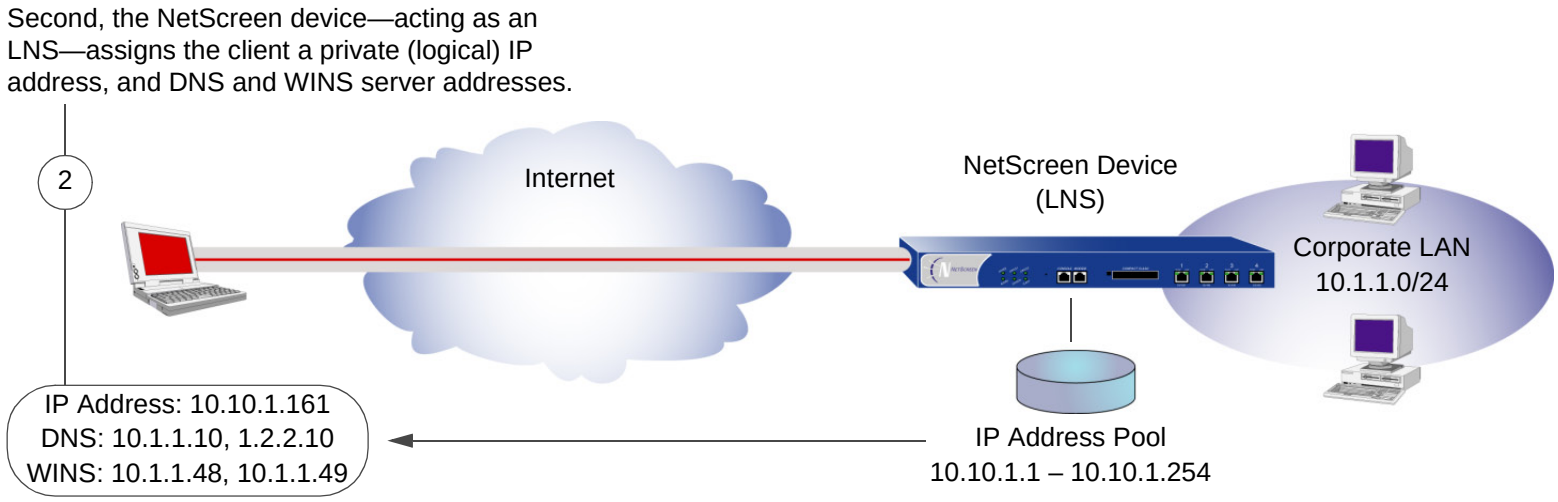

Note: *The IP addresses assigned to the L2TP client must be in a different subnet from the IP addresses in the corporate LAN.*

The current version of ScreenOS provides the following L2TP support:

- L2TP tunnels originating from a host running Windows 2000[1]
- A combination of L2TP and IPSec in transport mode (L2TP-over-IPSec)
  - For NetScreen-Remote: L2TP-over-IPSec with Main mode negotiations using certificates, and Aggressive mode using either a preshared key or certificates
  - For Windows 2000: L2TP-over-IPSec with Main mode negotiations using certificates
- User authentication using either the Password Authentication Protocol (PAP) and Challenge Handshake Authentication Protocol (CHAP) from the local database or an external auth server (RADIUS, SecurID, or LDAP)

  Note: *The local database and RADIUS servers support both PAP and CHAP. SecurID and LDAP servers support PAP only.*

- The assignment of dialup users' IP address, Domain Name System (DNS) servers, and Windows Internet Naming Service (WINS) servers from either the local database or a RADIUS server
- L2TP tunnels and L2TP-over-IPSec tunnels for the root system and virtual systems

  Note: *To use L2TP, the NetScreen device must be operating at Layer 3, with security zone interfaces in NAT or Route mode. When the NetScreen device is operating at Layer 2, with security zone interfaces in Transparent mode, no L2TP-related material appears in the WebUI, and L2TP-related CLI commands elicit error messages.*

1. By default, Windows 2000 performs L2TP-over-IPSec. To force it to use L2TP only, you must navigate to the ProhibitIPSec key in the registry and change **0** (L2TP-over-IPSec) to **1** (L2TP only). (Before performing this, NetScreen recommends that you backup your registry.) Click **Start** > **Run**: Type **regedit**. Double-click **HKEY_LOCAL_MACHINE** > **System** > **CurrentControlSet** > **Services** > **RasMan** > **Parameters**. Double-click **ProhibitIPSec**: Type **1** in the Value data field, select **Hexadecimal** as the base value, and then click **OK**. Reboot. (If you do not find such an entry in the registry, see Microsoft WIndows documentation for information on how to create one.)

PACKET ENCAPSULATION AND DECAPSULATION

L2TP employs encapsulation of packets as the means for transporting PPP frames from the LAC to the LNS. Before looking at specific examples for setting up L2TP and L2TP-over-IPSec, an overview of the encapsulation and decapsulation involved in the L2TP process is presented.

Encapsulation

When a dialup user on an IP network sends data over an L2TP tunnel, the LAC encapsulates the IP packet within a series of layer 2 frames, layer 3 packets, and layer 4 segments. Assuming that the dialup user connects to the local ISP over a PPP link, the encapsulation proceeds as follows:

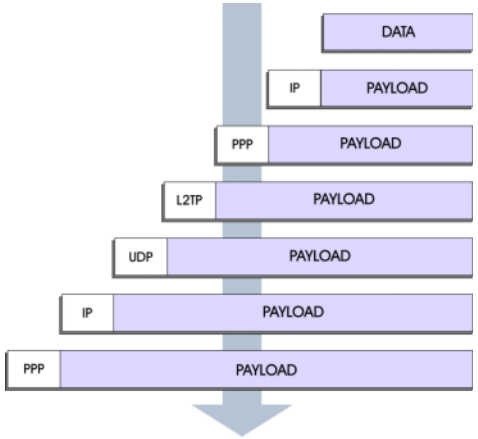

1. The data is placed in an IP payload.
2. The IP packet is encapsulated in a PPP frame.
3. The PPP frame is encapsulated in an L2TP frame.
4. The L2TP frame is encapsulated in a UDP segment.
5. The UDP segment is encapsulated in an IP packet.
6. The IP packet is encapsulated in a PPP frame to make the physical connection between the dialup user and the ISP.

Decapsulation

When the LAC initiates the PPP link to the ISP, the decapsulation and forwarding of the nested contents proceed as follows:

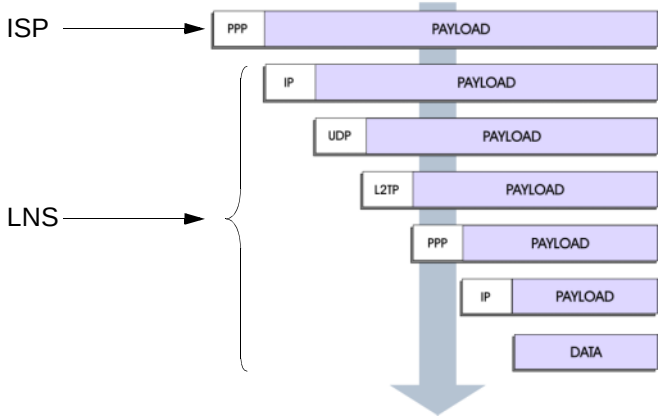

1. The ISP completes the PPP link and assigns the user's computer an IP address.
   Inside the PPP payload is an IP packet.
2. The ISP removes the PPP header and forwards the IP packet to the LNS.
3. The LNS removes the IP header.
   Inside the IP payload is a UDP segment specifying port 1701, the port number reserved for L2TP.
4. The LNS removes the UDP header.
   Inside the UDP payload is an L2TP frame.
5. The LNS processes the L2TP frame, using the tunnel ID and call ID in the L2TP header to identify the specific L2TP tunnel. The LNS then removes the L2TP header.
   Inside the L2TP payload is a PPP frame.
6. The LNS processes the PPP frame, assigning the user's computer a logical IP address.
   Inside the PPP payload is an IP packet.
7. The LNS routes the IP packet to its ultimate destination, where the IP header is removed and the data in the IP packet is extracted.

L2TP PARAMETERS

The LNS uses L2TP to provide the PPP settings for a dial-up user that typically come from an ISP. These settings are as follows:

- IP address – The NetScreen device selects an address from a pool of IP addresses and assigns it to the dial-up user's computer. The selection process operates cyclically through the IP address pool; that is, in a pool from 10.10.1.1 to 10.10.1.3, the addresses are selected in the following cycle: 10.10.1.1 – 10.10.1.2 – 10.10.1.3 – 10.10.1.1 – 10.10.1.2 …
- DNS primary and secondary server IP addresses – The NetScreen device provides these addresses for the dial-up user's computer to use.
- WINS primary and secondary server IP addresses – The NetScreen device also provides these addresses for the dial-up user's computer to use.

The LNS also authenticates the user through a user name and password. You can enter the user in the local database or in an external auth server (RADIUS, SecurID, or LDAP).

Note: *The RADIUS or SecurID server that you use for authenticating L2TP users can be the same server you use for network users, or it can be a different server.*

In addition, you can specify one of the following schemes for the PPP authentication:

- Challenge Handshake Authentication Protocol (CHAP), in which the NetScreen device sends a challenge (encryption key) to the dial-up user after he or she makes a PPP link request, and the user encrypts his or her login name and password with the key. The local database and RADIUS servers support CHAP.
- Password Authentication Protocol (PAP), which sends the dial-up user's password in the clear along with the PPP link request. The local database and RADIUS, SecurID, and LDAP servers support PAP.
- "ANY", meaning that the NetScreen device negotiates CHAP, and then if that fails, PAP.

You can apply to dial-up users and dialup user groups the default L2TP parameters that you configure on the L2TP Default Configuration page (VPNs > L2TP > Default Settings) or with the **set l2tp default** command. You can also apply L2TP parameters that you configure specifically for L2TP users on the User Configuration page (Users > Users > Local > New) or with the **set user** *name_str* **remote-settings** command. The user-specific L2TP settings supersede the default L2TP settings.

Example: Configuring an IP Pool and L2TP Default Settings

In this example, you define an IP address pool with addresses ranging from 10.1.3.40 to 10.1.3.100. You specify DNS server IP addresses 1.1.1.2 (primary) and 1.1.1.3 (secondary). The NetScreen device performs PPP authentication using CHAP.

Note: *You specify the auth server on a per-L2TP tunnel basis.*

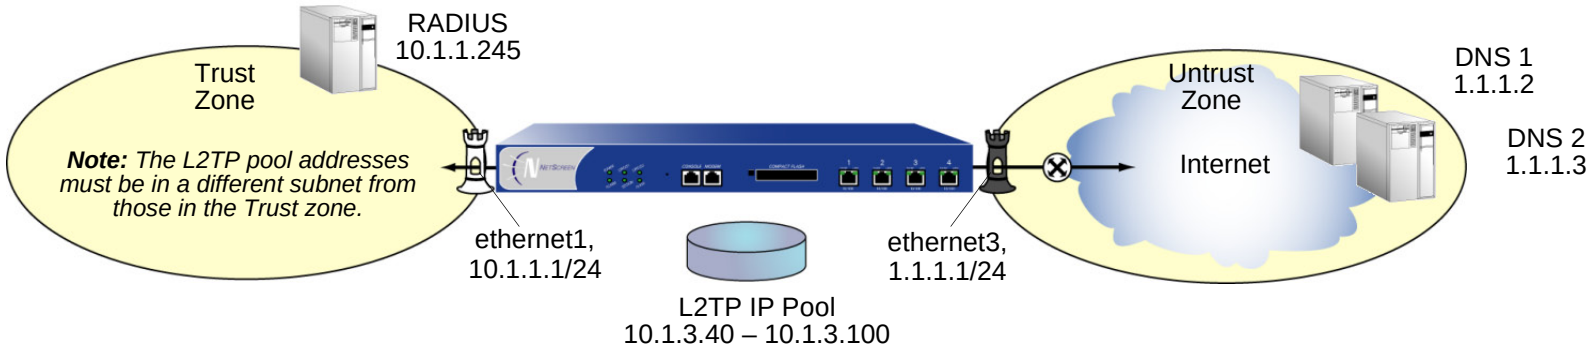

WebUI

1. IP Pool

   Objects > IP Pools > New: Enter the following, and then click **OK**:

   IP Pool Name: Sutro

   Start IP: 10.1.3.40

   End IP: 10.1.3.100

2. Default L2TP Settings

VPNs > L2TP > Default Settings: Enter the following, and then click **Apply**:

IP Pool Name: Sutro

PPP Authentication: CHAP

DNS Primary Server IP: 1.1.1.2

DNS Secondary Server IP: 1.1.1.3

WINS Primary Server IP: 0.0.0.0

WINS Secondary Server IP: 0.0.0.0

CLI

1. IP Pool

```
set ippool sutro 10.1.3.40 10.1.3.100
```

2. Default L2TP Settings

```
set l2tp default ippool sutro
set l2tp default ppp-auth chap
set l2tp default dns1 1.1.1.2
set l2tp default dns2 1.1.1.3
save
```

L2TP AND L2TP-OVER-IPSEC

Although the dial-up user can be authenticated using CHAP or PAP, an L2TP tunnel is not encrypted, and therefore is not a true VPN tunnel. The purpose of L2TP is simply to permit the administrator of the local NetScreen device to assign IP addresses to remote dial-up users. These addresses can then be referenced in policies.

To encrypt an L2TP tunnel, you need to apply an encryption scheme to the L2TP tunnel. Because L2TP assumes that the network between the LAC and the LNS is IP, you can employ IPSec to provide encryption. This combination is called L2TP-over-IPSec. L2TP-over-IPSec requires setting up both an L2TP tunnel and an IPSec tunnel with the same endpoints, and then linking them together in a policy. L2TP-over-IPSec requires that the IPSec tunnel be in transport mode so that the tunnel endpoint addresses remain in the clear. (For information about transport mode and tunnel mode, see “Modes” on page 4.)

You can create an L2TP tunnel between a NetScreen device and a host running Windows 2000 if you change the Windows 2000 registry settings. (For instructions on how to change the registry, see the footnote on page 273.)

You can create an L2TP-over-IPSec tunnel between a NetScreen device and either of the following VPN clients:

- A host running NetScreen-Remote on a Windows 2000 or Windows NT operating system
- A host running Windows 2000 (without NetScreen-Remote)

Example: Configuring L2TP

In this example, you create a dialup user group called “fs” (for “field-sales”) and configure an L2TP tunnel called “sales_corp,” using ethernet3 (Untrust zone) as the outgoing interface for the L2TP tunnel. The NetScreen device applies the following default L2TP tunnel settings to the dialup user group:

- The L2TP users are authenticated via the local database.
- PPP authentication uses CHAP.
- The range of addresses in the IP pool (named “global”) is from 10.10.2.100 to 10.10.2.180[2].
- The DNS servers are 1.1.1.2 (primary) and 1.1.1.3 (secondary)

Note: *An L2TP-only configuration is not secure. It is recommended only for debugging purposes.*

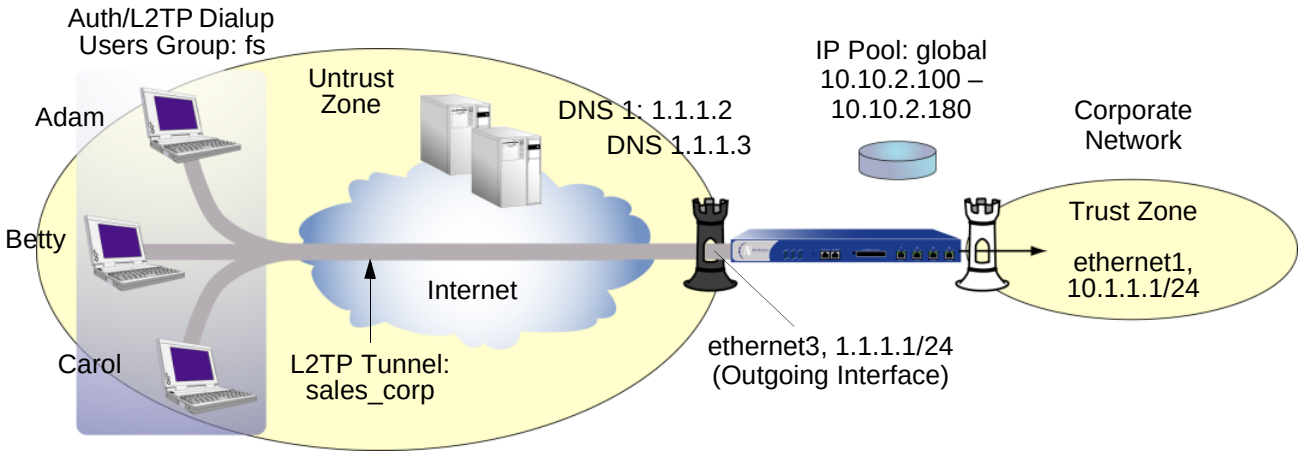

The remote L2TP clients are on Windows 2000 operating systems. For information on how to configure L2TP on the remote clients, refer to Windows 2000 documentation. Only the configuration for the NetScreen device end of the L2TP tunnel is provided below.

2. The addresses in the L2TP IP pool must be in a different subnet than the addresses in the corporate network.

WebUI

1. L2TP Users

Objects > Users > Local > New: Enter the following, and then click **OK**:

User Name: Adam

Status: Enable

L2TP User: (select)

User Password: AJbioJ15

Confirm Password: AJbioJ15

Objects > Users > Local > New: Enter the following, and then click **OK**:

User Name: Betty

Status: Enable

L2TP User: (select)

User Password: BviPsoJ1

Confirm Password: BviPsoJ1

Objects > Users > Local > New: Enter the following, and then click **OK**:

User Name: Carol

Status: Enable

L2TP User: (select)

User Password: Cs10kdD3

Confirm Password: Cs10kdD3

2. L2TP User Group

Objects > User Groups > Local > New: Type **fs** in the Group Name field, do the following, and then click **OK**:

Select **Adam** and use the **<<** button to move him from the Available Members column to the Group Members column.

Select **Betty** and use the **<<** button to move her from the Available Members column to the Group Members column.

Select **Carol** and use the **<<** button to move her from the Available Members column to the Group Members column.

3. Default L2TP Settings

Objects > IP Pools > New: Enter the following, and then click **OK**:

IP Pool Name: global

Start IP: 10.10.2.100

End IP: 10.10.2.180

VPNs > L2TP > Default Settings: Enter the following, and then click **OK**:

IP Pool Name: global

PPP Authentication: CHAP

DNS Primary Server IP: 1.1.1.2

DNS Secondary Server IP: 1.1.1.3

WINS Primary Server IP: 0.0.0.0

WINS Secondary Server IP: 0.0.0.0

4. L2TP Tunnel

VPNs > L2TP > Tunnel > New: Enter the following, and then click **OK**:

Name: sales_corp
Use Custom Settings: (select)
Authentication Server: Local
Dialup Group: Local Dialup Group - fs
Outgoing Interface: ethernet3
Peer IP: 0.0.0.0[3]
Host Name (optional): Enter the name of the computer acting as the LAC[4].
Secret (optional): Enter a secret shared between the LAC and the LNS.

Note: *To add a secret to the LAC for authenticating the L2TP tunnel, you must modify the Windows 2000 registry as follows:*

*(1) Click **Start > Run**, and then type **regedit**. The Registry Editor opens.*
*(2) Click **HKEY_LOCAL_MACHINE**.*
*(3) Right-click **SYSTEM**, and then select **Find** from the pop-up menu that appears.*
*(4) Type **ms_l2tpminiport**, and then click **Find Next**.*
*(5) In the Edit menu, highlight **New**, and then select **String Value**.*
*(6) Type **Password**.*
*(7) Double-click **Password**. The Edit String dialog box appears.*
(8) Type the password in the Value data field. This must be the same as the word in the L2TP Tunnel Configuration Secret field on the NetScreen device.
(9) Reboot the computer running Windows 2000.

When using L2TP-over-IPSec, which is the Windows 2000 default, tunnel authentication is unnecessary; all L2TP messages are encrypted and authenticated inside IPSec.

Keep Alive: 60[5]

3. Because the peer's ISP dynamically assigns it an IP address, enter **0.0.0.0** here.
4. To find the name of a computer running Windows 2000, do the following: Click **Start > Settings > Control Panel > System**. The System Properties dialog box appears. Click the **Network Identification** tab, and see entry following **Full computer name**.

5. Policy

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Dial-Up VPN

Destination Address:

Address Book Entry: (select), Any

NAT: Off

Service: ANY

Action: Tunnel

Tunnel L2TP: sales_corp

Position at Top: (select)

CLI

1. Dialup Users

```
set user adam type l2tp
set user adam password AJbioJ15
unset user adam type auth[6]
set user betty type l2tp
set user betty password BviPsoJ1
unset user betty type auth
set user carol type l2tp
set user carol password Cs10kdD3
unset user carol type auth
```

5. The Keep Alive value is the number of seconds of inactivity before the NetScreen device sends an L2TP hello signal to the LAC.

6. Defining a password for a user automatically classifies the user as an auth user. Therefore, to define the user type strictly as L2TP, you must unset the auth user type.

2. L2TP User Group

```
set user-group fs location local
set user-group fs user adam
set user-group fs user betty
set user-group fs user carol
```

3. Default L2TP Settings

```
set ippool global 10.10.2.100 10.10.2.180
set l2tp default ippool global
set l2tp default auth server Local
set l2tp default ppp-auth chap
set l2tp default dns1 1.1.1.2
set l2tp default dns2 1.1.1.3
```

4. L2TP Tunnel

```
set l2tp sales_corp outgoing-interface ethernet3
set l2tp sales_corp auth server Local user-group fs
```

5. Policy

```
set policy top from untrust to trust “Dial-Up VPN” any any tunnel l2tp
    sales_corp
save
```

Example: Configuring L2TP-over-IPSec

This example uses the same L2TP tunnel created in the previous example (“Example: Configuring L2TP” on page 280). Additionally, you overlay an IPSec tunnel onto the L2TP tunnel to provide encryption. The IPSec tunnel negotiates Phase 1 in Aggressive Mode using a previously loaded RSA certificate, 3DES encryption and SHA-1 authentication. The certificate authority (CA) is Verisign. (For information on obtaining and loading certificates, see Chapter 2, “Public Key Cryptography” on page 15.)The Phase 2 negotiation uses the security level predefined as “Compatible” for Phase 2 proposals. The IPSec tunnel is in transport mode.

The predefined Trust zone and the user-defined Dialup zone are in the trust-vr routing domain. The interfaces for the Dialup and Trust zones are ethernet2 (1.3.3.1/24) and ethernet1 (10.1.1.1/24) respectively. The Trust zone is in NAT mode.

The dialup users Adam, Betty, and Carol use NetScreen-Remote clients on a Windows 2000 operating system[7]. The NetScreen-Remote configuration for dialup user Adam is also included below. (The NetScreen-Remote configuration for the other two dialup users is the same as that for Adam.)

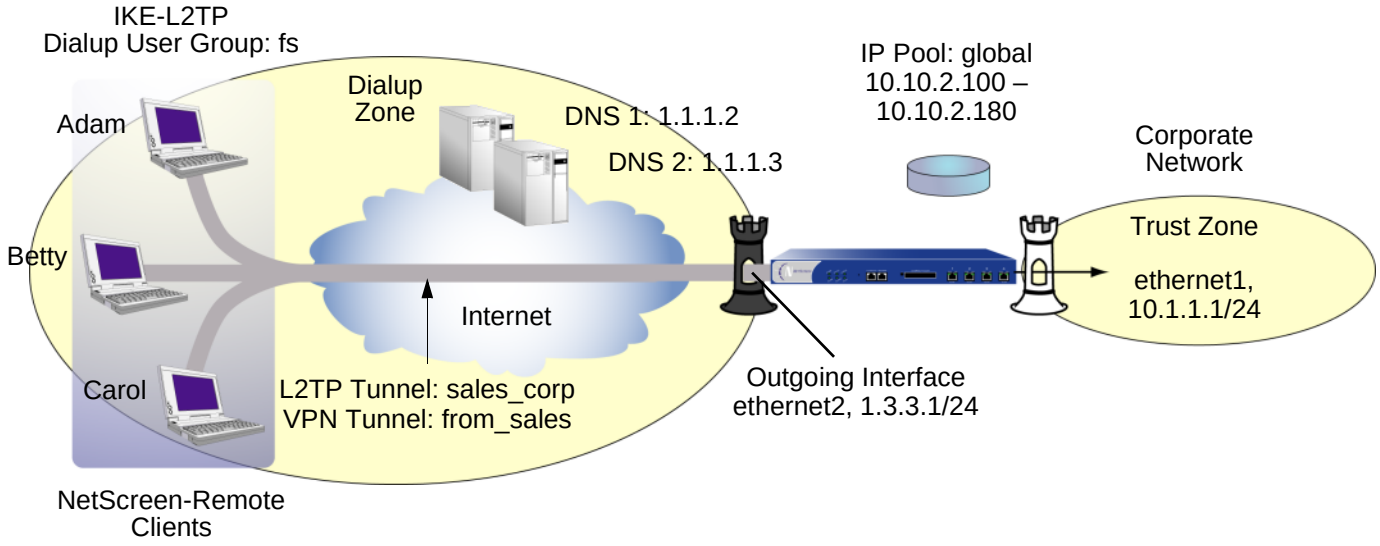

7. To configure an L2TP-over-IPSec tunnel for Windows 2000 (without the NetScreen-Remote), the Phase 1 negotiations must be in Main mode and the IKE ID type must be ASN1-DN.

WebUI

1. User-Defined Zone

Network > Zones > New: Enter the following, and then click **OK**:

Zone Name: Dialup

Virtual Router Name: trust-vr

Zone Type: Layer 3 (select)

Block Intra-Zone Traffic: (select)

TCP/IP Reassembly for ALG: (clear)

Note: *The Trust zone is preconfigured. You do not need to create it.*

2. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet2): Enter the following, and then click **OK**:

Zone Name: Dialup

Static IP: (select this option when present)

IP Address/Netmask: 1.3.3.1/24

3. IKE/L2TP Users

Objects > Users > Local > New: Enter the following, and then click **OK**:

User Name: Adam

Status: Enable

IKE User: (select)

Simple Identity: (select)[8]

IKE Identity: ajackson@abc.com

L2TP User: (select)

User Password: AJbioJ15

Confirm Password: AJbioJ15

Objects > Users > Local > New: Enter the following, and then click **OK**:

User Name: Betty

Status: Enable

IKE User: (select)

Simple Identity: (select)

IKE Identity: bdavis@abc.com

L2TP User: (select)

User Password: BviPsoJ1

Confirm Password: BviPsoJ1

8. The IKE ID that you enter must be the same as the one that the NetScreen-Remote client sends, which is the e-mail address that appears in the certificate that the client uses for authentication.

Objects > Users > Local > New: Enter the following, and then click **OK**:

User Name: Carol

Status: Enable

IKE User: (select)

Simple Identity: (select)

IKE Identity: cburnet@abc.com

L2TP User: (select)

User Password: Cs10kdD3

Confirm Password: Cs10kdD3

4. IKE/L2TP User Group

Objects > User Groups > Local > New: Type **fs** in the Group Name field, do the following, and then click **OK**:

Select **Adam** and use the **<<** button to move him from the Available Members column to the Group Members column.

Select **Betty** and use the **<<** button to move her from the Available Members column to the Group Members column.

Select **Carol** and use the **<<** button to move her from the Available Members column to the Group Members column.

5. IP Pool

Objects > IP Pools > New: Enter the following, and then click **OK**:

IP Pool Name: global

Start IP: 10.10.2.100

End IP: 10.10.2.180

6. Default L2TP Settings

VPNs > L2TP > Default Settings: Enter the following, and then click **Apply**:

IP Pool Name: global

PPP Authentication: CHAP

DNS Primary Server IP: 1.1.1.2

DNS Secondary Server IP: 1.1.1.3

WINS Primary Server IP: 0.0.0.0

WINS Secondary Server IP: 0.0.0.0

7. L2TP Tunnel

VPNs > L2TP > Tunnel > New: Enter the following, and then click **OK**:

Name: sales_corp

Dialup Group: (select), Local Dialup Group - fs

Authentication Server: Local

Outgoing Interface: ethernet2

Peer IP: 0.0.0.0[9]

Host Name (optional): If you want to restrict the L2TP tunnel to a specific host, enter the name of the computer acting as the LAC[10].

Secret (optional): Enter a secret shared between the LAC and the LNS[11]

Note: *The host name and secret settings can usually be ignored. Only advanced users are recommended to use these settings.*

Keep Alive: 60[12]

9. Because the IP address of the peer is dynamic, enter **0.0.0.0** here.
10. To find the name of a computer running Windows 2000, do the following: Click **Start > Settings > Control Panel > System**. The System Properties dialog box appears. Click the **Network Identification** tab, and see entry following **Full computer name**.
11. To add a secret to the LAC for authenticating the L2TP tunnel, you must modify the Windows 2000 registry. See the note in the previous example.
12. The Keep Alive value is the number of seconds of inactivity before the NetScreen device sends an L2TP hello signal to the LAC.

8. VPN Tunnel

VPNs > AutoKey Advanced > Gateway > New: Enter the following, and then click **OK**:

Gateway Name: field

Security Level: Custom

Remote Gateway Type:

Dialup User Group: (select), Group: fs

Outgoing Interface: ethernet2

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Gateway configuration page:

Security Level: User Defined: Custom

Phase 1 Proposal: rsa-g2-3des-sha

Mode (Initiator): Aggressive[13]

Preferred Certificate (Optional):

Peer CA: Verisign

Peer Type: X509-SIG

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

Name: from_sales

Security Level: Compatible

Remote Gateway: Predefined: field

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Security Level: Compatible

Transport Mode: (select)

13. Windows 2000 (without NetScreen-Remote) supports Main mode negotiations only.

9. Policy

Policies > (From: Dialup, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Dial-Up VPN

Destination Address:

Address Book Entry: (select), Any

Service: ANY

Action: Tunnel

Tunnel VPN: from_sales

Modify matching bidirectional VPN policy: (clear)

L2TP: sales_corp

Position at Top: (select)

CLI

1. User-Defined Zone

```
set zone name dialup
set zone dialup vrouter trust-vr
set zone dialup block
```

2. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet2 zone dialup
set interface ethernet2 ip 1.3.3.1/24
```

3. L2TP/IKE Users

```
set user adam type ike l2tp
set user adam password AJbioJ15
unset user adam type auth
set user adam ike-id u-fqdn ajackson@abc.com
set user betty type ike l2tp
set user betty password BviPsoJ1
unset user betty type auth
set user betty ike-id u-fqdn bdavis@abc.com
set user carol type ike l2tp
set user carol password Cs10kdD3
unset user carol type auth
set user carol ike-id u-fqdn cburnet@abc.com
```

4. IKE/L2TP User Group

```
set user-group fs location Local
set user-group fs user adam
set user-group fs user betty
set user-group fs user carol
```

5. IP Pool

```
set ippool global 10.10.2.100 10.10.2.180
```

6. Default L2TP Settings

```
set l2tp default ippool global
set l2tp default ppp-auth chap
set l2tp default dns1 1.1.1.2
set l2tp default dns2 1.1.1.3
```

7. L2TP Tunnel

```
set l2tp sales_corp outgoing-interface ethernet2
set l2tp sales_corp auth server Local user-group fs
```

8. VPN Tunnel

```
set ike gateway field dialup fs aggressive[14] outgoing-interface ethernet2
    proposal rsa-g2-3des-sha
set ike gateway field cert peer-ca1[15]
set ike gateway field cert peer-cert-type x509-sig
set vpn from_sales gateway field transport sec-level compatible
```

9. Policy

```
set policy top from dialup to trust “Dial-Up VPN” any any tunnel vpn from_sales
    l2tp sales_corp
save
```

14. Windows 2000 (without NetScreen-Remote) supports Main mode negotiations only.

15. The number 1 is the CA ID number. To discover the CA’s ID number, use the following command: **get pki x509 list ca-cert**.

NetScreen-Remote Security Policy Editor (Adam[16])

1. Click **Options > Secure > Specified Connections**.
2. Click **Add a new connection**, and type **AJ** next to the new connection icon that appears.
3. Configure the connection options:

   Connection Security: Secure

   Remote Party ID Type: IP Address

   IP Address: 1.3.3.1

   Protocol: UDP

   Port: L2TP

   Connect using Secure Gateway Tunnel: (clear)
4. Click the **PLUS** symbol, located to the left of the AJ icon, to expand the connection policy.
5. Click **My Identity**, and configure the following:

   Select the certificate with the e-mail address specified as the user's IKE ID on the NetScreen device from the Select Certificate drop-down list

   ID Type: E-mail Address[17]

   Port: L2TP
6. Click the **Security Policy** icon, and select **Aggressive Mode**.
7. Click the **PLUS** symbol, located to the left of the Security Policy icon, and then the **PLUS** symbol to the left of Authentication (Phase 1) and Key Exchange (Phase 2) to expand the policy further.

16. To configure L2TP-over-IPSec tunnels for Betty and Carol's NetScreen-Remote clients, follow the same procedure as that provided here for Adam.
17. The e-mail address from the certificate appears in the identifier field automatically.

8. Click **Authentication (Phase 1)** > **Proposal 1**: Select the following Encryption and Data Integrity Algorithms:

   Encrypt Alg: Triple DES

   Hash Alg: SHA-1

   Key Group: Diffie-Hellman Group 2

9. Click **Key Exchange (Phase 2)** > **Proposal 1**: Select the following IPSec Protocols:

   Encapsulation Protocol (ESP): (select)

   Encrypt Alg: Triple DES

   Hash Alg: SHA-1

   Encapsulation: Transport

10. Click **Key Exchange (Phase 2)** > **Create New Proposal**: Select the following IPSec Protocols:

    Encapsulation Protocol (ESP): (select)

    Encrypt Alg: Triple DES

    Hash Alg: MD5

    Encapsulation: Transport

11. Click **Key Exchange (Phase 2)** > **Create New Proposal**: Select the following IPSec Protocols:

    Encapsulation Protocol (ESP): (select)

    Encrypt Alg: DES

    Hash Alg: SHA-1

    Encapsulation: Transport

12. Click **Key Exchange (Phase 2)** > **Create New Proposal**: Select the following IPSec Protocols:

    Encapsulation Protocol (ESP): (select)

    Encrypt Alg: DES

    Hash Alg: MD5

    Encapsulation: Transport

13. Click **Save**.
14. You also need to set up the network connection for your Windows 2000 operating system using the Network Connection Wizard.

Note: *When configuring the Network Connection Wizard, you must enter a destination host name or IP address. Enter 1.3.3.1. Later, when initiating a connection and are prompted for a user name and password, enter adam, AJbioJ15. For more information, consult Microsoft Windows 2000 documentation.*

7

Advanced VPN Features

The material in this chapter covers the following more advanced uses of VPN technology:

- "IPSec NAT Traversal" on page 301
  - "Traversing a NAT Device" on page 302
  - "UDP Checksum" on page 303
  - "The Keepalive Frequency Value" on page 303
  - "IPSec NAT-Traversal and Initiator/Responder Symmetry" on page 304
- "VPN Monitoring" on page 307
  - "Rekey and Optimization Options" on page 307
  - "Source Interface and Destination Address" on page 308
  - "Policy Considerations" on page 310
  - "Configuring the VPN Monitoring Feature" on page 310
  - "Security Consideration for a Route-Based VPN Design" on page 323
  - "SNMP VPN Monitoring Objects and Traps" on page 325
- "Multiple Tunnels per Tunnel Interface" on page 326
  - "Route-to-Tunnel Mapping" on page 327
  - "Remote Peers' Addresses" on page 328
  - "Manual and Automatic Table Entries" on page 330
- "Redundant VPN Gateways" on page 382
  - "VPN Groups" on page 383
  - "Monitoring Mechanisms" on page 384
  - "TCP SYN-Flag Checking" on page 388

- "Back-to-Back VPNs" on page 401
  - "Example: Back-to-Back VPNs" on page 402
- "Hub-and-Spoke VPNs" on page 412
  - "Example: Hub-and-Spoke VPNs" on page 413

IPSEC NAT TRAVERSAL

Network Address Translation (NAT) and Network Address Port Translation (NAPT) are Internet standards that allow a local-area network (LAN) to use one set of IP addresses for internal traffic and a second set of addresses for external traffic. NAT devices generate these external addresses from predetermined pools of IP addresses.

When setting up an IPSec tunnel, the presence of a NAT device along the data path has no effect on Phase 1 and Phase 2 IKE negotiations, which always encapsulate IKE packets within User Datagram Protocol (UDP) packets. However, after the Phase 2 negotiations are completed, performing NAT on the IPSec packets causes the tunnel to fail. Of the many reasons why NAT causes disruption to IPSec[1], one reason is that, for the Encapsulating Security Protocol (ESP), NAT devices cannot discern the location of the Layer 4 header (because it is encrypted) for port translation. For the Authentication Header (AH) protocol, NAT devices can modify the port number, but the authentication check, which includes the entire IPSec packet, fails.

To solve this problem, NetScreen devices (with ScreenOS 3.0.0 or later) and the NetScreen-Remote client (version 6.0 or later) can apply the NAT-traversal (NAT-T) feature. NAT-T adds a layer of UDP encapsulation after detecting one or more NAT devices along the data path during Phase 1 exchanges.

Note: *NetScreen does not support NAT-T for Manual Key tunnels. NetScreen only supports NAT-T for AutoKey IKE tunnels using the Encapsulating Security Protocol (ESP).*

1. For a list of IPSec/NAT incompatibilities, see *draft-ietf-ipsec-nat-regts-00.txt* by Bernard Aboba.

Traversing a NAT Device

In the following illustration, a NAT device at the perimeter of a hotel LAN receives a packet from a VPN dialup client with IP address 200.1.1.1, assigned by its ISP. For all outbound traffic, the NAT device replaces the original source IP address in the outer header with a new address 210.2.2.2. During Phase 1 negotiations, the VPN client and the NetScreen device detect that both VPN participants support NAT-T, that a NAT device is present along the data path, and that it is located in front of the VPN client.

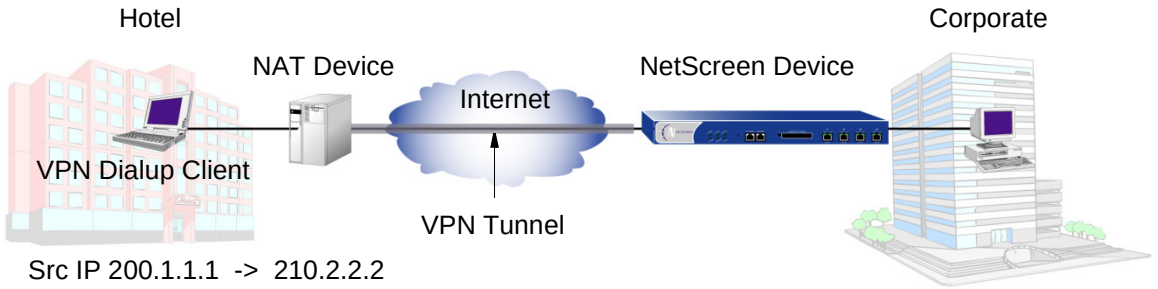

Encapsulating the IPSec packets within UDP packets—which both the VPN client and the NetScreen device do—solves the problem of the authentication check failure. The NAT device processes them as UDP packets, changing the source port in the UDP header and leaving the SPI in the AH or ESP header unmodified. The VPN participants strip off the UDP layer and process the IPSec packets, which pass the authentication check because none of the authenticated content has been changed.

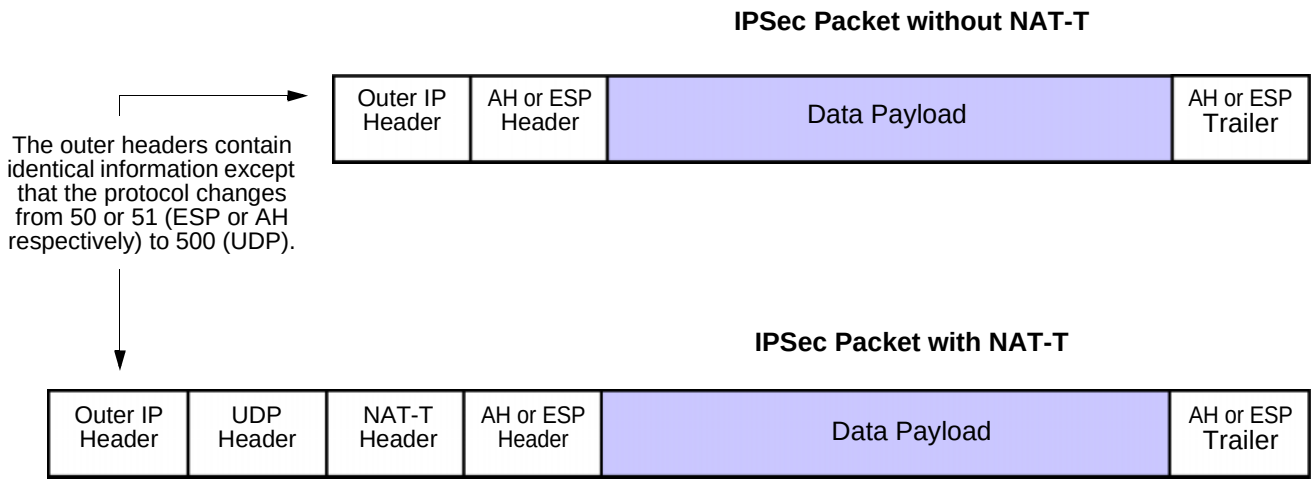

Note: *When NAT-T is enabled, the NetScreen device applies it only when necessary; that is, when it detects a NAT device between the remote host and the NetScreen device.*

UDP Checksum

All UDP packets contain a UDP checksum, a calculated value that ensures UDP packets are free of transmission errors. A NetScreen device does not require use of the UDP checksum for NAT-T, so the WebUI and CLI present the checksum as an optional setting. Even so, some NAT devices require a checksum, so you might have to enable this setting.

The Keepalive Frequency Value

When a NAT device assigns an IP address to a host, the NAT device determines how long the new address remains valid when no traffic occurs. For example, a NAT device might invalidate any generated IP address that remains unused for 20 seconds. Therefore, it is usually necessary for the IPSec participants to send periodic keepalive packets—empty UDP packets—through the NAT device, so that the NAT mapping does not change until the Phase 1 and Phase 2 SAs expire.

Note: *NAT devices have different session timeout intervals, depending on the manufacturer and model. It is important to determine what the interval is for the NAT device, and to set the keepalive frequency value below that.*

IPSec NAT-Traversal and Initiator/Responder Symmetry

When two NetScreen devices establish a tunnel in the absence of a NAT device, either device can serve as initiator or responder. However, if either host resides behind a NAT device, such initiator/responder symmetry might be impossible. This happens whenever the NAT device generates IP addresses dynamically.

Note: *Security zones depicted below are from the perspective of NetScreen B.*

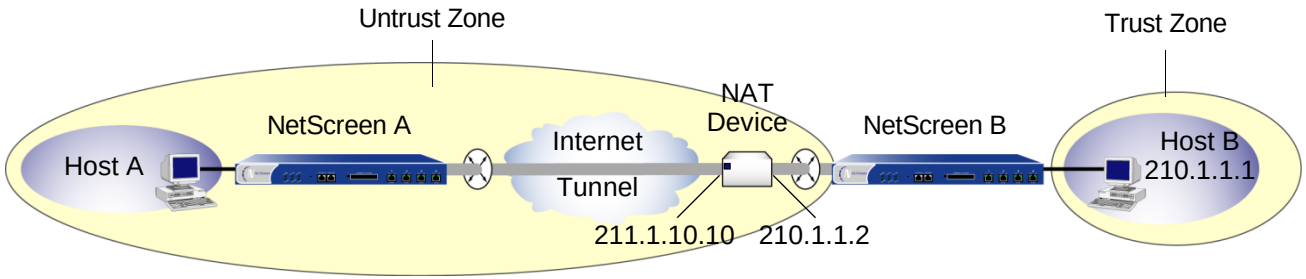

In the above illustration, NetScreen B resides in a subnet located behind a NAT device. If the NAT device generates the new IP address (210.1.1.1) dynamically from a pool of IP addresses, NetScreen A cannot unambiguously identify NetScreen B. Therefore, NetScreen A cannot successfully initiate a tunnel with NetScreen B. NetScreen A must be the responder, NetScreen B must be the initiator, and they must perform Phase 1 negotiations in Aggressive mode.

However, if the NAT device generates the new IP address using a mapped IP (MIP) address, or some other one-to-one addressing method, NetScreen A can unambiguously identify NetScreen B. Consequently, either NetScreen A or NetScreen B can be the initiator, and both can use Main mode or Aggressive mode for Phase 1.

Note: *If you enable NAT-T on a NetScreen device acting as the responder and configure it to perform IKE negotiations in Main mode, then that device and all its peers of the following types that are configured on the same outgoing interface must use the same Phase 1 proposals presented in the same order as each other:*

- *Dynamic peer (peers with dynamically assigned IP addresses)*
- *Dialup VPN users*
- *Peers with static IP addresses behind a NAT device*

Because it is not possible to know the identity of a peer when negotiating Phase 1 in Main mode until the last two messages, the Phase 1 proposals must all be the same so that IKE negotiations can proceed.

The NetScreen device automatically checks that all Phase 1 proposals are the same and in the same order when you configure IKE in Main mode to one of the above peer types on the same outgoing interface. If the proposals are different, the NetScreen device generates an error message.

Example: Enabling NAT-Traversal

In the following example, a NAT device at the perimeter of a hotel LAN assigns an address to the VPN dialup client used by Michael Smith, a salesman attending a convention. For Michael Smith to reach the corporate LAN via a dialup VPN tunnel, you must enable NAT-T for the remote gateway “msmith,” configured on the NetScreen device, and for the remote gateway configured on the VPN dialup client. You also enable the NetScreen device to include a UDP checksum in its transmissions, and you set the keepalive frequency to 8 seconds.

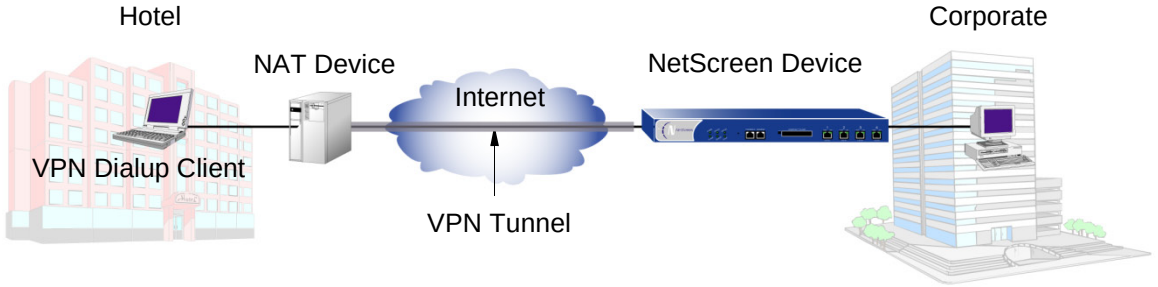

WebUI

VPNs > AutoKey Advanced > Gateway > New: Enter the necessary parameters for the new tunnel gateway as described in Chapter 4, “Site-to-Site VPNs” on page 69 or Chapter 5, “Dialup VPNs” on page 199, enter the following, and then click **OK**:

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Gateway configuration page:

Enable Nat-Traversal: (select)

UDP Checksum: Enable

Keepalive Frequency: 8

Note: *The NetScreen device automatically enables NAT traversal for dial-up VPNs.*

CLI

```
set ike gateway msmith nat-traversal
set ike gateway msmith nat-traversal enable-udp-checksum
set ike gateway msmith nat-traversal keepalive-frequency 8
save
```

VPN MONITORING

When you enable VPN monitoring for a specific tunnel, the NetScreen device sends ICMP echo requests (or "pings") through the tunnel at specified intervals (configurable in seconds) to monitor network connectivity through the tunnel.[2] If the ping activity indicates that the VPN monitoring status has changed, the NetScreen device triggers one of the following Simple Network Management Protocol (SNMP) traps:

- **Up to Down:** This trap occurs when the state of VPN monitoring for the tunnel is up, but a specified consecutive number of ICMP echo requests does not elicit a reply and there is no other incoming VPN traffic.[3] Then the state changes to down.
- **Down to Up:** When the state of VPN monitoring for the tunnel is down, but the ICMP echo request elicits a single response, then the state changes to up. The down-to-up trap occurs only if you have disabled the rekey option and the Phase 2 SA is still active when an ICMP echo request elicits a reply through the tunnel.

Note: *For more information about the SNMP data that VPN monitoring provides, see* *"SNMP VPN Monitoring Objects and Traps" on page 325**.*

Rekey and Optimization Options

If you enable the rekey option, the NetScreen device starts sending ICMP echo requests immediately upon completion of the tunnel configuration and continues to send them indefinitely. The echo requests trigger an attempt to initiate IKE negotiations to establish a VPN tunnel until the state of VPN monitoring for the tunnel is up. The NetScreen device then uses the pings for VPN monitoring purposes. If the state of VPN monitoring for the tunnel changes from up to down, the NetScreen device deactivates its Phase 2 security association (SA) for that peer. The NetScreen device continues to send echo requests to its peer at defined intervals, triggering attempts to reinitiate IKE Phase 2 negotiations—and Phase 1 negotiations, if necessary—until it succeeds. At that point, the NetScreen device reactivates the Phase 2 SA, generates a new key, and reestablishes the tunnel. A message appears in the event log stating that a successful rekey operation has occurred[4].

2. To change the ping interval, you can use the following CLI command: **set vpnmonitor interval** *number*. The default is 10 seconds.
3. To change the threshold for the number of consecutive unsuccessful ICMP echo requests, you can use the following CLI command: **set vpnmonitor threshold** *number*. The default is 10 consecutive requests.
4. If a NetScreen device is a DHCP client, a DHCP update to a different address causes IKE to rekey. However, a DHCP update to the same address does not provoke the IKE rekey operation.

You can use the rekey option to ensure that an AutoKey IKE tunnel is always up, perhaps to monitor devices at the remote site or to allow dynamic routing protocols to learn routes at a remote site and transmit messages through the tunnel. Another use to which you can apply VPN monitoring with the rekey option is for automatic population of the next-hop tunnel binding table (NHTB table) and the route table when multiple VPN tunnels are bound to a single tunnel interface. For an example of this last use, see "Multiple Tunnels per Tunnel Interface" on page 326.

If you disable the rekey option, the NetScreen device performs VPN monitoring only when the tunnel is active with user-generated traffic.

By default, VPN monitoring optimization is disabled. If you enable it (**set vpn** *name* **monitor optimized**), the VPN monitoring behavior changes as follows:

- The NetScreen device considers incoming traffic through the VPN tunnel to be the equivalent of ICMP echo replies. Accepting incoming traffic as a substitute for ICMP echo replies can reduce false alarms that might occur when traffic through the tunnel is heavy and the echo replies do not get through.
- If there is both incoming and outgoing traffic through the VPN tunnel, the NetScreen device suppresses VPN monitoring pings altogether. Doing so can help reduce network traffic.

Although VPN monitoring optimization offers some benefits, be aware that VPN monitoring can no longer provide accurate SNMP statistics, such as VPN network delay time, when the optimization option is active. Also, if you are using VPN monitoring to track the availability of a particular destination IP address at the remote end of a tunnel, the optimization feature can produce misleading results.

Source Interface and Destination Address

By default, the VPN monitoring feature uses the IP address of the local outgoing interface as the source address and the IP address of the remote gateway as the destination address. If the remote peer is a VPN dialup client—such as the NetScreen-Remote—that has an internal IP address, the NetScreen device automatically detects its internal address and uses that as the destination. The VPN client can be an XAuth user with an assigned internal IP address, or a dialup VPN user or a member of a dialup VPN group with an internal IP address. You can also specify the use of other source and destination IP addresses for VPN monitoring—mainly to provide support for VPN monitoring when the other end of a VPN tunnel is not a NetScreen device.

Because VPN monitoring operates independently at the local and remote sites, the source address configured on the device at one end of a tunnel does not have to be the destination address configured on the device at the other end. In fact, you can enable VPN monitoring at both ends of a tunnel or only at one end.

NetScreen-A –> NetScreen-B

NetScreen-A pings from its outgoing interface to the remote gateway; that is, from the Untrust zone interface on NetScreen-A to the Untrust zone interface on NetScreen-B.

(Default Behavior)

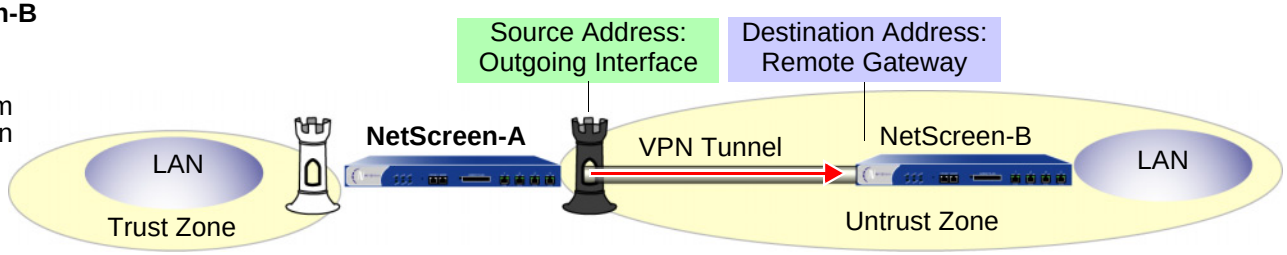

NetScreen-A –> NetScreen-Remote

NetScreen-A pings from its Trust zone interface to the NetScreen-Remote. The NetScreen-Remote requires a policy permitting inbound ICMP traffic from an address beyond the remote gateway; that is, from beyond the Untrust zone interface of NetScreen-A.

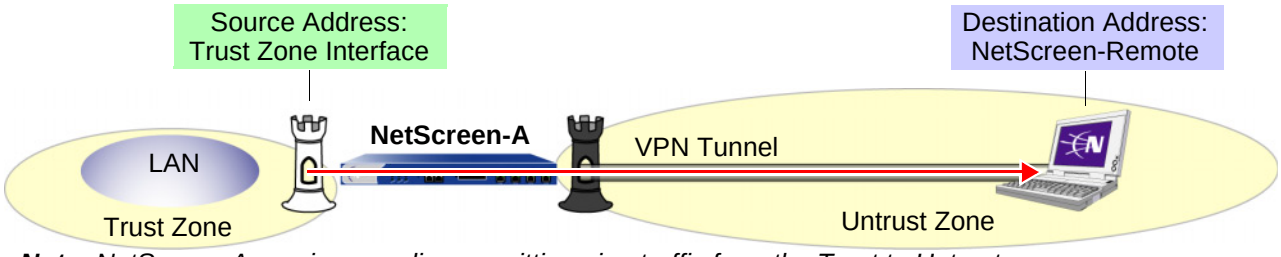

Note: *NetScreen-A requires a policy permitting ping traffic from the Trust to Untrust zones.*

NetScreen-A –> Third-Party VPN Terminator

NetScreen-A pings from its Trust zone interface to a device beyond the remote gateway. This might be necessary if the remote peer does not respond to pings but can support policies permitting inbound ping traffic.

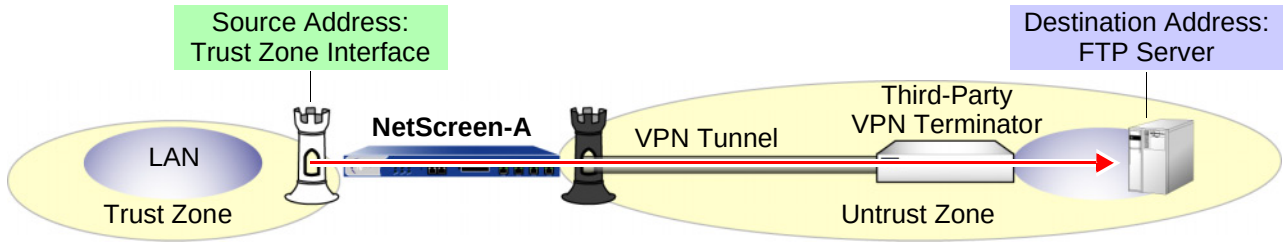

Note: *NetScreen-A requires a policy permitting ping traffic from the Trust to Untrust zones.*

Note: *If the other end of a tunnel is the NetScreen-Remote VPN client that receives its address through XAuth, then the NetScreen device, by default, uses the XAuth-assigned IP address as the destination for VPN monitoring. For information about XAuth, see "XAuth Users and User Groups" on page* ***2*** *-436.*

Policy Considerations

You must create a policy on the sending device to permit pings from the zone containing the source interface to pass through the VPN tunnel to the zone containing the destination address if:

- The source interface is in a different zone from the destination address.
- The source interface is in the same zone as the destination address, and intrazone blocking is enabled.

Likewise, you must create a policy on the receiving device to permit pings from the zone containing the source address to pass through the VPN tunnel to the zone containing the destination address if:

- The destination address is in a different zone from the source address.
- The destination address is in the same zone as the source address, and intrazone blocking is enabled.

Note: *If the receiving device is a third-party product that does not respond to the ICMP echo requests, change the destination to an internal host in the remote peer's LAN that does respond. The remote peer's firewall must have a policy permitting the ICMP echo requests to pass through it.*

Configuring the VPN Monitoring Feature

To enable VPN monitoring, do the following:

WebUI

VPNs > AutoKey IKE > New: Configure the VPN, click **Advanced**, enter the following, click **Return** to go back to the basic VPN configuration page, and then click **OK**:

VPN Monitor: Select to enable VPN monitoring of this VPN tunnel.

Source Interface: Choose an interface from the drop-down list. If you choose "default", the NetScreen device uses the outgoing interface.

Destination IP: Enter a destination IP address. If you do not enter anything, the NetScreen device uses the remote gateway IP address.

Rekey: Select this option if you want the NetScreen device to attempt IKE Phase 2 negotiations—and IKE Phase 1 negotiations if necessary—if the tunnel status changes from up to down. When you select this option, the

NetScreen device attempts IKE negotiations to set up the tunnel and begin VPN monitoring immediately after you finish configuring the tunnel.

Clear this option if you do not want the NetScreen device to attempt IKE negotiations if the tunnel status changes from up to down. When the rekey option is disabled, VPN monitoring begins after user-generated traffic has triggered IKE negotiations and stops when the tunnel status changes from up to down.

(Or)

VPNs > Manual Key > New: Configure the VPN, click **Advanced**, enter the following, click **Return** to go back to the basic VPN configuration page, and then click **OK**:

VPN Monitor: Select to enable VPN monitoring of this VPN tunnel.

Source Interface: Choose an interface from the drop-down list. If you choose "default", the NetScreen device uses the outgoing interface.

Destination IP: Enter a destination IP address. If you do not enter anything, the NetScreen device uses the remote gateway IP address.

CLI

```
set vpnmonitor frequency number[5]
set vpnmonitor threshold number[6]
set vpn name_str monitor [ source-interface interface[7] [ destination-ip
    ip_addr[8] ] ] [optimized] [ rekey[9] ]
save
```

5. The VPN monitoring frequency is in seconds. The default setting is 10-second intervals.
6. The VPN monitoring threshold number is the consecutive number of successful or unsuccessful ICMP echo requests that determines whether the remote gateway is reachable through the VPN tunnel or not. The default threshold is 10 consecutive successful or 10 consecutive unsuccessful ICMP echo requests.
7. If you do not choose a source interface, the NetScreen device uses the outgoing interface as the default.
8. If you do not choose a destination IP address, the NetScreen device uses the IP address for the remote gateway.
9. The rekey option is not available for Manual Key VPN tunnels.

Example: Specifying Source and Destination Addresses for VPN Monitoring

In this example, you configure an AutoKey IKE VPN tunnel between two NetScreen devices (NetScreen-A and NetScreen-B). On device A, you set up VPN monitoring from its Trust zone interface (ethernet1) to the Trust zone interface (10.2.1.1/24) on NetScreen-B. On the NetScreen-B, you set up VPN monitoring from its Trust zone interface (ethernet1) to a corporate intranet server (10.1.1.5) behind NetScreen-A.

NetScreen-A	NetScreen-B
Zones and Interfaces	
• ethernet1 - Zone: Trust - IP address: 10.1.1.1/24 - Interface mode: NAT	• ethernet1 - Zone: Trust - IP address: 10.2.1.1/24 - Interface mode: NAT
• ethernet3 - Zone: Untrust - IP address: 1.1.1.1/24	• ethernet3 - Zone: Untrust - IP address: 2.2.2.2/24
Route-Based AutoKey IKE Tunnel Parameters	
• Phase 1 - Gateway name: gw1 - Gateway static IP address: 2.2.2.2 - Security level: Compatible* - Preshared Key: Ti82g4aX - Outgoing interface: ethernet3 - Mode: Main	• Phase 1 - Gateway name: gw1 - Gateway static IP address: 1.1.1.1 - Proposals: Compatible - Preshared Key: Ti82g4aX - Outgoing interface: ethernet3 - Mode: Main
• Phase 2 - VPN tunnel name: vpn1 - Security level: Compatible† - VPN Monitoring: src = ethernet1; dst = 10.2.1.1 - Bound to interface: tunnel.1	• Phase 2 - VPN tunnel name: vpn1 - Security level: Compatible - VPN Monitoring: src = ethernet1; dst = 10.1.1.5 - Bound to interface: tunnel.1

* A Phase 1 security level of Compatible includes the following proposals: pre-g2-3des-sha, pre-g2-3des-md5, pre-g2-des-sha, and pre-g2-des-md5.

† A Phase 1 security level of Compatible includes the following proposals: nopfs-esp-3des-sha, nopfs-esp-3des-md5, nopfs-esp-des-sha, and nopfs-esp-des-md5.

NetScreen-A	NetScreen-B
Routes	
To 0.0.0.0/0, use ethernet3, gateway 1.1.1.250	To 0.0.0.0/0, use ethernet3, gateway 2.2.2.250
To 10.2.1.0/0, use tunnel.1, no gateway	To 10.1.1.0/0, use tunnel.1, no gateway

Because both devices ping from an interface in their Trust zone to an address in their Untrust zone, the admins at both ends of the VPN tunnel must define policies permitting pings to pass from zone to zone.

Note: *Because both VPN terminators are NetScreen devices in this example, you can use the default source and destination addresses for VPN monitoring. The use of other options is included purely to illustrate how you can configure a NetScreen device to use them.*

WebUI (NetScreen-A)

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Enter the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

Network > Interfaces > Tunnel IF New: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.1

Zone (VR): Trust (trust-vr)

Unnumbered: (select)

Interface: ethernet1(trust-vr)

2. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Trust_LAN

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.0/24

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Remote_LAN

IP Address/Domain Name:

IP/Netmask: (select), 10.2.1.0/24

Zone: Untrust

3. VPN

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: vpn1

Security Level: Compatible

Remote Gateway:

Create a Simple Gateway: (select)

Gateway Name: gw1

Type:

Static IP: (select), Address/Hostname: 2.2.2.2

Preshared Key: Ti82g4aX

Security Level: Compatible

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Bind to: Tunnel Interface, tunnel.1

Proxy-ID: (select)

Local IP / Netmask: 10.1.1.0/24

Remote IP / Netmask: 10.2.1.0/24

Service: ANY

VPN Monitor: (select)

Source Interface: ethernet1

Destination IP: 10.2.1.1

Rekey: (clear)

4. Routes

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 1.1.1.250

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 10.2.1.0/24

Gateway: (select)

Interface: Tunnel.1

Gateway IP Address: 0.0.0.0

5. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Trust_LAN

Destination Address:

Address Book Entry: (select), Remote_LAN

Service: ANY

Action: Permit

Position at Top: (select)

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Remote_LAN

Destination Address:

Address Book Entry: (select), Trust_LAN

Service: Any

Action: Permit

Position at Top: (select)

WebUI (NetScreen-B)

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.2.1.1/24

Enter the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 2.2.2.2/24

Network > Interfaces > Tunnel IF New: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.1

Zone (VR): Trust (trust-vr)

Unnumbered: (select)

Interface: ethernet1(trust-vr)

2. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Trust_LAN

IP Address/Domain Name:

IP/Netmask: (select), 10.2.1.0/24

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Remote_LAN

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.0/24

Zone: Untrust

3. VPN

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: vpn1

Security Level: Compatible

Remote Gateway:

Create a Simple Gateway: (select)

Gateway Name: gw1

Type:

Static IP: (select), Address/Hostname: 1.1.1.1

Preshared Key: Ti82g4aX

Security Level: Compatible

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Bind to: Tunnel Interface, tunnel.1

Proxy-ID: (select)

Local IP / Netmask: 10.2.1.0/24

Remote IP / Netmask: 10.1.1.0/24

Service: ANY

VPN Monitor: (select)

Source Interface: ethernet1

Destination IP: 10.1.1.5

Rekey: (clear)

4. Routes

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 2.2.2.250

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 10.1.1.0/24

Gateway: (select)

Interface: Tunnel.1

Gateway IP Address: 0.0.0.0

5. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Trust_LAN

Destination Address:

Address Book Entry: (select), Remote_LAN

Service: ANY

Action: Permit

Position at Top: (select)

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Remote_LAN

Destination Address:

Address Book Entry: (select), Trust_LAN

Service: Any

Action: Permit

Position at Top: (select)

CLI (NetScreen-A)

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat
set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
set interface tunnel.1 zone trust
set interface tunnel.1 ip unnumbered interface ethernet1
```

2. Addresses

```
set address trust Trust_LAN 10.1.1.0/24
set address untrust Remote_LAN 10.2.1.0/24
```

3. VPN

```
set ike gateway gw1 address 2.2.2.2 main outgoing-interface ethernet3 preshare
    Ti82g4aX sec-level compatible
set vpn vpn1 gateway gw1 sec-level compatible
set vpn vpn1 bind interface tunnel.1
set vpn vpn1 proxy-id local-ip 10.1.1.0/24 remote-ip 10.2.1.0/24 any
set vpn vpn1 monitor source-interface ethernet1 destination-ip 10.2.1.1
```

4. Routes

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
set vrouter trust-vr route 10.2.1.0/24 interface tunnel.1
```

5. Policies

```
set policy top from trust to untrust Trust_LAN Remote_LAN any permit
set policy top from untrust to trust Remote_LAN Trust_LAN any permit
save
```

CLI (NetScreen-B)

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.2.1.1/24
set interface ethernet1 nat
set interface ethernet3 zone untrust
set interface ethernet3 ip 2.2.2.2/24
set interface tunnel.1 zone trust
set interface tunnel.1 ip unnumbered interface ethernet1
```

2. Addresses

```
set address trust Trust_LAN 10.2.1.0/24
set address untrust Remote_LAN 10.1.1.0/24
```

3. VPN

```
set ike gateway gw1 address 1.1.1.1 main outgoing-interface ethernet3 preshare
    Ti82g4aX sec-level compatible
set vpn vpn1 gateway gw1 sec-level compatible
set vpn vpn1 bind interface tunnel.1
set vpn vpn1 proxy-id local-ip 10.2.1.0/24 remote-ip 10.1.1.0/24 any
set vpn vpn1 monitor source-interface ethernet1 destination-ip 10.1.1.5
```

4. Routes

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 2.2.2.250
set vrouter trust-vr route 10.1.1.0/24 interface tunnel.1
```

5. Policies

```
set policy top from trust to untrust Trust_LAN Remote_LAN any permit
set policy top from untrust to trust Remote_LAN Trust_LAN any permit
save
```

Security Consideration for a Route-Based VPN Design

When using VPN monitoring with a route-based VPN tunnel configuration, the state of a tunnel might change from up to down. When this occurs, all route table entries referencing that interface change to inactive. Then, when the NetScreen device does a route lookup for traffic originally intended to be encrypted and sent through a VPN tunnel bound to that tunnel interface, it bypasses the route referencing the tunnel interface and searches for a route with the next longest match. The route that it finds might be the default route. Using this route, the NetScreen device would then send the traffic unencrypted out through a non-tunnel interface to the public WAN.

To avoid rerouting traffic originally intended for a tunnel interface to a non-tunnel interface, you can configure the NetScreen device to drop such traffic instead of sending it out unencrypted. To accomplish this, use either of the following work-arounds:

- Decoy Tunnel Interface

  1. Create a second tunnel interface, but do not bind it to a VPN tunnel. Instead, bind it to a tunnel zone that is in the same virtual routing domain as the first tunnel interface[10].
  2. Define a second route to the same destination using this second tunnel interface, and assign it a high metric.

     Then, when the state of the functioning tunnel interface changes from up to down and the route table entry referencing that interface becomes inactive, all subsequent route lookups find this second route to the nonfunctioning tunnel interface. The NetScreen device forwards traffic to the second tunnel interface and because it is not bound to a VPN tunnel, the device drops the traffic.

10. If a tunnel interface is bound to a tunnel zone, its status is always up.

- Virtual Router for Tunnel Interfaces

  1. Create a separate virtual router to use for all routes pointing to tunnel interfaces and name it, for example, “VR-VPN”.
  2. Create a security zone—named, for example, “VPN zone”—and bind it to VR-VPN.
  3. Bind all tunnel interfaces to the VPN zone, and also put all addresses for remote sites that you want to reach through VPN tunnels in this zone.
  4. Configure static routes in all other virtual routers to VR-VPN for traffic that you want encrypted and sent through the tunnels. If necessary, define static routes for decrypted traffic from VR-VPN to the other virtual routers. Such routes are necessary to allow inbound VPN traffic through the tunnel if it is initiated from the remote site.

     If the state of a tunnel interface changes from up to down, the NetScreen device still forwards traffic to VR-VPN, where—because the state of the route to that interface is now inactive and there are no other matching routes—the NetScreen device drops the traffic.

SNMP VPN Monitoring Objects and Traps

ScreenOS provides the ability to determine the status and condition of active VPNs through the use of Simple Network Management Protocol (SNMP) VPN monitoring objects and traps. The VPN monitoring MIB notes whether each ICMP echo request elicits a reply, a running average of successful replies, the latency of the reply, and the average latency over the last 30 attempts.

Note: *To enable your SNMP manager application to recognize the VPN monitoring MIBs, you must import the NetScreen-specific MIB extension files into the application. You can find the MIB extension files on the NetScreen documentation CD that shipped with your NetScreen device.*

By enabling the VPN monitoring feature on an AutoKey IKE or Manual Key VPN tunnel, the NetScreen device activates its SNMP VPN monitoring objects, which include data on the following:

- The total number of active VPN sessions
- The time each session started
- The Security Association (SA) elements for each session:
  - ESP encryption (DES or 3DES) and authentication algorithm (MD5 or SHA-1) types
  - AH algorithm type (MD5 or SHA-1)
  - Key exchange protocol (AutoKey IKE or Manual Key)
  - Phase 1 authentication method (Preshared Key or certificates)
  - VPN type (dialup or peer-to-peer)
  - Peer and local gateway IP addresses
  - Peer and local gateway IDs
  - Security Parameter Index (SPI) numbers
- Session status parameters
  - VPN monitoring status (up or down)
  - Tunnel status (up or down)
  - Phase 1 and 2 status (inactive or active)
  - Phase 1 and 2 lifetime (time in seconds before rekeying; Phase 2 lifetime is also reported in remaining bytes before rekeying)

MULTIPLE TUNNELS PER TUNNEL INTERFACE

You can bind multiple IPSec VPN tunnels to a single tunnel interface. To link a specific destination to one of a number of VPN tunnels bound to the same tunnel interface, the NetScreen device uses two tables: the route table and the next-hop tunnel binding (NHTB). The NetScreen device maps the next-hop gateway IP address specified in the route table entry to a particular VPN tunnel specified in the NHTB table. With this technique, a single tunnel interface can support many VPN tunnels. (See "Route-to-Tunnel Mapping" on page 327.)

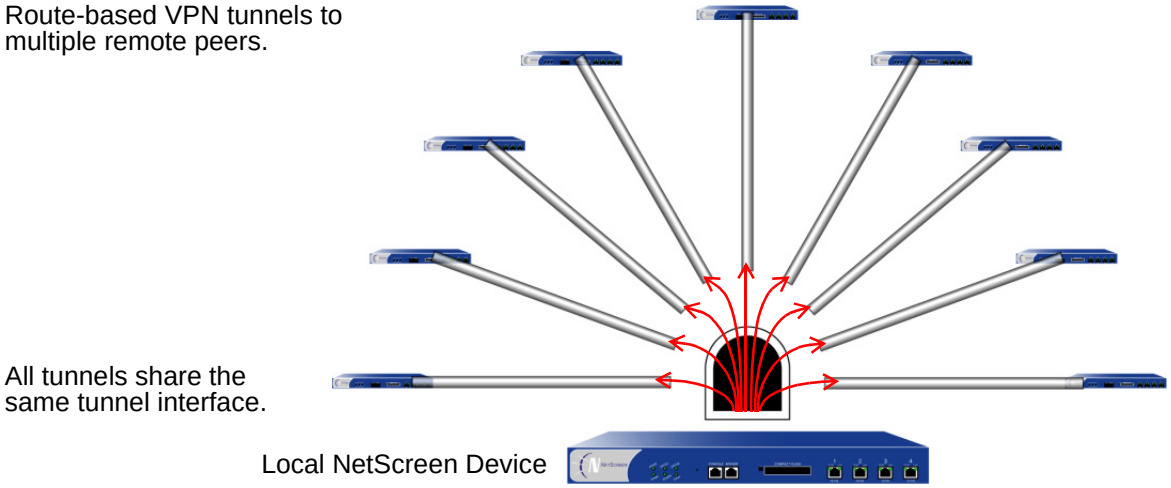

The NetScreen device can sort VPN traffic sent through a single tunnel interface to as many VPN tunnels as the route table or VPN tunnel capacity—whichever is lower—can support.

The maximum number of VPN tunnels is not limited by the number of tunnel interfaces that you can create, but by either route table capacity or the maximum number of dedicated VPN tunnels allowed—whichever is lower. For instance, if your NetScreen device supports 4000 routes and 1000 dedicated VPN tunnels, you can create 1000 VPN tunnels and bind them to a single tunnel interface. If your NetScreen device supports 8192 routes and 10,000 dedicated VPN tunnels, then you can create over 8000 VPN tunnels and bind them to a single tunnel interface[11]. To see the maximum route and tunnel capacities for your NetScreen device, refer to the relevant product data sheet.

11. If route table capacity is the limiting factor, you must subtract the routes automatically generated by security zone interfaces and any other static routes—such as the route to the default gateway—that you might need to define from the total available for route-based VPN tunnels.

Route-to-Tunnel Mapping

To sort traffic among multiple VPN tunnels bound to the same tunnel interface, the NetScreen device maps the next-hop gateway IP address specified in the route to a particular VPN tunnel name. The mapping of entries in the route table to entries in the NHTB table is shown below. In the following illustration, the local NetScreen device routes traffic sent from 10.2.1.5 to 10.1.1.5 through the tunnel.1 interface and then through vpn2.

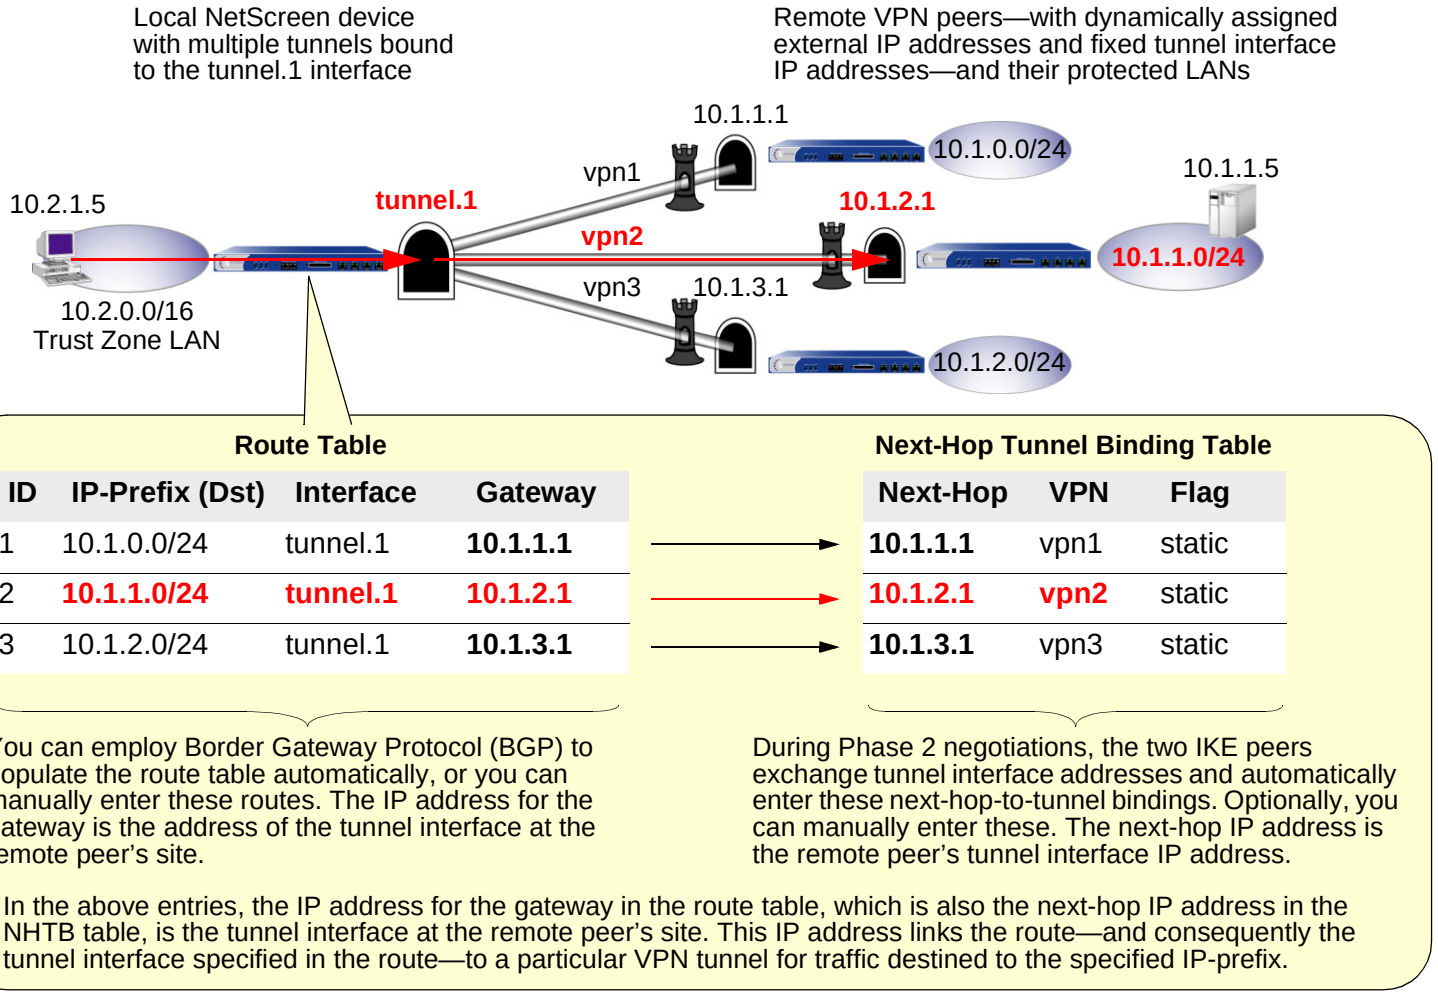

ID	IP-Prefix (Dst)	Interface	Gateway
1	10.1.0.0/24	tunnel.1	**10.1.1.1**
2	**10.1.1.0/24**	**tunnel.1**	**10.1.2.1**
3	10.1.2.0/24	tunnel.1	**10.1.3.1**

Next-Hop	VPN	Flag
10.1.1.1	vpn1	static
10.1.2.1	**vpn2**	static
10.1.3.1	vpn3	static

You can employ Border Gateway Protocol (BGP) to populate the route table automatically, or you can manually enter these routes. The IP address for the gateway is the address of the tunnel interface at the remote peer's site.

During Phase 2 negotiations, the two IKE peers exchange tunnel interface addresses and automatically enter these next-hop-to-tunnel bindings. Optionally, you can manually enter these. The next-hop IP address is the remote peer's tunnel interface IP address.

In the above entries, the IP address for the gateway in the route table, which is also the next-hop IP address in the NHTB table, is the tunnel interface at the remote peer's site. This IP address links the route—and consequently the tunnel interface specified in the route—to a particular VPN tunnel for traffic destined to the specified IP-prefix.

The NetScreen device uses the IP address of the remote peer's tunnel interface as the gateway and next-hop IP address. You can enter the route manually, or you can allow Border Gateway Protocol (BGP) to enter a route referencing the peer's tunnel interface IP address as the gateway in the route table automatically[12]. The same IP address must also be entered as the next hop, along with the appropriate VPN tunnel name, in the NHTB table. Again, there are two options: you can either enter it manually, or you can allow the NetScreen device to obtain it from the remote peer during Phase 2 negotiations and enter it automatically.

The NetScreen device uses the gateway IP address in the route table entry and the next-hop IP address in the NHTB table entry as the common element to link the tunnel interface with the corresponding VPN tunnel. The NetScreen device can then direct traffic destined for the IP-prefix specified in the route with the correct VPN tunnel specified in the NHTB table.

Remote Peers' Addresses

The internal addressing scheme for all remote peers reached through route-based VPNs must be unique among each other. One way to accomplish this is for each remote peer to perform network address translation (NAT) for the source and destination addresses. In addition, the tunnel interface IP addresses must also be unique among all remote peers. If you intend to connect to large numbers of remote sites, an address plan becomes imperative. The following is a possible addressing plan for up to 1000 VPN tunnels:

Dst in Local Route Table	Local Tunnel Interface	Gateway/Next-Hop (Peer's Tunnel Interface)	VPN Tunnel
10.0.3.0/24	tunnel.1	10.0.2.1/24	vpn1
10.0.5.0/24	tunnel.1	10.0.4.1/24	vpn2
10.0.7.0/24	tunnel.1	10.0.6.1/24	vpn3
…	…	…	…
10.0.251.0/24	tunnel.1	10.0.250.1/24	vpn125

12. Because a tunnel interface bound to multiple tunnels cannot send dynamic routing protocol broadcasts and multicasts, it cannot support the Open Shortest Path First (OSPF) and Routing Information Protocol (RIP). See "Automatic Table Entries" on page 331.

Dst in Local Route Table	Local Tunnel Interface	Gateway/Next-Hop (Peer's Tunnel Interface)	VPN Tunnel
10.1.3.0/24	tunnel.1	10.1.2.1/24	vpn126
10.1.5.0/24	tunnel.1	10.1.4.1/24	vpn127
10.1.7.0/24	tunnel.1	10.1.6.1/24	vpn128
…	…	…	…
10.1.251.0/24	tunnel.1	10.1.250.1/24	vpn250
10.2.3.0/24	tunnel.1	10.2.2.1/24	vpn251
…	…	…	…
10.2.251.0/24	tunnel.1	10.2.250.1/24	vpn375
…	…	…	…
10.7.3.0/24	tunnel.1	10.7.2.1/24	vpn876
…	…	…	…
10.7.251.0/24	tunnel.1	10.7.250.1/24	vpn1000

The tunnel interface on the local NetScreen device: is 10.0.0.1/24. On all remote hosts, there is a tunnel interface with an IP address, which appears as the gateway/next-hop IP address in the local route table and NHTB table.

For an example illustrating multiple tunnels bound to a single tunnel interface with address translation, see "Example: Multiple VPNs on One Tunnel Interface to Overlapping Subnets" on page 333.

The local NetScreen device and all its peers perform NAT-dst with IP shifting on inbound VPN traffic and NAT-src from the egress tunnel interface IP address with port translation on outbound VPN traffic. For more information about NAT-src and NAT-dst, see “Address Translation” on page **2** -245.

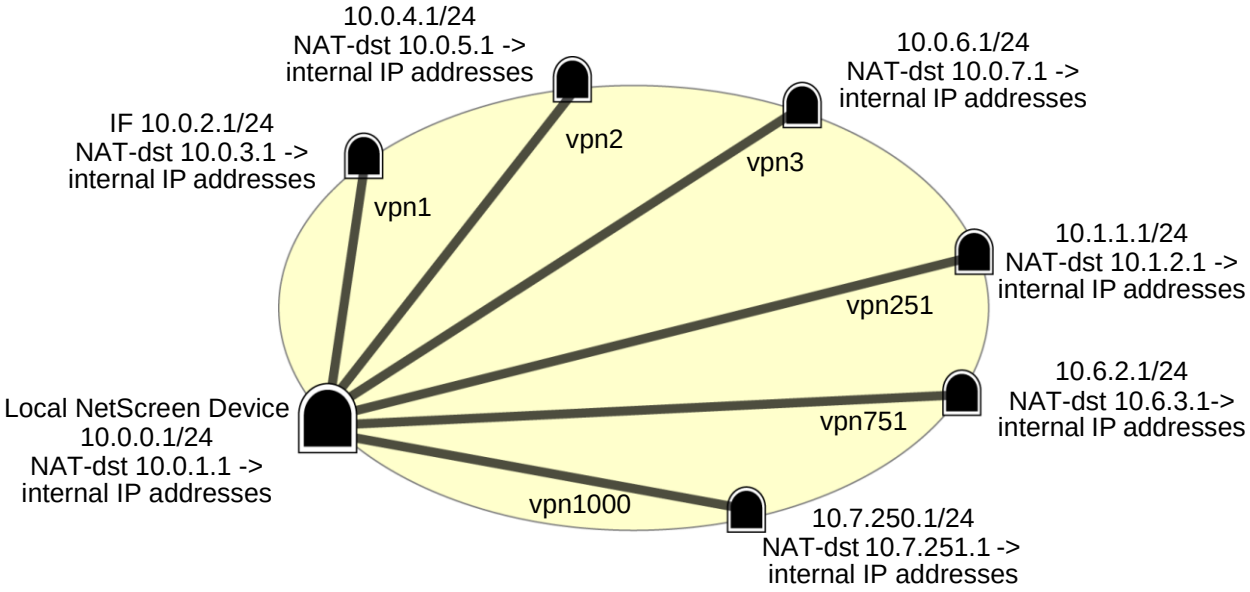

Manual and Automatic Table Entries

You can make entries in the NHTB and route tables manually. You can also automate the populating of the NHTB and route tables. For a small number of tunnels bound to a single tunnel interface, the manual method works well. For a large number of tunnels, the automatic method reduces administrative setup and maintenance as the routes dynamically self-adjust if tunnels or interfaces become unavailable on the tunnel interface at the hub site.

Manual Table Entries

You can manually map a VPN tunnel to the IP address of a remote peer’s tunnel interface in the next-hop tunnel binding (NHTB) table. First, you must contact the remote admin and learn the IP address used for the tunnel interface at that end of a tunnel. Then, you can associate that address with the VPN tunnel name in the NHTB table with the following command:

set interface tunnel.1 nhtb *peer’s_tunnel_interface_addr* **vpn** *name_str*

After that, you can enter a static route in the route table that uses that tunnel interface IP address as the gateway. You can enter the route either through the WebUI or through the following CLI command:

set vrouter *name-str* **route** *dst_addr* **interface tunnel.1 gateway** *peer's_tunnel_interface_addr*

Automatic Table Entries

To make the population of both the NHTB and route tables automatic, the following conditions must be met:

- The remote peers for all VPN tunnels bound to a single local tunnel interface must be NetScreen devices running ScreenOS 5.0.0.
- Each remote peer must bind its tunnel to a tunnel interface, and that interface must have an IP address unique among all peer tunnel interface addresses.
- At both ends of each VPN tunnel, enable VPN monitoring with the rekey option, or enable the IKE heartbeat reconnect option for each remote gateway[13].
- The local and remote peers must have an instance of Border Gateway Protocol (BGP) enabled on their connecting tunnel interfaces.

The use of VPN monitoring with the rekey option allows the NetScreen devices at both ends of a tunnel to set up the tunnel without having to wait for user-originated VPN traffic[14]. After you enable VPN monitoring with the rekey option at both ends of a VPN tunnel, the two NetScreen devices perform Phase 1 and Phase 2 IKE negotiations to establish the tunnel. (For more information, see "VPN Monitoring" on page 307.)

During Phase 2 negotiations, the NetScreen devices exchange tunnel interface IP addresses with each other. Each IKE module can then automatically enter the tunnel interface IP address and its corresponding VPN tunnel name in the NHTB table.

13. If you are running BGP on the tunnel interfaces, traffic generated by the protocol can trigger IKE negotiations even without enabling VPN monitoring with the rekey option or enabling the IKE heartbeat reconnect option. Still, NetScreen recommends that you not rely on dynamic routing traffic to trigger IKE negotiations. Instead use VPN monitoring with the rekey option or the IKE heartbeat reconnect option.

14. For remote peers with a dynamically assigned external IP address or with a fully-qualified domain name (FQDN) mapped to a dynamic IP address, the remote peer must first initiate IKE negotiations. However, because the Phase 2 SA on the local NetScreen device caches the remote peer's dynamically assigned IP address, either peer can reinitiate IKE negotiations to reestablish a tunnel whose VPN monitoring state has changed from up to down.

To enable the local NetScreen device to enter routes to remote destinations automatically in its route table, you must enable an instance of BGP on the local and remote tunnel interfaces. The basic steps are as follows:

1. Create a BGP routing instance on the virtual router that contains the tunnel interface to which you have bound multiple VPN tunnels.
2. Enable the routing instance on the virtual router.
3. Enable the routing instance on the tunnel interface leading to the BPG peers.

The remote peers also perform these steps.

On the local (or hub) device, you must also define a default route and a static route to each peer's tunnel interface IP address. Static routes to the peers' tunnel interfaces are necessary for the hub device to reach its BGP neighbors initially through the correct VPN tunnel.

After establishing communications, the BGP neighbors exchange routing information so that they can each automatically populate their route tables. After the two peers establish a VPN tunnel between themselves, the remote peers can send and receive routing information to and from the local device. When the dynamic routing instance on the local NetScreen device learns a route to a peer through a local tunnel interface, it includes the IP address of the remote peer's tunnel interface as the gateway in the route.

For an example illustrating the configuration of multiple tunnels bound to a single tunnel interface where the "hub" device populates its NHTB and route tables automatically, see "Example: Automatic Route and NHTB Table Entries" on page 364.

Example: Multiple VPNs on One Tunnel Interface to Overlapping Subnets

In this example, you bind three route-based AutoKey IKE VPN tunnels—vpn1, vpn2, and vpn3—to a single tunnel interface—tunnel.1. The tunnels lead from NetScreen-A to three remote peers—peer1, peer2, and peer3. You manually add both the route table and NHTB table entries on NetScreen-A for all three peers.

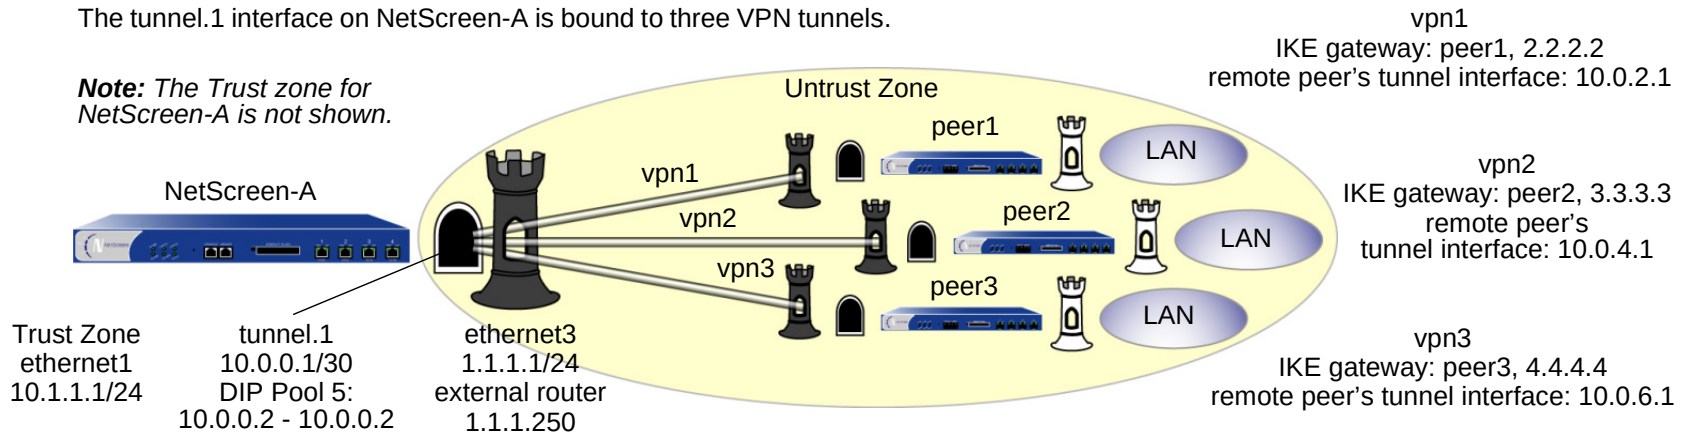

The VPN tunnel configurations at both ends of each tunnel use the following parameters: AutoKey IKE, preshared key (peer1: "netscreen1", peer2: "netscreen2", peer3: "netscreen3"), and the security level predefined as "Compatible" for both Phase 1 and Phase 2 proposals. (For details about these proposals, see "Tunnel Negotiation" on page 11.)

All security zones and interfaces on each device are in the trust-vr virtual routing domain for that device.

This example uses the same address space—10.1.1.0/24—for every LAN to show how you can use source and destination network address translation (NAT-src and NAT-dst) to overcome addressing conflicts among IPSec peers. For more information about NAT-src and NAT-dst, see "Address Translation" on page **2**-245.

WebUI (NetScreen-A)

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Enter the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

Network > Interfaces > New Tunnel IF: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.1

Zone (VR): Untrust (trust-vr)

Fixed IP: (select)

IP Address / Netmask: 10.0.0.1/30

Network > Interfaces > Edit (for tunnel.1) > DIP > New: Enter the following, and then click **OK**:

ID: 5

IP Address Range: (select), 10.0.0.2 ~ 10.0.0.2

Port Translation: (select)

In the same subnet as the interface IP or its secondary IPs: (select)

2. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: corp

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.0/24

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: oda1

IP Address/Domain Name:

IP/Netmask: (select), 10.0.1.0/24

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: peers

IP Address/Domain Name:

IP/Netmask: (select), 10.0.0.0/16

Zone: Untrust

3. VPNs

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: vpn1

Security Level: Compatible

Remote Gateway: Create a Simple Gateway: (select)

Gateway Name: peer1

Type: Static IP: (select), Address/Hostname: 2.2.2.2

Preshared Key: netscreen1

Security Level: Compatible

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Bind to: Tunnel Interface, tunnel.1

Proxy-ID: (select)

Local IP / Netmask: 0.0.0.0/0

Remote IP / Netmask: 0.0.0.0/0

Service: ANY

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: vpn2

Security Level: Compatible

Remote Gateway: Create a Simple Gateway: (select)

Gateway Name: peer2

Type: Static IP: (select), Address/Hostname: 3.3.3.3

Preshared Key: netscreen2

Security Level: Compatible

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Bind to: Tunnel Interface, tunnel.1

Proxy-ID: (select)

Local IP / Netmask: 0.0.0.0/0

Remote IP / Netmask: 0.0.0.0/0

Service: ANY

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: vpn3

Security Level: Compatible

Remote Gateway: Create a Simple Gateway: (select)

Gateway Name: peer3

Type: Static IP: (select), Address/Hostname: 4.4.4.4

Preshared Key: netscreen3

Security Level: Compatible

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Bind to: Tunnel Interface, tunnel.1

Proxy-ID: (select)

Local IP / Netmask: 0.0.0.0/0

Remote IP / Netmask: 0.0.0.0/0

Service: ANY

4. Routes

Network > Routing > Routing Entries > (trust-vr) New: Enter the following, and then click **OK**:

Network Address / Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 1.1.1.250

Network > Routing > Routing Entries > (trust-vr) New: Enter the following, and then click **OK**:

Network Address / Netmask: 10.0.1.0/24

Gateway: (select)

Interface: ethernet1

Gateway IP Address: 0.0.0.0

Network > Routing > Routing Entries > (trust-vr) New: Enter the following, and then click **OK**:

Network Address / Netmask: 10.0.3.0/24

Gateway: (select)

Interface: tunnel.1

Gateway IP Address: 10.0.2.1

Network > Routing > Routing Entries > (trust-vr) New: Enter the following, and then click **OK**:

Network Address / Netmask: 10.0.2.2/32

Gateway: (select)

Interface: tunnel.1

Gateway IP Address: 10.0.2.1

Network > Routing > Routing Entries > (trust-vr) New: Enter the following, and then click **OK**:

Network Address / Netmask: 10.0.5.0/24

Gateway: (select)

Interface: tunnel.1

Gateway IP Address: 10.0.4.1

Network > Routing > Routing Entries > (trust-vr) New: Enter the following, and then click **OK**:

Network Address / Netmask: 10.0.4.2/32

Gateway: (select)

Interface: tunnel.1

Gateway IP Address: 10.0.4.1

Network > Routing > Routing Entries > (trust-vr) New: Enter the following, and then click **OK**:

Network Address / Netmask: 10.0.7.0/24

Gateway: (select)

Interface: tunnel.1

Gateway IP Address: 10.0.6.1

Network > Routing > Routing Entries > (trust-vr) New: Enter the following, and then click **OK**:

Network Address / Netmask: 10.0.6.2/32

Gateway: (select)

Interface: tunnel.1

Gateway IP Address: 10.0.6.1

Network > Interfaces > Edit (for tunnel.1) > NHTB > New: Enter the following, and then click **Add**:

New Next Hop Entry:

IP Address: 10.0.2.1

VPN: vpn1

Network > Interfaces > Edit (for tunnel.1) > NHTB: Enter the following, and then click **Add**:

New Next Hop Entry:

IP Address: 10.0.4.1

VPN: vpn2

Network > Interfaces > Edit (for tunnel.1) > NHTB: Enter the following, and then click **Add**:

New Next Hop Entry:

IP Address: 10.0.6.1

VPN: vpn3

5. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book: (select), corp

Destination Address:

Address Book: (select), peers

Service: Any

Action: Permit

Position at Top: (select)

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Policy configuration page:

NAT:

Source Translation: (select)

DIP On: 5 (10.0.0.2–10.0.0.2)/X-late

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), peers

Destination Address:

Address Book Entry: (select), oda1

Service: Any

Action: Permit

Position at Top: (select)

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Policy configuration page:

NAT:

Destination Translation: (select)

Translate to IP Range: (select), 10.1.1.0 - 10.1.1.254

CLI (NetScreen-A)

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
set interface tunnel.1 zone untrust
set interface tunnel.1 ip 10.0.0.1/30
set interface tunnel.1 dip 5 10.0.0.2 10.0.0.2
```

2. Addresses

```
set address trust corp 10.1.1.0/24
set address trust oda1 10.0.1.0/24
set address untrust peers 10.0.0.0/16
```

3. VPNs

```
set ike gateway peer1 address 2.2.2.2 outgoing-interface ethernet3 preshare
    netscreen1 sec-level compatible
set vpn vpn1 gateway peer1 sec-level compatible
set vpn vpn1 bind interface tunnel.1
set vpn vpn1 proxy-id local-ip 0.0.0.0/0 remote-ip 0.0.0.0/0 any
set ike gateway peer2 address 3.3.3.3 outgoing-interface ethernet3 preshare
    netscreen2 sec-level compatible
```

```
set vpn vpn2 gateway peer2 sec-level compatible
set vpn vpn2 bind interface tunnel.1
set vpn vpn2 proxy-id local-ip 0.0.0.0/0 remote-ip 0.0.0.0/0 any
set ike gateway peer3 address 4.4.4.4 outgoing-interface ethernet3 preshare
    netscreen3 sec-level compatible
set vpn vpn3 gateway peer3 sec-level compatible
set vpn vpn3 bind interface tunnel.1
set vpn vpn3 proxy-id local-ip 0.0.0.0/0 remote-ip 0.0.0.0/0 any
```

4. Routes

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
set vrouter trust-vr route 10.0.1.0/24 interface ethernet1
set vrouter trust-vr route 10.0.3.0/24 interface tunnel.1 gateway 10.0.2.1
set vrouter trust-vr route 10.0.2.2/32 interface tunnel.1 gateway 10.0.2.1
set vrouter trust-vr route 10.0.5.0/24 interface tunnel.1 gateway 10.0.4.1
set vrouter trust-vr route 10.0.4.2/32 interface tunnel.1 gateway 10.0.4.1
set vrouter trust-vr route 10.0.7.0/24 interface tunnel.1 gateway 10.0.6.1
set vrouter trust-vr route 10.0.6.2/32 interface tunnel.1 gateway 10.0.6.1
set interface tunnel.1 nhtb 10.0.2.1 vpn vpn1
set interface tunnel.1 nhtb 10.0.4.1 vpn vpn2
set interface tunnel.1 nhtb 10.0.6.1 vpn vpn3
```

5. Policies

```
set policy from trust to untrust corp peers any nat src dip-id 5 permit
set policy from untrust to trust peers oda1 any nat dst ip 10.1.1.0 10.1.1.254
    permit
save
```

Peer1

The following configuration is what the remote admin for the NetScreen device at the peer1 site must enter to create a VPN tunnel to NetScreen-A at the corporate site. The remote admin configures the NetScreen device to perform source and destination NAT (NAT-src and NAT-dst) because the internal addresses are in the same address space as those in the corporate LAN: 10.1.1.0/24. Peer1 performs NAT-src using DIP pool 6 to translate all internal source addresses to 10.0.2.2 when sending traffic through VPN1 to NetScreen-A. Peer1 performs NAT-dst on VPN traffic sent from NetScreen-A, translating addresses from 10.0.3.0/24 to 10.1.1.0/24 with address shifting in effect.

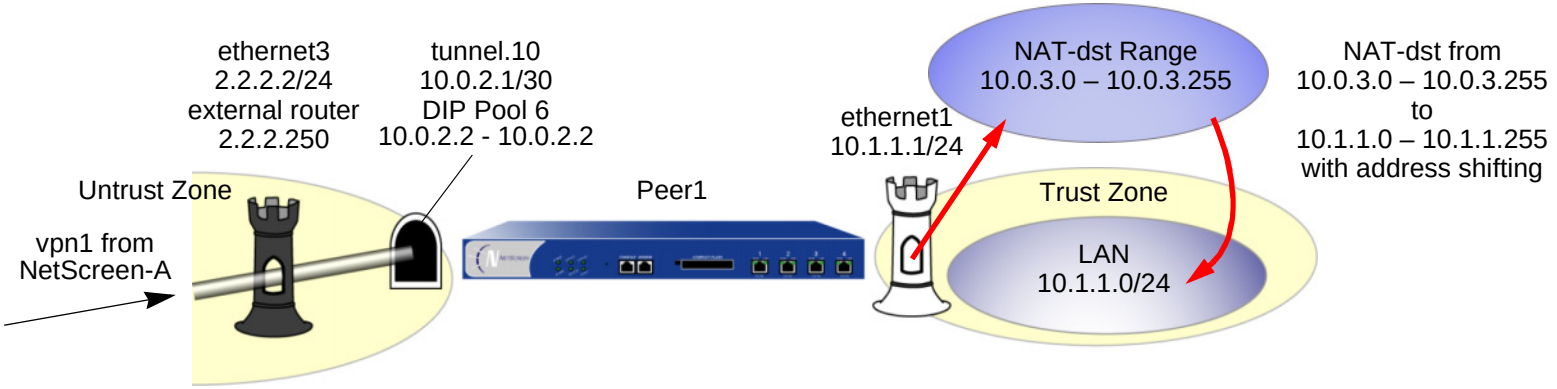

Note: *For more information about NAT-src and NAT-dst, see "Address Translation" on page* ***2****-245.*

WebUI (Peer1)

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Enter the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 2.2.2.2/24

Network > Interfaces > New Tunnel IF: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.10

Zone (VR): Untrust (trust-vr)

Fixed IP: (select)

IP Address / Netmask: 10.0.2.1/30

Network > Interfaces > Edit (for tunnel.10) > DIP > New: Enter the following, and then click **OK**:

ID: 6

IP Address Range: (select), 10.0.2.2 ~ 10.0.2.2

Port Translation: (select)

In the same subnet as the interface IP or its secondary IPs: (select)

2. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: lan

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.0/24

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: oda2

IP Address/Domain Name:

IP/Netmask: (select), 10.0.3.0/24

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: to_corp

IP Address/Domain Name:

IP/Netmask: (select), 10.0.1.0/24

Zone: Untrust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: fr_corp

IP Address/Domain Name:

IP/Netmask: (select), 10.0.0.2/32

Zone: Untrust

3. VPN

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: vpn1

Security Level: Compatible

Remote Gateway: Create a Simple Gateway: (select)

Gateway Name: corp

Type: Static IP: (select), Address/Hostname: 1.1.1.1

Preshared Key: netscreen1

Security Level: Compatible

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Bind to: Tunnel Interface, tunnel.10

Proxy-ID: (select)

Local IP / Netmask: 0.0.0.0/0

Remote IP / Netmask: 0.0.0.0/0

Service: ANY

4. Routes

Network > Routing > Routing Entries > (trust-vr) New: Enter the following, and then click **OK**:

Network Address / Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 2.2.2.250

Metric: 1

Network > Routing > Routing Entries > (trust-vr) New: Enter the following, and then click **OK**:

Network Address / Netmask: 10.0.3.0/24

Gateway: (select)

Interface: ethernet1

Gateway IP Address: 0.0.0.0

Metric: 1

Network > Routing > Routing Entries > (trust-vr) New: Enter the following, and then click **OK**:

Network Address / Netmask: 10.0.0.0/8

Gateway: (select)

Interface: tunnel.10

Gateway IP Address: 0.0.0.0

Metric: 10

5. Policies

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), fr_corp

Destination Address:

Address Book Entry: (select), oda2

Service: Any

Action: Permit

Position at Top: (select)

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Policy configuration page:

NAT:

Destination Translation: (select)

Translate to IP Range: (select), 10.1.1.0 - 10.1.1.254

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), lan

Destination Address:

Address Book Entry: (select), to_corp

Service: Any

Action: Permit

Position at Top: (select)

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Policy configuration page:

NAT:

Source Translation: (select)

DIP On: 6 (10.0.2.2–10.0.2.2)/X-late

CLI (Peer1)

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet3 zone untrust
set interface ethernet3 ip 2.2.2.2/24
set interface tunnel.10 zone untrust
set interface tunnel.10 ip 10.0.2.1/30
set interface tunnel.10 dip 6 10.0.2.2 10.0.2.2
```

2. Addresses

```
set address trust lan 10.1.1.0/24
set address trust oda2 10.0.3.0/24
set address untrust to_corp 10.0.1.0/24
set address untrust fr_corp 10.0.0.2/32
```

3. VPN

```
set ike gateway corp address 1.1.1.1 outgoing-interface ethernet3 preshare
    netscreen1 sec-level compatible
set vpn vpn1 gateway corp sec-level compatible
set vpn vpn1 bind interface tunnel.10
set vpn vpn1 proxy-id local-ip 0.0.0.0/0 remote-ip 0.0.0.0/0 any
```

4. Routes

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 2.2.2.250 metric 1
set vrouter trust-vr route 10.0.3.0/24 interface ethernet1 metric 1
set vrouter trust-vr route 10.0.0.0/8 interface tunnel.10 metric 10
```

5. Policies

```
set policy from trust to untrust lan to_corp any nat src dip-id 6 permit
set policy from untrust to trust fr_corp oda2 any nat dst ip 10.1.1.0
    10.1.1.254 permit
save
```

Peer2

The following configuration is what the remote admin for the NetScreen device at the peer2 site must enter to create a VPN tunnel to NetScreen-A at the corporate site. The remote admin configures the NetScreen device to perform source and destination NAT (NAT-src and NAT-dst) because the internal addresses are in the same address space as those in the corporate LAN: 10.1.1.0/24. Peer2 performs NAT-src using DIP pool 7 to translate all internal source addresses to 10.0.4.2 when sending traffic through VPN2 to NetScreen-A. Peer2 performs NAT-dst on VPN traffic sent from NetScreen-A, translating addresses from 10.0.5.0/24 to 10.1.1.0/24 with address shifting in effect.

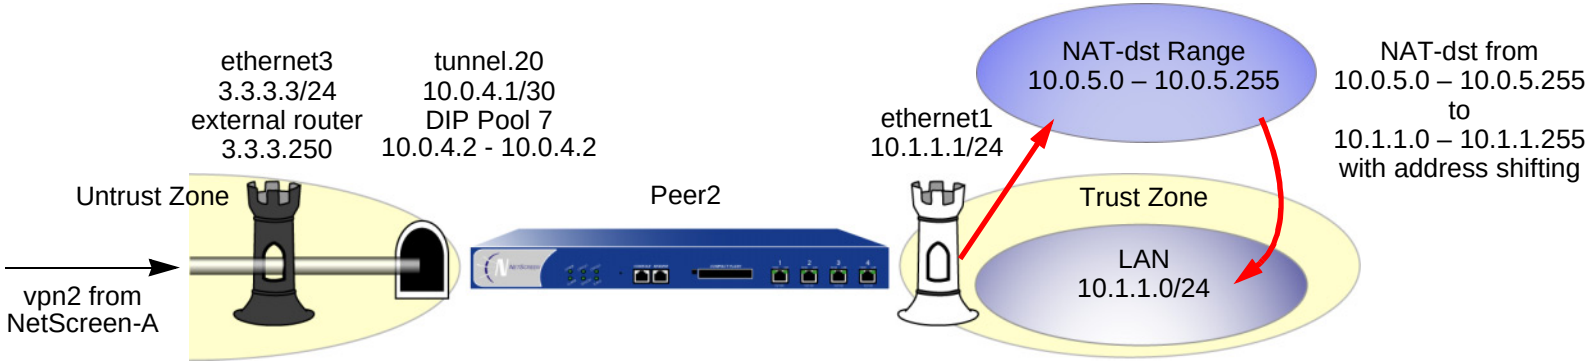

Note: *For more information about NAT-src and NAT-dst, see "Address Translation" on page* ***2****-245.*

WebUI (Peer2)

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Enter the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 3.3.3.3/24

Network > Interfaces > New Tunnel IF: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.20

Zone (VR): Untrust (trust-vr)

Fixed IP: (select)

IP Address / Netmask: 10.0.4.1/30

Network > Interfaces > Edit (for tunnel.20) > DIP > New: Enter the following, and then click **OK**:

ID: 7

IP Address Range: (select), 10.0.4.2 ~ 10.0.4.2

Port Translation: (select)

In the same subnet as the interface IP or its secondary IPs: (select)

2. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: lan

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.0/24

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: oda3

IP Address/Domain Name:

IP/Netmask: (select), 10.0.5.0/24

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: to_corp

IP Address/Domain Name:

IP/Netmask: (select), 10.0.1.0/24

Zone: Untrust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: fr_corp

IP Address/Domain Name:

IP/Netmask: (select), 10.0.0.2/32

Zone: Untrust

3. VPN

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: vpn2

Security Level: Compatible

Remote Gateway: Create a Simple Gateway: (select)

Gateway Name: corp

Type: Static IP: (select), Address/Hostname: 1.1.1.1

Preshared Key: netscreen2

Security Level: Compatible

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Bind to: Tunnel Interface, tunnel.20

Proxy-ID: (select)

Local IP / Netmask: 0.0.0.0/0

Remote IP / Netmask: 0.0.0.0/0

Service: ANY

4. Routes

Network > Routing > Routing Entries > (trust-vr) New: Enter the following, and then click **OK**:

Network Address / Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 3.3.3.250

Metric: 1

Network > Routing > Routing Entries > (trust-vr) New: Enter the following, and then click **OK**:

Network Address / Netmask: 10.0.5.0/24

Gateway: (select)

Interface: ethernet1

Gateway IP Address: 0.0.0.0

Metric: 1

Network > Routing > Routing Entries > (trust-vr) New: Enter the following, and then click **OK**:

Network Address / Netmask: 10.0.1.0/24

Gateway: (select)

Interface: tunnel.20

Gateway IP Address: 0.0.0.0

Metric: 10

5. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), lan

Destination Address:

Address Book Entry: (select), to_corp

Service: Any

Action: Permit

Position at Top: (select)

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Policy configuration page:

NAT:

Source Translation: (select)

DIP On: 7 (10.0.4.2–10.0.4.2)/X-late

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), fr_corp

Destination Address:

Address Book Entry: (select), oda3

Service: Any

Action: Permit

Position at Top: (select)

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Policy configuration page:

NAT:

Destination Translation: (select)

Translate to IP Range: (select), 10.1.1.0 - 10.1.1.254

CLI (Peer2)

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet3 zone untrust
set interface ethernet3 ip 3.3.3.3/24
set interface tunnel.20 zone untrust
set interface tunnel.20 ip 10.0.4.1/30
set interface tunnel.20 dip 7 10.0.4.2 10.0.4.2
```

2. Addresses

```
set address trust lan 10.1.1.0/24
set address trust oda3 10.0.5.0/24
set address untrust to_corp 10.0.1.0/24
set address untrust fr_corp 10.0.0.2/32
```

3. VPN

```
set ike gateway corp address 1.1.1.1 outgoing-interface ethernet3 preshare
    netscreen2 sec-level compatible
set vpn vpn2 gateway corp sec-level compatible
set vpn vpn2 bind interface tunnel.20
set vpn vpn2 proxy-id local-ip 0.0.0.0/0 remote-ip 0.0.0.0/0 any
```

4. Routes

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 3.3.3.250 metric 1
set vrouter trust-vr route 10.0.5.0/24 interface ethernet1 metric 1
set vrouter trust-vr route 10.0.0.0/8 interface tunnel.20 metric 10
```

5. Policies

```
set policy from trust to untrust lan to_corp any nat src dip-id 7 permit
set policy from untrust to trust fr_corp oda3 any nat dst ip 10.1.1.0
    10.1.1.254 permit
save
```

Peer3

The following configuration is what the remote admin for the NetScreen device at the peer3 site must enter to create a VPN tunnel to NetScreen-A at the corporate site. The remote admin configures the NetScreen device to perform source and destination NAT (NAT-src and NAT-dst) because the internal addresses are in the same address space as those in the corporate LAN: 10.1.1.0/24. Peer3 performs NAT-src using DIP pool 8 to translate all internal source addresses to 10.0.6.2 when sending traffic through VPN3 to NetScreen-A. Peer3 performs NAT-dst on VPN traffic sent from NetScreen-A, translating addresses from 10.0.7.0/24 to 10.1.1.0/24 with address shifting in effect.

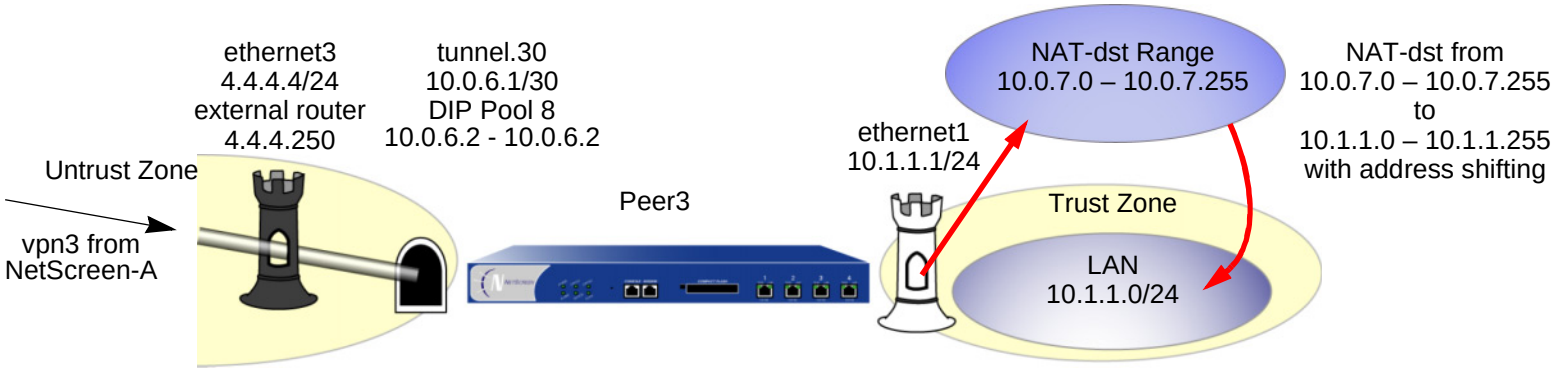

Note: *For more information about NAT-dst, see "Destination Network Address Translation" on page* ***2*** *-276.*

WebUI (Peer3)

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Enter the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 4.4.4.4/24

Network > Interfaces > New Tunnel IF: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.30

Zone (VR): Untrust (trust-vr)

Fixed IP: (select)

IP Address / Netmask: 10.0.6.1/30

Network > Interfaces > Edit (for tunnel.320) > DIP > New: Enter the following, and then click **OK**:

ID: 7

IP Address Range: (select), 10.0.6.2 ~ 10.0.6.2

Port Translation: (select)

In the same subnet as the interface IP or its secondary IPs: (select)

2. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: lan

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.0/24

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: oda4

IP Address/Domain Name:

IP/Netmask: (select), 10.0.7.0/24

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: to_corp

IP Address/Domain Name:

IP/Netmask: (select), 10.0.1.0/24

Zone: Untrust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: fr_corp

IP Address/Domain Name:

IP/Netmask: (select), 10.0.0.2/32

Zone: Untrust

3. VPN

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: vpn3

Security Level: Compatible

Remote Gateway: Create a Simple Gateway: (select)

Gateway Name: corp

Type: Static IP: (select), Address/Hostname: 1.1.1.1

Preshared Key: netscreen3

Security Level: Compatible

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Bind to: Tunnel Interface, tunnel.30

Proxy-ID: (select)

Local IP / Netmask: 0.0.0.0/0

Remote IP / Netmask: 0.0.0.0/0

Service: ANY

4. Routes

Network > Routing > Routing Entries > (trust-vr) New: Enter the following, and then click **OK**:

Network Address / Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 4.4.4.250

Metric: 1

Network > Routing > Routing Entries > (trust-vr) New: Enter the following, and then click **OK**:

Network Address / Netmask: 10.0.7.0/24

Gateway: (select)

Interface: ethernet1

Gateway IP Address: 0.0.0.0

Metric: 1

Network > Routing > Routing Entries > (trust-vr) New: Enter the following, and then click **OK**:

Network Address / Netmask: 10.0.0.0/8

Gateway: (select)

Interface: tunnel.20

Gateway IP Address: 10.0.0.1

Metric: 10

5. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), lan

Destination Address:

Address Book Entry: (select), to_corp

Service: Any

Action: Permit

Position at Top: (select)

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Policy configuration page:

NAT:

Source Translation: (select)

DIP On: 8 (10.0.6.2–10.0.6.2)/X-late

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), fr_corp

Destination Address:

Address Book Entry: (select), oda4

Service: Any

Action: Permit

Position at Top: (select)

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Policy configuration page:

NAT:

Destination Translation: (select)

Translate to IP Range: (select), 10.1.1.0 - 10.1.1.254

CLI (Peer3)

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet3 zone untrust
set interface ethernet3 ip 4.4.4.4/24
set interface tunnel.30 zone untrust
set interface tunnel.30 ip 10.0.6.1/30
set interface tunnel.30 dip 8 10.0.6.2 10.0.6.2
```

2. Addresses

```
set address trust lan 10.1.1.0/24
set address trust oda4 10.0.7.0/24
set address untrust to_corp 10.0.1.0/24
set address untrust fr_corp 10.0.0.2/32
```

3. VPN

```
set ike gateway corp address 1.1.1.1 outgoing-interface ethernet3 preshare
    netscreen3 sec-level compatible
set vpn vpn3 gateway corp sec-level compatible
set vpn vpn3 bind interface tunnel.30
set vpn vpn3 proxy-id local-ip 0.0.0.0/0 remote-ip 0.0.0.0/0 any
```

4. Routes

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 3.3.3.250 metric 1
set vrouter trust-vr route 10.0.7.0/24 interface ethernet1 metric 1
set vrouter trust-vr route 10.0.0.0/8 interface tunnel.30 metric 10
```

5. Policies

```
set policy from trust to untrust lan to_corp any nat src dip-id 8 permit
set policy from untrust to trust fr_corp oda4 any nat dst ip 10.1.1.0
    10.1.1.254 permit
save
```

Example: Automatic Route and NHTB Table Entries

In this example, you bind two route-based AutoKey IKE VPN tunnels—vpn1, vpn2—to a single tunnel interface—tunnel.1 on NetScreen-A at the corporate site. The network that each remote peer protects has multiple routes behind the connected route. Using Border Gateway Protocol (BGP), the peers communicate their routes to NetScreen-A. This example permits VPN traffic from the corporate site behind NetScreen-A to the peer sites.

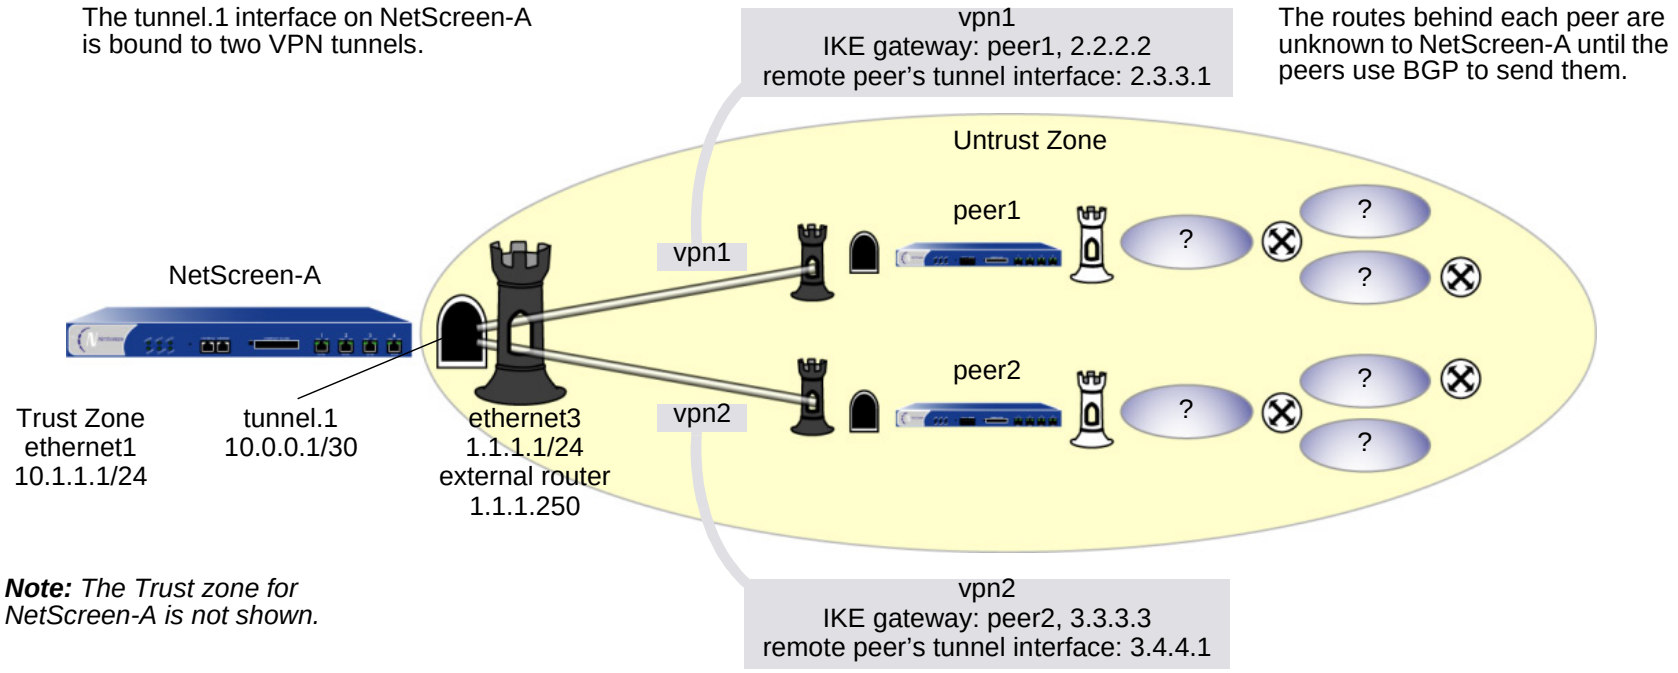

The VPN tunnel configurations at both ends of each tunnel use the following parameters: AutoKey IKE, preshared key (peer1: "netscreen1", peer2: "netscreen2"), and the security level predefined as "Compatible" for both Phase 1 and Phase 2 proposals. (For details about these proposals, see "Tunnel Negotiation" on page 11.)

By configuring the following two features, you can enable NetScreen-A to populate its NHTB and route tables automatically[15]:

- VPN monitoring with the rekey option (or the IKE heartbeats reconnect option)[16]
- BGP dynamic routing on tunnel.1

When you enable VPN monitoring with the rekey option for an AutoKey IKE VPN tunnel, NetScreen-A establishes a VPN connection with its remote peer as soon as you and the admin at the remote site finish configuring the tunnel. The devices do not wait for user-generated VPN traffic to perform IKE negotiations. During Phase 2 negotiations, the NetScreen devices exchange their tunnel interface IP address, so that NetScreen-A can automatically make a VPN-to-next-hop mapping in its NHTB table.

The rekey option ensures that when the Phase 1 and Phase 2 key lifetimes expire, the devices automatically negotiate the generation of new keys without the need for human intervention. VPN monitoring with the rekey option enabled essentially provides a means for keeping a VPN tunnel up continually, even when there is no user-generated traffic. This is necessary so that the BGP dynamic routing instances that you and the remote admins create and enable on the tunnel interfaces at both ends of the tunnels can send routing information to NetScreen-A and automatically populate its route table with the routes it needs to direct traffic through the VPN tunnel before those routes are required for user-generated traffic. (The admins at the peer sites still need to enter a single static route to the rest of the virtual private network through the tunnel interface at each respective site.)

You enter a default route and static routes on NetScreen-A to reach its BGP neighbors through the correct VPN tunnels. All security zones and interfaces on each device are in the trust-vr virtual routing domain for that device.

15. If you are running a dynamic routing protocol on the tunnel interfaces, traffic generated by the protocol can trigger IKE negotiations even without enabling VPN monitoring with the rekey option or enabling the IKE heartbeat reconnect option. Still, NetScreen recommends that you not rely on dynamic routing traffic to trigger IKE negotiations. Instead use VPN monitoring with the rekey option or the IKE heartbeat reconnect option.

16. If you are running BGP on the tunnel interfaces, the BGP-generated traffic can trigger IKE negotiations even without enabling VPN monitoring with the rekey option or enabling the IKE heartbeat reconnect option. Still, NetScreen recommends that you not rely on BGP traffic to trigger IKE negotiations. Instead, use VPN monitoring with the rekey option or the IKE heartbeat reconnect option.

WebUI (NetScreen-A)

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Enter the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

Network > Interfaces > New Tunnel IF: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.1

Zone (VR): Untrust (trust-vr)

Fixed IP: (select)

IP Address / Netmask: 10.0.0.1/30

2. VPNs

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: vpn1

Security Level: Compatible

Remote Gateway: Create a Simple Gateway: (select)

Gateway Name: peer1

Type: Static IP: (select), Address/Hostname: 2.2.2.2

Preshared Key: netscreen1

Security Level: Compatible

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Bind to: Tunnel Interface, tunnel.1

Proxy-ID: (select)

Local IP / Netmask: 0.0.0.0/0

Remote IP / Netmask: 0.0.0.0/0

Service: ANY

VPN Monitor[17]: (select)

Rekey: (select)

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: vpn2

Security Level: Compatible

Remote Gateway: Create a Simple Gateway: (select)

Gateway Name: peer2

Type: Static IP: (select), Address/Hostname: 3.3.3.3

Preshared Key: netscreen2

Security Level: Compatible

Outgoing Interface: ethernet3

17. Leave the Source Interface and Destination IP options at their default settings. For information about these options, see "VPN Monitoring" on page 307.

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Bind to: Tunnel Interface, tunnel.1

Proxy-ID: (select)

Local IP / Netmask: 0.0.0.0/0

Remote IP / Netmask: 0.0.0.0/0

Service: ANY

VPN Monitor[18]: (select)

Rekey: (select)

3. Static Route

Network > Routing > Routing Entries > (trust-vr) New: Enter the following, and then click **OK**:

Network Address / Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 1.1.1.250

Network > Routing > Routing Entries > (trust-vr) New: Enter the following, and then click **OK**:

Network Address / Netmask: 2.3.3.1/32

Gateway: (select)

Interface: tunnel.1

Gateway IP Address: 2.3.3.1

18. Leave the Source Interface and Destination IP options at their default settings. For information about these options, see "VPN Monitoring" on page 307.

Network > Routing > Routing Entries > (trust-vr) New: Enter the following, and then click **OK**:

Network Address / Netmask: 3.4.4.1/32

Gateway: (select)

Interface: tunnel.1

Gateway IP Address: 3.4.4.1

4. Dynamic Routing

Network > Routing > Virtual Routers > Edit (for trust-vr) > Create BGP Instance: Enter the following, and then click **OK**:

AS Number (required): 99

BGP Enabled: (select)

Network > Interfaces > Edit (for tunnel.1) > BGP: Select the **Protocol BGP** check box, and then click **OK**.

Network > Routing > Virtual Routers > Edit (for trust-vr) > Edit BGP Instance > Neighbors: Enter the following, and then click **Add**:

AS Number: 99

Remote IP: 2.3.3.1

Outgoing Interface: tunnel.1

Network > Routing > Virtual Routers > Edit (for trust-vr) > Edit BGP Instance > Neighbors: Enter the following, and then click **Add**:

AS Number: 99

Remote IP: 3.4.4.1

Outgoing Interface: tunnel.1

5. Policy

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book: (select), Any

Destination Address:

Address Book: (select), Any

Service: ANY

Action: Permit

CLI (NetScreen-A)

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24

set interface tunnel.1 zone untrust
set interface tunnel.1 ip 10.0.0.1/30
```

2. VPNs

```
set ike gateway peer1 address 2.2.2.2 outgoing-interface ethernet3 preshare
    netscreen1 sec-level compatible
set vpn vpn1 gateway peer1 sec-level compatible
set vpn vpn1 bind interface tunnel.1
set vpn vpn1 proxy-id local-ip 0.0.0.0/0 remote-ip 0.0.0.0/0 any
set vpn vpn1 monitor rekey
```

```
set ike gateway peer2 address 3.3.3.3 outgoing-interface ethernet3 preshare
    netscreen2 sec-level compatible
set vpn vpn2 gateway peer2 sec-level compatible
set vpn vpn2 bind interface tunnel.1
set vpn vpn2 proxy-id local-ip 0.0.0.0/0 remote-ip 0.0.0.0/0 any
set vpn vpn2 monitor rekey
```

3. Static Routes

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
set vrouter trust-vr route 2.3.3.1/32 interface tunnel.1 gateway 2.3.3.1
set vrouter trust-vr route 2.4.4.1/32 interface tunnel.1 gateway 2.4.4.1
```

4. Dynamic Routing

```
ns-> set vrouter trust-vr protocol bgp 99
ns-> set vrouter trust-vr protocol bgp enable
ns-> set interface tunnel.1 protocol bgp
ns-> set vrouter trust-vr
ns(trust-vr)-> set protocol bgp
ns(trust-vr/bgp)-> set neighbor 2.3.3.1 remote-as 99 outgoing interface
    tunnel.1
ns(trust-vr/bgp)-> set neighbor 2.3.3.1 enable
ns(trust-vr/bgp)-> set neighbor 3.4.4.1 remote-as 99 outgoing interface
    tunnel.1
ns(trust-vr/bgp)-> set neighbor 3.4.4.1 enable
ns(trust-vr/bgp)-> exit
ns(trust-vr)-> exit
```

5. Policy

```
set policy from trust to untrust any any any permit
save
```

Peer1

The following configuration is what the remote admin for the NetScreen device at the peer1 site must enter to create a VPN tunnel to NetScreen-A at the corporate site. The remote admin configures the NetScreen device to permit inbound traffic from the corporate site. He also configures the NetScreen device to communicate internal routes to its BGP neighbor through vpn1.

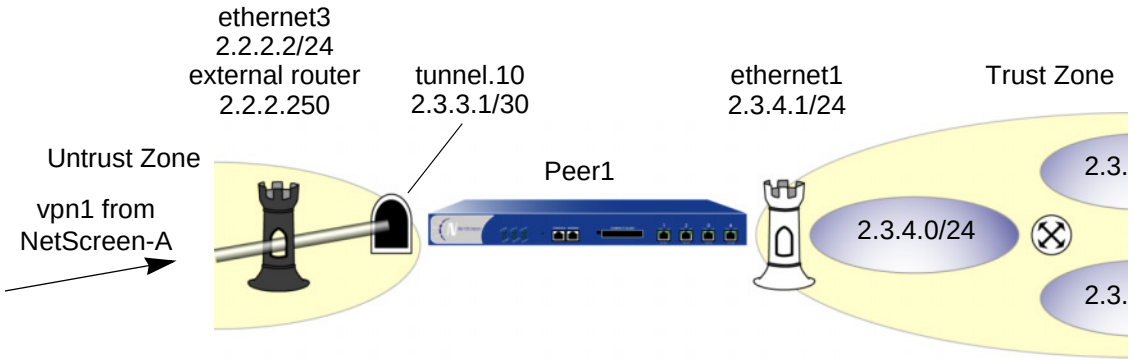

WebUI (Peer 1)

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **OK**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 2.3.4.1/24

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 2.2.2.2/24

Network > Interfaces > New Tunnel IF: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.10

Zone (VR): Untrust (trust-vr)

Fixed IP: (select)

IP Address / Netmask: 2.3.3.1/30

2. Address

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: corp

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.0/24

Zone: Untrust

3. VPN

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: vpn1

Security Level: Compatible

Remote Gateway: Create a Simple Gateway: (select)

Gateway Name: corp

Type: Static IP: (select), Address/Hostname: 1.1.1.1

Preshared Key: netscreen1

Security Level: Compatible

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Bind to: Tunnel Interface, tunnel.10

Proxy-ID: (select)

Local IP / Netmask: 0.0.0.0/0

Remote IP / Netmask: 0.0.0.0/0

Service: ANY

4. Static Routes

Network > Routing > Routing Entries > (trust-vr) New: Enter the following, and then click **OK**:

Network Address / Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 2.2.2.250

Metric: 1

Network > Routing > Routing Entries > (trust-vr) New: Enter the following, and then click **OK**:

Network Address / Netmask: 10.1.1.0/24

Gateway: (select)

Interface: tunnel.10

Gateway IP Address: 0.0.0.0

Metric: 1

5. Dynamic Routing

Network > Routing > Virtual Routers > Edit (for trust-vr) > Create BGP Instance: Enter the following, and then click **OK**:

AS Number (required): 99

BGP Enabled: (select)

Network > Interfaces > Edit (for tunnel.10) > BGP: Select the **Protocol BGP** check box, and then click **OK**.

Network > Routing > Virtual Routers > Edit (for trust-vr) > Edit BGP Instance > Neighbors: Enter the following, and then click **Add**:

AS Number: 99

Remote IP: 10.0.0.1

Outgoing Interface: tunnel.10

6. Policy

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), corp

Destination Address:

Address Book Entry: (select), Any

Service: ANY

Action: Permit

CLI (Peer1)

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 2.3.4.1/24

set interface ethernet3 zone untrust
set interface ethernet3 ip 2.2.2.2/24

set interface tunnel.10 zone untrust
set interface tunnel.10 ip 2.3.3.1/30
```

2. Address

```
set address untrust corp 10.1.1.0/24
```

3. VPN

```
set ike gateway corp address 1.1.1.1 outgoing-interface ethernet3 preshare
    netscreen1 sec-level compatible
set vpn vpn1 gateway corp sec-level compatible
set vpn vpn1 bind interface tunnel.10
set vpn vpn1 proxy-id local-ip 0.0.0.0/0 remote-ip 0.0.0.0/0 any
```

4. Static Routes

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 2.2.2.250 metric 1
set vrouter trust-vr route 10.1.1.0/24 interface tunnel.10 metric 1
```

5. Dynamic Routing

```
ns-> set vrouter trust-vr protocol bgp 99
ns-> set vrouter trust-vr protocol bgp enable
ns-> set interface tunnel.10 protocol bgp
ns-> set vrouter trust-vr
ns(trust-vr)-> set protocol bgp
ns(trust-vr/bgp)-> set neighbor 10.0.0.1 remote-as 99 outgoing interface
    tunnel.10
ns(trust-vr/bgp)-> set neighbor 10.0.0.1 enable
ns(trust-vr/bgp)-> exit
ns(trust-vr)-> exit
```

6. Policy

```
set policy from untrust to trust corp any any permit
save
```

Peer2

The following configuration is what the remote admin for the NetScreen device at the peer2 site must enter to create a VPN tunnel to NetScreen-A at the corporate site. The remote admin configures the NetScreen device to permit inbound traffic from the corporate site. He also configures the NetScreen device to communicate internal routes to its BGP neighbor through vpn2.

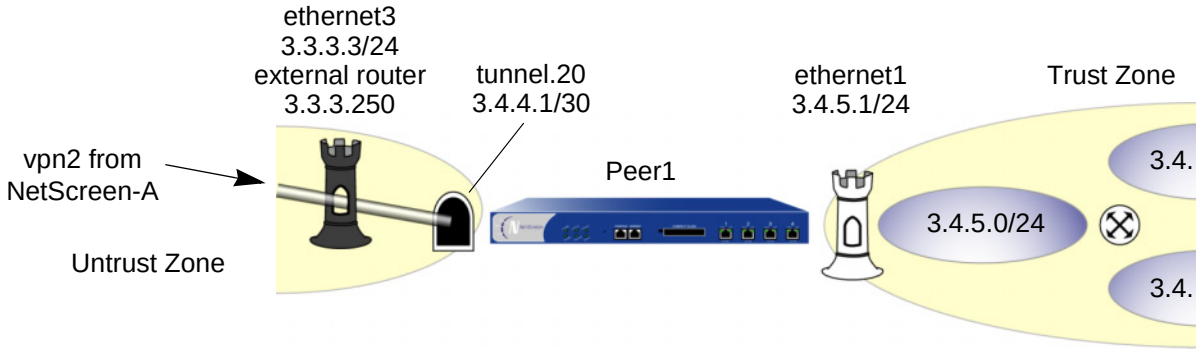

WebUI (Peer 2)

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **OK**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 2.3.4.1/24

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 3.3.3.3/24

Network > Interfaces > New Tunnel IF: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.20

Zone (VR): Untrust (trust-vr)

Fixed IP: (select)

IP Address / Netmask: 3.4.4.1/30

2. Address

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: corp

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.0/24

Zone: Untrust

3. VPN

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: vpn2

Security Level: Compatible

Remote Gateway: Create a Simple Gateway: (select)

Gateway Name: corp

Type: Static IP: (select), Address/Hostname: 1.1.1.1

Preshared Key: netscreen2

Security Level: Compatible

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Bind to: Tunnel Interface, tunnel.20

Proxy-ID: (select)

Local IP / Netmask: 0.0.0.0/0

Remote IP / Netmask: 0.0.0.0/0

Service: ANY

4. Static Routes

Network > Routing > Routing Entries > (trust-vr) New: Enter the following, and then click **OK**:

Network Address / Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 3.3.3.250

Metric: 1

Network > Routing > Routing Entries > (trust-vr) New: Enter the following, and then click **OK**:

Network Address / Netmask: 10.1.1.0/24

Gateway: (select)

Interface: tunnel.20

Gateway IP Address: 0.0.0.0

Metric: 1

5. Dynamic Routing

Network > Routing > Virtual Routers > Edit (for trust-vr) > Create BGP Instance: Enter the following, and then click **OK**:

AS Number (required): 99

BGP Enabled: (select)

Network > Interfaces > Edit (for tunnel.20) > BGP: Select the **Protocol BGP** check box, and then click **OK**.

Network > Routing > Virtual Routers > Edit (for trust-vr) > Edit BGP Instance > Neighbors: Enter the following, and then click **Add**:

AS Number: 99

Remote IP: 10.0.0.1

Outgoing Interface: tunnel.20

6. Policy

Policies > (From: Untrust, To: Trust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), corp

Destination Address:

Address Book Entry: (select), Any

Service: ANY

Action: Permit

CLI (Peer2)

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 3.4.5.1/24
```

```
set interface ethernet3 zone untrust
set interface ethernet3 ip 3.3.3.3/24

set interface tunnel.20 zone untrust
set interface tunnel.20 ip 3.4.4.1/30
```

2. Address

```
set address untrust corp 10.1.1.0/24
```

3. VPN

```
set ike gateway corp address 1.1.1.1 outgoing-interface ethernet3 preshare
    netscreen2 sec-level compatible
set vpn vpn1 gateway corp sec-level compatible
set vpn vpn1 bind interface tunnel.20
set vpn vpn1 proxy-id local-ip 0.0.0.0/0 remote-ip 0.0.0.0/0 any
```

4. Static Routes

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 3.3.3.250 metric 1
set vrouter trust-vr route 10.1.1.0/24 interface tunnel.20 metric 1
```

5. Dynamic Routing

```
ns-> set vrouter trust-vr protocol bgp 99
ns-> set vrouter trust-vr protocol bgp enable
ns-> set interface tunnel.20 protocol bgp
ns-> set vrouter trust-vr
ns(trust-vr)-> set protocol bgp
ns(trust-vr/bgp)-> set neighbor 10.0.0.1 remote-as 99 outgoing interface
    tunnel.20
ns(trust-vr/bgp)-> set neighbor 10.0.0.1 enable
ns(trust-vr/bgp)-> exit
ns(trust-vr)-> exit
```

6. Policy

```
set policy from untrust to trust corp any any permit
save
```

REDUNDANT VPN GATEWAYS

The NetScreen redundant gateway feature provides a solution for continuous VPN connectivity during and after a site-to-site failover. You can create a VPN group to provide a set of up to four redundant gateways to which policy-based site-to-site or site-to-site dynamic peer AutoKey IKE IPSec[19] VPN tunnels can connect. When the NetScreen device first receives traffic matching a policy referencing a VPN group, it performs Phase 1 and Phase 2 IKE negotiations with all members in that group. The NetScreen device sends data through the VPN tunnel to the gateway with the highest priority, or "weight", in the group. For all other gateways in the group, the NetScreen device maintains the Phase 1 and 2 SAs and keeps the tunnels active by sending IKE keepalive packets through them. If the active VPN tunnel fails, the tunnel can fail over to the tunnel and gateway with the second highest priority in the group.

Note: *This scheme assumes that the sites behind the redundant gateways are connected so that data is mirrored among hosts at all sites. Furthermore, each site—being dedicated to high availability (HA)—has a redundant cluster of NetScreen devices operating in HA mode. Therefore, the VPN failover threshold must be set higher than the device failover threshold or VPN failovers might occur unnecessarily.*

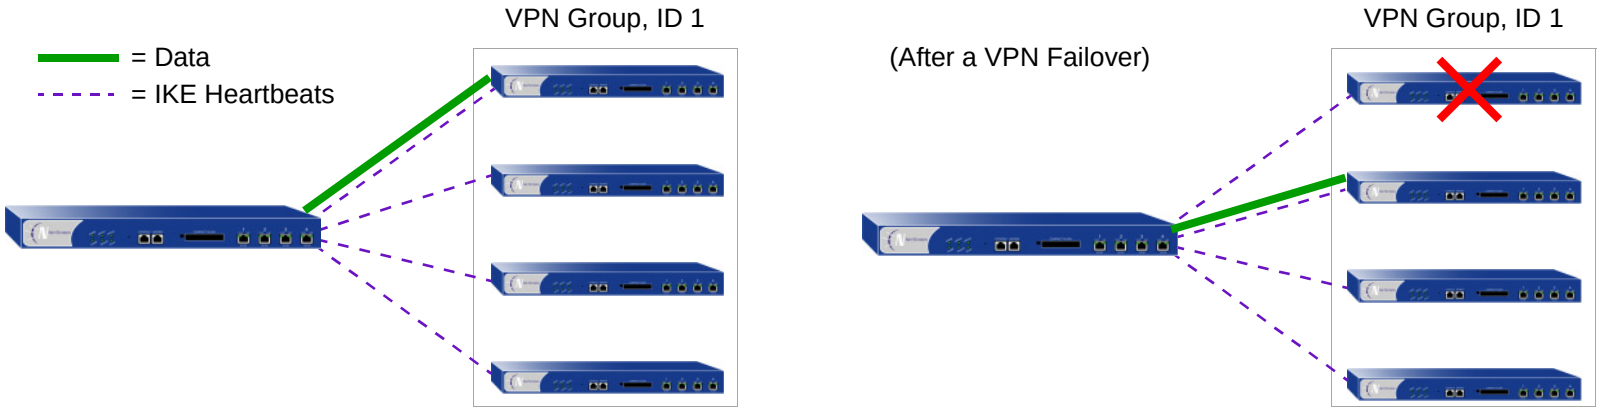

19. VPN groups do not support L2TP, L2TP-over-IPSec, dialup, Manual Key, or route-based VPN tunnel types. In a Site-to-Site Dynamic Peer arrangement, the NetScreen device monitoring the VPN group must be the one whose untrust IP address is dynamically assigned, while the untrust IP addresses of the VPN group members must be static.

VPN Groups

A VPN group is a set of VPN tunnel configurations for up to four targeted remote gateways. The Phase 1 and Phase 2 security association (SA) parameters for each tunnel in a group can be different or identical (except for the IP address of the remote gateway, which obviously must be different). The VPN group has a unique ID number, and each member in the group is assigned a unique weight to indicate its place in rank of preference to be the active tunnel. A value of 1 indicates the lowest, or least preferred, ranking.

Note: *In this illustration, the shading symbolizes the weight of each tunnel. The darker the tunnel is shaded, the higher its priority.*

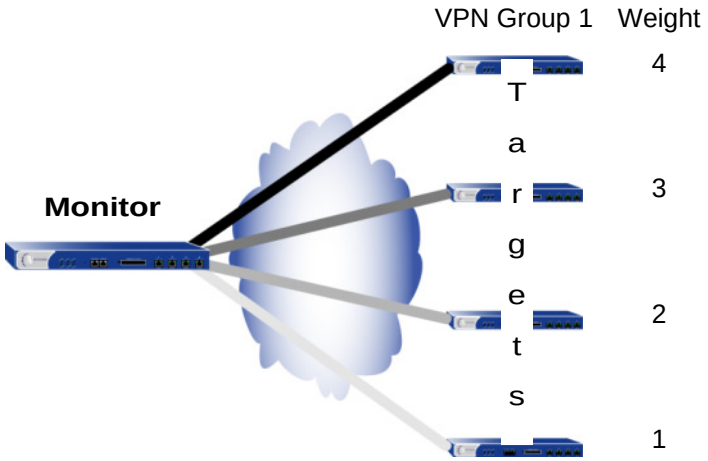

The NetScreen device communicating with VPN group members and the members themselves have a monitor-to-target relationship. The monitoring device continually monitors the connectivity and wellbeing of each targeted device. The tools that the monitor uses to do this are as follows:

- IKE heartbeats
- IKE recovery attempts

Both tools are presented in the next section, “Monitoring Mechanisms” on page 384.

Note: *The monitor-to-target relationship need not be one way. The monitoring device might also be a member of a VPN group and thus be the target of another monitoring device.*

Monitoring Mechanisms

NetScreen uses two mechanisms to monitor members of a VPN group to determine their ability to terminate VPN traffic:

- IKE heartbeats
- IKE recovery attempts

Using these two tools, plus the TCP application failover option (see “TCP SYN-Flag Checking” on page 388), NetScreen devices can detect when a VPN failover is required and shift traffic to the new tunnel without disrupting VPN service.

IKE Heartbeats

IKE heartbeats are hello messages that IKE peers send to each other under the protection of an established Phase 1 security association (SA) to confirm the connectivity and wellbeing of the other. If, for example, device_m (the “monitor”) does not receive a specified number of heartbeats (the default is 5) from device_t (the “target”), device_m concludes that device_t is down. Device_m clears the corresponding Phase 1 and Phase 2 security associations (SAs) from its SA cache and begins the IKE recovery procedure. (See “IKE Recovery Procedure” on page 385.) Device_t also clears its SAs.

Note: *The IKE heartbeats feature must be enabled on the devices at both ends of a VPN tunnel in a VPN group. If it is enabled on device_m but not on device_t, device_m suppresses IKE heartbeat transmission and generates the following message in the event log: “Heartbeats have been disabled because the peer is not sending them*.”

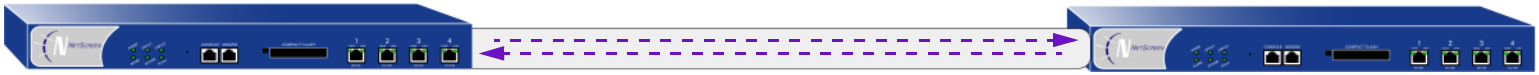

IKE Heartbeats must flow both ways through the VPN tunnel.

To define the IKE heartbeat interval and threshold for a specified VPN tunnel (the default is 5), do the following:

WebUI

VPNs > AutoKey Advanced > Gateway > Edit (for the gateway whose IKE heartbeat threshold you want to modify) > Advanced: Enter the new values in the Heartbeat Hello and Heartbeat Threshold fields, and then click **OK**.

CLI

set ike gateway *name_str* heartbeat hello *number*

set ike gateway *name_str* heartbeat threshold *number*

IKE Recovery Procedure

After the monitoring NetScreen device determines that a targeted device is down, the monitor stops sending IKE heartbeats and clears the SAs for that peer from its SA cache. After a defined interval, the monitor attempts to initiate Phase 1 negotiations with the failed peer. If the first attempt is unsuccessful, the monitor continues to attempt Phase 1 negotiations at regular intervals until negotiations are successful.

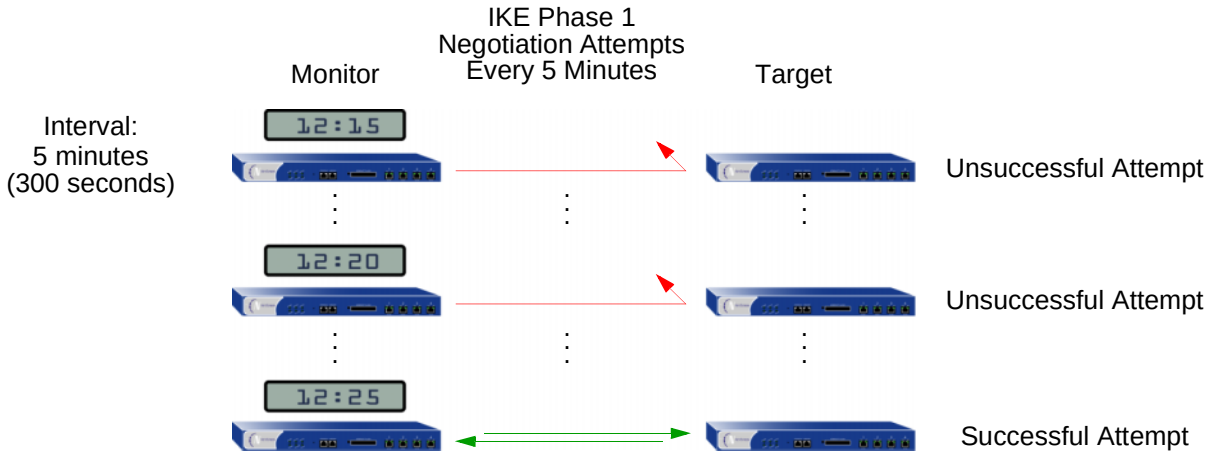

To define the IKE recovery interval for a specified VPN tunnel (the minimum setting is 60 seconds), do either of the following:

WebUI

VPNs > AutoKey Advanced > Gateway > Edit (for the gateway whose IKE reconnect interval you want to modify) > Advanced: Enter the value in seconds in the Heartbeat Reconnect field, and then click **OK**.

CLI

set ike gateway *name_str* heartbeat reconnect *number*

When a VPN group member with the highest weight fails over the tunnel to another group member and then reconnects with the monitoring device, the tunnel automatically fails back to the first member. The weighting system always causes the best ranking gateway in the group to handle the VPN data whenever it can do so.

The following illustration presents the process that a member of a VPN group undergoes when the missing heartbeats from a targeted gateway surpass the failure threshold.

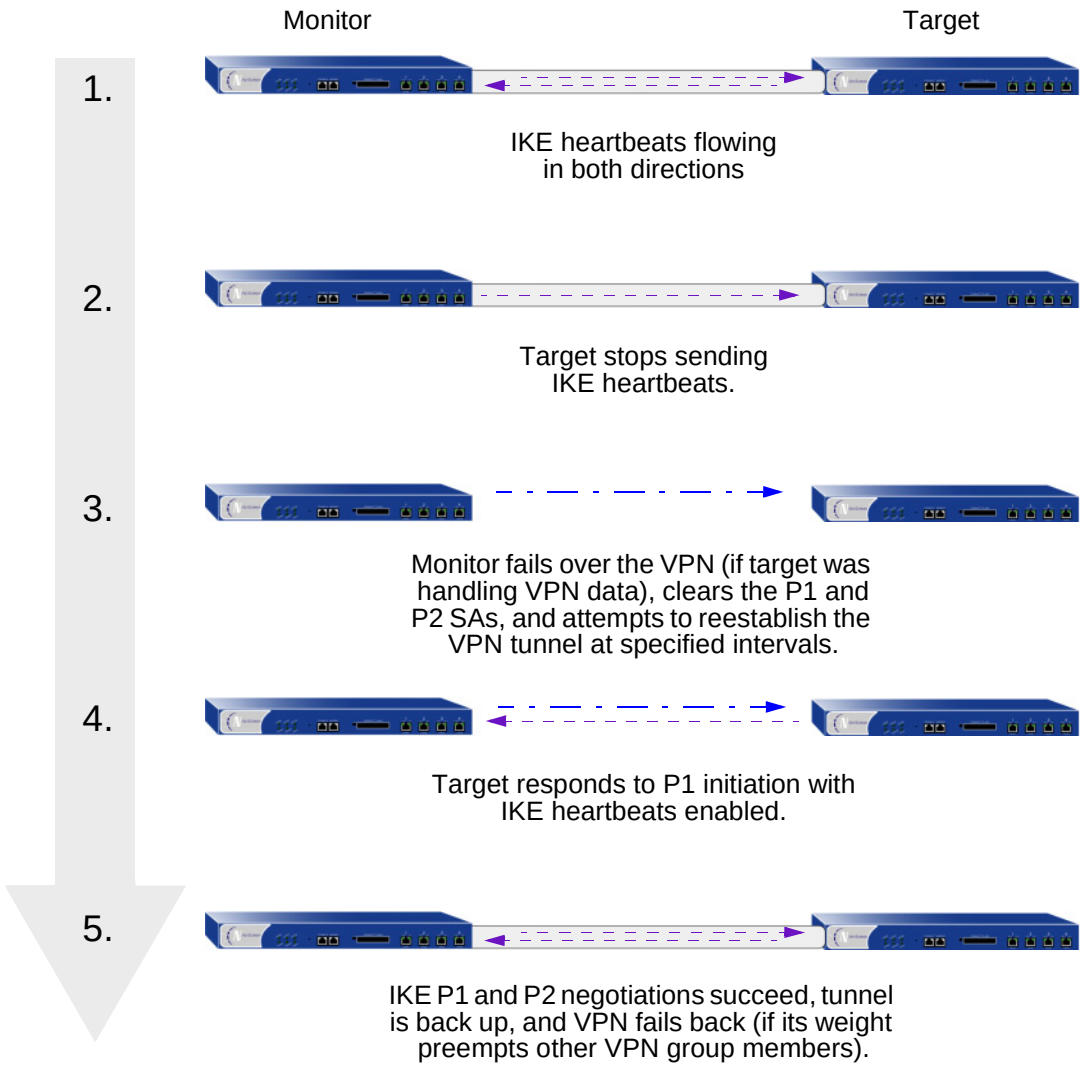

TCP SYN-Flag Checking

For a seamless VPN failover to occur, the handling of TCP sessions must be addressed. If, after a failover, the new active gateway receives a packet in an existing TCP session, the new gateway treats it as the first packet in a new TCP session and checks if the SYN flag is set in the packet header. Because this packet is really part of an existing session, it does not have the SYN flag set. Consequently, the new gateway rejects the packet. With TCP SYN flag checking enabled, all TCP applications have to reconnect after the failover occurs.

To resolve this, you can disable SYN-flag checking for TCP sessions in VPN tunnels, as follows:

WebUI

You cannot disable SYN-flag checking via the WebUI.

CLI

unset flow tcp-syn-check-in-tunnel

Note: *By default, SYN-flag checking is enabled.*

Example: Redundant VPN Gateways

In this example, a corporate site has one VPN tunnel to a data center and a second tunnel to a backup data center. All the data is mirrored via a leased line connection between the two data center sites. The data centers are physically separate to provide continuous service even in the event of a catastrophic failure such as an all-day power outage or a natural disaster.

The device location and name, the physical interfaces and their IP addresses for the Trust and Untrust zones, and the VPN group ID and weight for each NetScreen device are as follows:

Device Location	Device Name	Physical Interface and IP Address (Trust Zone)	Physical Interface, IP Address, Default Gateway (Untrust Zone)	VPN Group ID and Weight
Corporate	Monitor1	ethernet1, 10.10.1.1/24	ethernet3, 1.1.1.1/24, (GW) 1.1.1.2	– –
Data Center (Primary)	Target1	ethernet1, 10.1.1.1/16	ethernet3, 2.2.2.1/24, (GW) 2.2.2.2	ID = 1, Weight = 2
Data Center (Backup)	Target2	ethernet1, 10.1.1.1/16	ethernet3, 3.3.3.1/24, (GW) 3.3.3.2	ID = 1, Weight =1

Note: *The internal address space at both data center sites must be identical.*

All security zones are in the trust-vr routing domain. All the Site-to-Site AutoKey IKE tunnels use the security level predefined as “Compatible” for both Phase 1 and Phase 2 proposals. Preshared keys authenticate the participants.

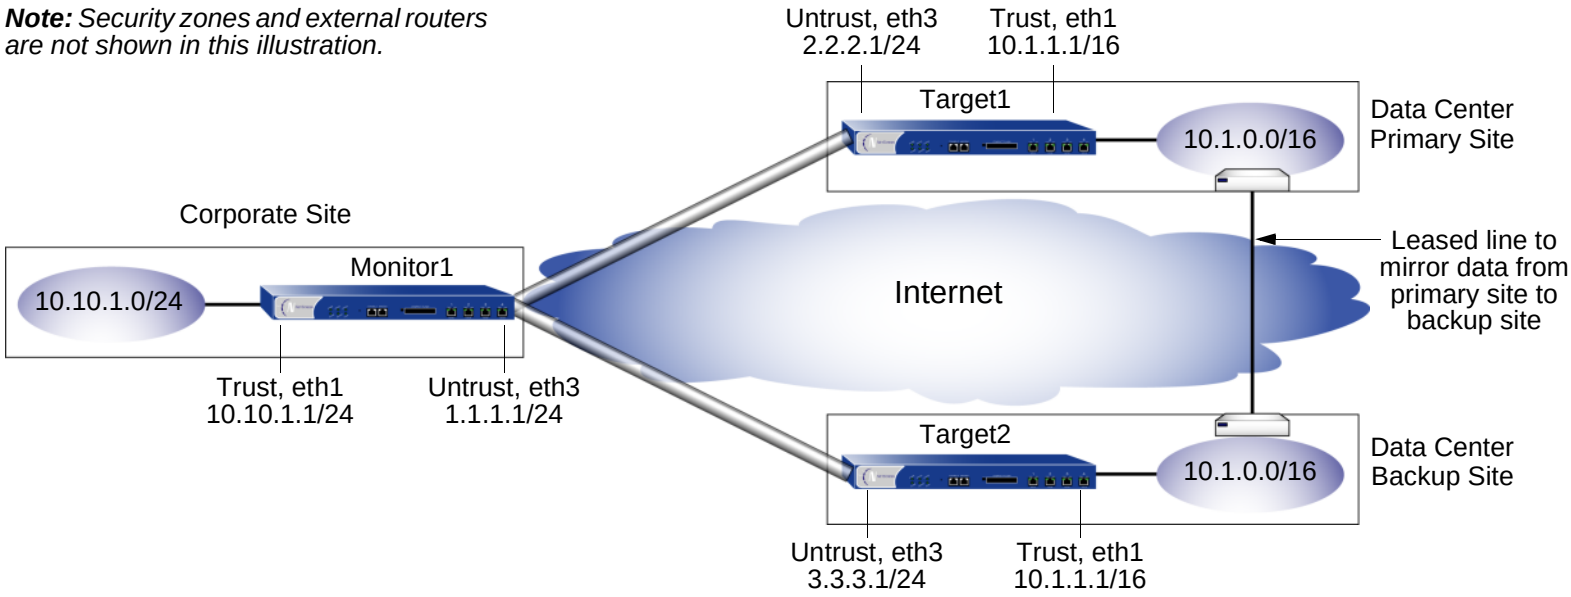

WebUI (Monitor1)

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.10.1.1/24

Enter the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

2. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: in_trust

IP Address/Domain Name:

IP/Netmask: (select), 10.10.1.0/24

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: data_ctr

IP Address/Domain Name:

IP/Netmask: (select), 10.1.0.0/16

Zone: Untrust

3. VPNs

VPNs > AutoKey Advanced > VPN Group: Enter 1 in the VPN Group ID field, and then click **Add**.

VPNs > AutoKey Advanced > Gateway > New: Enter the following, and then click **OK**:

Gateway Name: target1

Security Level: Compatible

Remote Gateway Type: Static IP Address: (select), IP Address: 2.2.2.1

Preshared Key: SLi1yoo129

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Gateway configuration page:

Security Level: Compatible

Mode (Initiator): Main (ID Protection)

Heartbeat:

Hello: 3 Seconds

Reconnect: 60 seconds

Threshold: 5

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: to_target1

Security Level: Compatible

Remote Gateway: Predefined: (select), target1

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

VPN Group: VPN Group -1

Weight: 2

VPNs > AutoKey Advanced > Gateway > New: Enter the following, and then click **OK**:

Gateway Name: target2

Security Level: Compatible

Remote Gateway Type: Static IP Address: (select), IP Address: 3.3.3.1

Preshared Key: CMFwb7oN23

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Gateway configuration page:

Security Level: Compatible

Mode (Initiator): Main (ID Protection)

Heartbeat:

Hello: 3 Seconds

Reconnect: 60 seconds

Threshold: 5

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: to_target2

Security Level: Compatible

Remote Gateway: Predefined: (select), target2

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

VPN Group: VPN Group -1

Weight: 1

4. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 1.1.1.2(untrust)

5. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), in_trust

Destination Address:

Address Book Entry: (select), data_ctr

Service: ANY

Action: Tunnel

VPN: VPN Group-1

Modify matching bidirectional VPN policy: (select)

Position at Top: (select)

WebUI (Target1)

1. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/16

Enter the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 2.2.2.1/24

2. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: in_trust

IP Address/Domain Name:

IP/Netmask: (select), 10.1.0.0/16

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: corp

IP Address/Domain Name:

IP/Netmask: (select), 10.10.1.0/24

Zone: Untrust

3. VPN

VPNs > AutoKey Advanced > Gateway > New: Enter the following, and then click **OK**:

Gateway Name: monitor1

Security Level: Compatible

Remote Gateway Type:

Static IP Address: (select), IP Address/Hostname: 1.1.1.1

Preshared Key: SLi1yoo129

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic Gateway configuration page:

Security Level: Compatible

Mode (Initiator): Main (ID Protection)

Heartbeat:

Hello: 3 Seconds

Reconnect: 0 seconds

VPN > AutoKey IKE > New: Enter the following, and then click **OK**:

Name: to_monitor1

Security Level: Compatible

Remote Gateway: Predefined: (select), monitor1

4. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 2.2.2.2

5. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), in_trust

Destination Address:

Address Book Entry: (select), corp

Service: ANY

Action: Tunnel

Tunnel VPN: monitor1

Modify matching bidirectional VPN policy: (select)

Position at Top: (select)

WebUI (Target2)

Note: *Follow the Target1 configuration steps to configure Target2, but define the Untrust zone interface IP address as 3.3.3.1/24, the default gateway IP address as 3.3.3.2, and use CMFwb7oN23 to generate the preshared key.*

CLI (Monitor1)

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.10.1.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
```

2. Addresses

```
set address trust in_trust 10.10.1.0/24
set address untrust data_ctr 10.1.0.0/16
```

3. VPNs

```
set ike gateway target1 address 2.2.2.1 main outgoing-interface ethernet3
    preshare SLi1yoo129 sec-level compatible
set ike gateway target1 heartbeat hello 3
set ike gateway target1 heartbeat reconnect 60
set ike gateway target1 heartbeat threshold 5
set vpn to_target1 gateway target1 sec-level compatible
set ike gateway target2 address 3.3.3.1 main outgoing-interface ethernet3
    preshare CMFwb7oN23 sec-level compatible
set ike gateway target2 heartbeat hello 3
set ike gateway target2 heartbeat reconnect 60
set ike gateway target2 heartbeat threshold 5
set vpn to_target2 gateway target2 sec-level compatible
set vpn-group id 1 vpn to_target1 weight 2
set vpn-group id 1 vpn to_target2 weight 1
unset flow tcp-syn-check-in-tunnel
```

4. Route

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.2
```

5. Policies

```
set policy top from trust to untrust in_trust data_ctr any tunnel “vpn-group 1”
set policy top from untrust to trust data_ctr in_trust any tunnel “vpn-group 1”
save
```

CLI (Target1)

1. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.1.1.1/16
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 2.2.2.1/24
```

2. Addresses

```
set address trust in_trust 10.1.0.0/16
set address untrust corp 10.10.1.0/24
```

3. VPN

```
set ike gateway monitor1 address 1.1.1.1 main outgoing-interface ethernet3
    preshare SLi1yoo129 sec-level compatible
set ike gateway monitor1 heartbeat hello 3
set ike gateway monitor1 heartbeat threshold 5
set vpn to_monitor1 gateway monitor1 sec-level compatible
```

4. Route

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 2.2.2.2
```

5. Policies

```
set policy top from trust to untrust in_trust corp any tunnel vpn to_monitor
set policy top from untrust to trust corp in_trust any tunnel vpn to_monitor
save
```

CLI (Target2)

Note: *Follow the Target1 configuration steps to configure Target2, but define the Untrust zone interface IP address as 3.3.3.1/24, the default gateway IP address as 3.3.3.2, and use CMFwb7oN23 to generate the preshared key.*

BACK-TO-BACK VPNS

You can enforce interzone policies at the hub site for traffic passing from one VPN tunnel to another by putting the spoke sites in different zones[20]. Because they are in different zones, the NetScreen device at the hub must do a policy lookup before routing the traffic from one tunnel to another. Thus you can control the traffic flowing via the VPN tunnels between the spoke sites. Such an arrangement is called back-to-back VPNs.

Back-to-Back VPNs

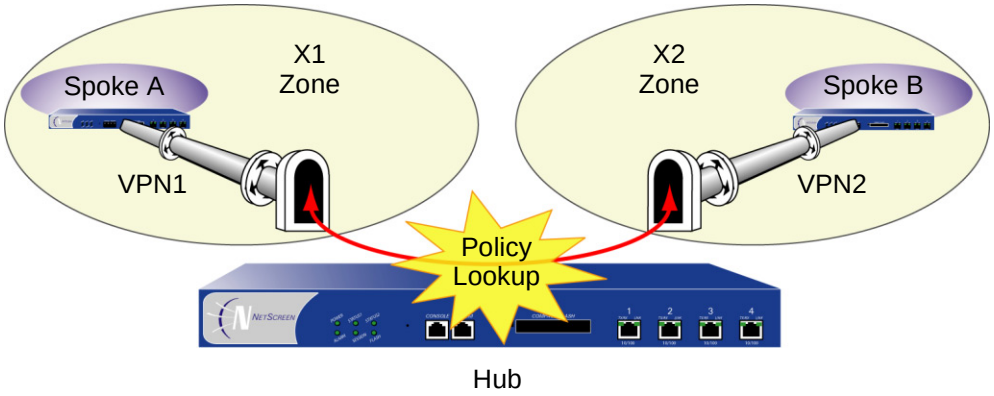

20. Optionally, you can enable intrazone blocking and define an intrazone policy to control traffic between the two tunnel interfaces within the same zone.

A few benefits of back-to-back VPNs:

- You can conserve the number of VPNs you need to create. For example, perimeter site A can link to the hub, and to perimeter sites B, C, D…, but A only has to set up one VPN tunnel. Especially for NetScreen-5XP users, who can use a maximum of ten VPN tunnels concurrently, applying the hub-and-spoke method dramatically increases their VPN options and capabilities.
- The administrator at the hub device can completely control VPN traffic between perimeter sites. For example,
  - He or she might permit only HTTP traffic to flow from sites A to B, but allow any kind of traffic to flow from B to A.
  - He or she can allow traffic originating from A to reach C, but deny traffic originating from C to reach A.
  - He or she can allow a specific host at A to reach the entire D network, while allowing only a specific host at D to reach a different host at A.
- The administrator at the hub device can completely control outbound traffic from all perimeter networks. At each perimeter site, there must first be a policy that tunnels all outbound traffic through the spoke VPNs to the hub; for example: **set policy top from trust to untrust any any any tunnel vpn** *name_str* (where *name_str* defines the specific VPN tunnel from each perimeter site to the hub). At the hub, the administrator can control Internet access, allowing certain kinds of traffic (such as HTTP only), performing URL blocking on undesirable Web sites, and so on.
- Regional hubs can be used and interconnected via spoke tunnels, allowing spoke sites in one region to reach spoke sites in another.

Example: Back-to-Back VPNs

The following example is similar to "Example: Hub-and-Spoke VPNs" on page 413 except that the NetScreen device at the hub site in New York performs policy checking on the traffic it routes between the two tunnels to the branch offices in Tokyo and Paris. By putting each remote site in a different zone, you control the VPN traffic at the hub.

The Tokyo LAN address is in the user-defined X1 zone, and the Paris LAN address is in the user-defined X2 zone. Both zones are in the Trust-VR routing domain.

Note: *To create user-defined zones, you must first obtain and load a zone software key on the NetScreen device.*

You bind the VPN1 tunnel to the tunnel.1 interface and the VPN2 tunnel to the tunnel.2 interface. Although you do not assign IP addresses to the X1 and X2 zone interfaces, you do give addresses to both tunnel interfaces. Routes for these interfaces automatically appear in the Trust-VR routing table. By putting the IP address for a tunnel interface in the same subnet as that of the destination, traffic destined for that subnet is routed to the tunnel interface.

The outgoing interface is ethernet3, which is bound to the Untrust zone. As you can see in the illustration below, both tunnels terminate in the Untrust zone; however, the endpoints for the traffic that makes use of the tunnels are in the X1 and X2 zones. The tunnels use AutoKey IKE, with preshared keys. You select the security level predefined as “Compatible” for both Phase 1 and Phase 2 proposals. You bind the Untrust zone to the untrust-vr. Because the tunnels are route-based (that is, the correct tunnel is determined by routing, not by a tunnel name specified in a policy), proxy IDs are included in the configuration of each tunnel.

Note: *Only the configuration for the NetScreen device at the hub site is provided below.*

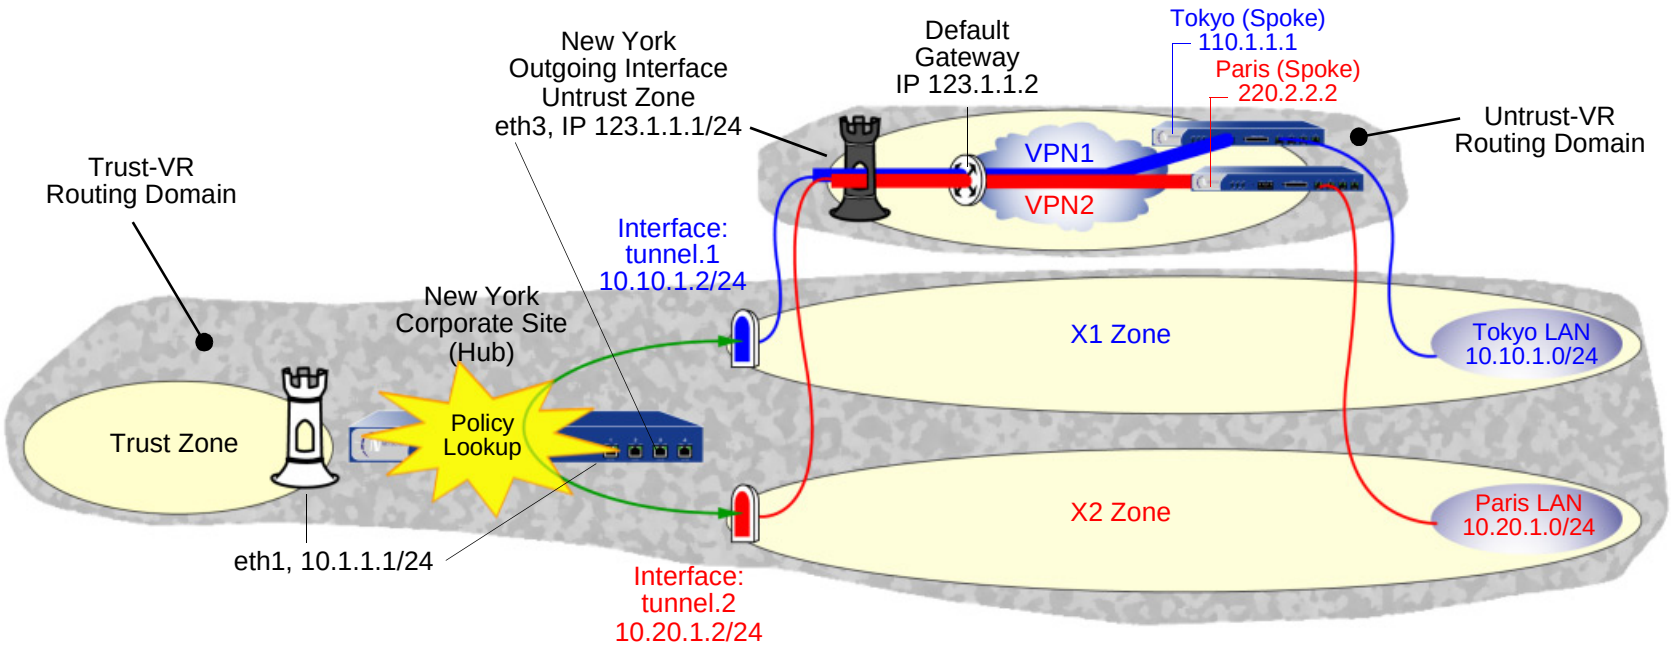

WebUI

1. Security Zones and Virtual Routers

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

IP Address/Netmask: 0.0.0.0/0

Manage IP: 0.0.0.0

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Null

Network > Zones > Edit (for Untrust): Enter the following, and then click **OK**:

Virtual Router Name: untrust-vr

Block Intra-Zone Traffic: (select)

Network > Zones > New: Enter the following, and then click **OK**:

Zone Name: X1

Virtual Router Name: trust-vr

Block Intra-Zone Traffic: (select)

Network > Zones > New: Enter the following, and then click **OK**:

Name: X2

Virtual Router Name: trust-vr

Block Intra-Zone Traffic: (select)

2. Interfaces

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 123.1.1.1/24

Network > Interfaces > New Tunnel IF: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.1

Zone (VR): X1 (trust-vr)

Fixed IP: (select)

IP Address / Netmask: 10.10.1.2/24

Network > Interfaces > New Tunnel IF: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.2

Zone (VR): X2 (trust-vr)

Fixed IP: (select)

IP Address / Netmask: 10.20.1.2/24

3. VPN for Tokyo Office

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: VPN1

Security Level: Compatible

Remote Gateway: Create a Simple Gateway: (select)

Gateway Name: Tokyo

Type: Static IP: (select), Address/Hostname: 110.1.1.1

Preshared Key: netscreen1

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Proxy-ID: (select)[21]

Local IP / Netmask: 10.20.1.0/24

Remote IP / Netmask: 10.10.1.0/24

Service: ANY

21. When configuring the VPN tunnel on the NetScreen device protecting the Tokyo and Paris offices, do either of the following:
(Route-based VPN) Select the **Enable Proxy-ID** check box and enter **10.10.1.0/24** (Tokyo) and **10.20.1.0/24** (Paris) for the Local IP and Netmask, and **10.20.1.0/24** (Tokyo) and **10.10.1.0/24** (Paris) for the Remote IP and Netmask.
(Policy-based VPN) Make an entry in the Trust zone address book for 10.10.1.0/24 (Tokyo) and 10.20.1.0/24 (Paris) and another in the Untrust zone address book for 10.20.1.0/24 (Tokyo) and 10.10.1.0/24 (Paris). Use those as the source and destination addresses in the policy referencing the VPN tunnel to the hub site.

4. VPN for Paris Office

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: VPN2

Security Level: Compatible

Remote Gateway: Create a Simple Gateway: (select)

Gateway Name: Paris

Type: Static IP: (select), Address/Hostname: 220.2.2.2

Preshared Key: netscreen2

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Proxy-ID: (select)

Local IP / Netmask: 10.10.1.0/24

Remote IP / Netmask: 10.20.1.0/24

Service: ANY

5. Routes

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 0.0.0.0/0

Next Hop Virtual Router Name: (select), untrust-vr

Network > Routing > Routing Entries > untrust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 123.1.1.2

6. Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Tokyo LAN

IP Address/Domain Name:

IP/Netmask: (select), 10.10.1.0/24

Zone: X1

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Paris LAN

IP Address/Domain Name:

IP/Netmask: (select), 10.20.1.0/24

Zone: X2

7. Policies

Policy > (From: X1, To: X2) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Tokyo LAN

Destination Address:

Address Book Entry: (select), Paris LAN

Service: ANY

Action: Permit

Position at Top: (select)

Policy > (From: X2, To: X1) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Paris LAN

Destination Address:

Address Book Entry: (select), Tokyo LAN

Service: ANY

Action: Permit

Position at Top: (select)

CLI

1. Security Zones and Virtual Routers

```
unset interface ethernet3 ip
unset interface ethernet3 zone
set zone untrust vrouter untrust-vr
set zone untrust block
set zone name X1
set zone x1 vrouter trust-vr
set zone x1 block
set zone name x2
set zone x2 vrouter trust-vr
set zone x2 block
```

2. Interfaces

```
set interface ethernet3 zone untrust
set interface ethernet3 ip 123.1.1.1/24
set interface tunnel.1 zone x1
set interface tunnel.1 ip 10.10.1.2/24
set interface tunnel.2 zone x2
set interface tunnel.2 ip 10.20.1.2/24
```

3. VPN for Tokyo Office

```
set ike gateway Tokyo address 110.1.1.1 outgoing-interface ethernet3 preshare
    netscreen1 sec-level compatible
set vpn VPN1 gateway Tokyo sec-level compatible
set vpn VPN1 bind interface tunnel.1
set vpn VPN1 proxy-id local-ip 10.20.1.0/24 remote-ip 10.10.1.0/24 any[22]
```

4. VPN for Paris Office

```
set ike gateway Paris address 220.2.2.2 outgoing-interface ethernet3 preshare
    netscreen2 sec-level compatible
```

```
set vpn VPN2 gateway Paris sec-level compatible
set vpn VPN2 bind interface tunnel.2
set vpn VPN2 proxy-id local-ip 10.10.1.0/24 remote-ip 10.20.1.0/24 any
```

5. Routes

```
set vrouter trust-vr route 0.0.0.0/0 vrouter untrust-vr
set vrouter untrust-vr route 0.0.0.0/0 interface ethernet3 gateway 123.1.1.2
```

6. Addresses

```
set address x1 “Tokyo LAN” 10.10.1.0/24
set address x2 “Paris LAN” 10.20.1.0/24
```

7. Policies

```
set policy top from x1 to x2 “Tokyo LAN” “Paris LAN” any permit[23]
set policy top from x2 to x1 “Paris LAN” “Tokyo LAN” any permit
save
```

22. When configuring the VPN tunnel on the NetScreen device protecting the Tokyo and Paris offices, do either of the following:
(Route-based VPN) Enter the following commands: **set vpn VPN1 proxy-id local-ip 10.20.1.0/24 remote-ip 10.10.1.0/24** (Tokyo) and **set vpn VPN1 proxy-id local-ip 10.10.1.0/24 remote-ip 10.20.1.0/24** (Paris).
(Policy-based VPN) Make an entry in the Trust zone address book for 10.10.1.0/24 (Tokyo) and 10.20.1.0/24 (Paris) and another in the Untrust zone address book for 10.20.1.0/24 (Tokyo) and 10.10.1.0/24 (Paris). Use those as the source and destination addresses in the policies referencing the VPN tunnel to the hub site.

23. You can ignore the following message, which appears because tunnel interfaces are in NAT mode: *Warning: Some interfaces in the <zone_name> zone are in NAT mode. Traffic might not pass through them!*

HUB-AND-SPOKE VPNS

If you create two VPN tunnels that terminate at a NetScreen device, you can set up a pair of routes so that the NetScreen device directs traffic exiting one tunnel to the other tunnel. If both tunnels are contained within a single zone, you do not need to create a policy to permit the traffic to pass from one tunnel to the other. You only need to define the routes. Such an arrangement is known as hub-and-spoke VPNs.

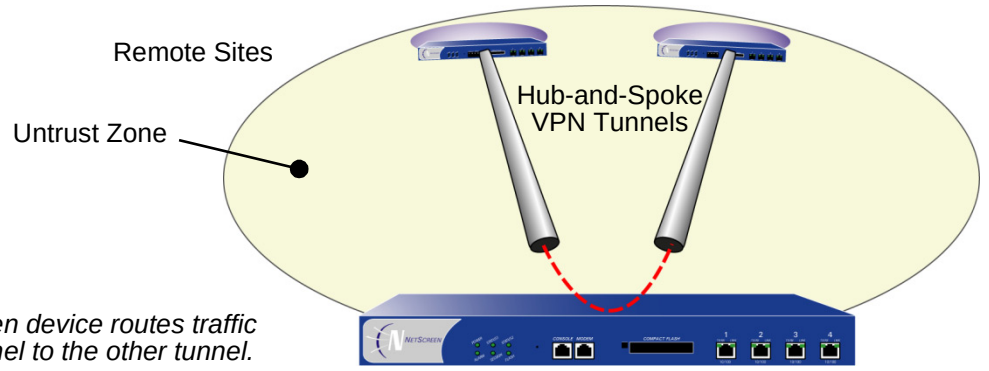

The NetScreen device routes traffic from one tunnel to the other tunnel.

You can also configure multiple VPNs in one zone and route traffic between any two tunnels.

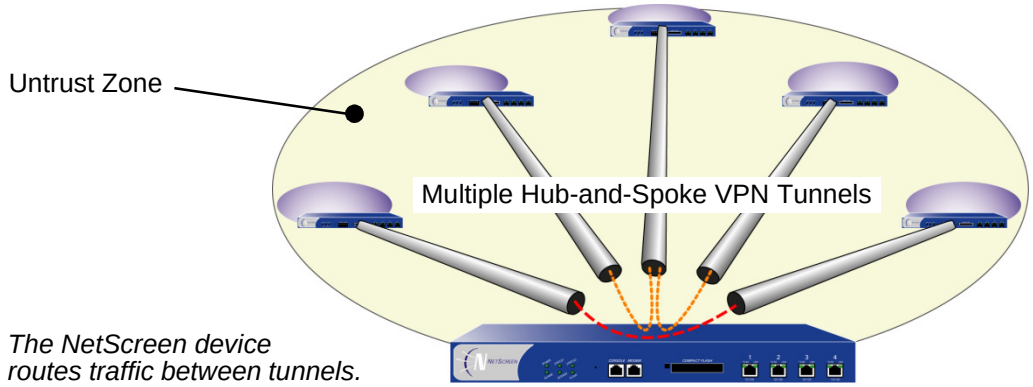

The NetScreen device routes traffic between tunnels.

Example: Hub-and-Spoke VPNs

In this example, two branch offices in Tokyo and Paris communicate with each other via a pair of VPN tunnels—VPN1 and VPN2. Each tunnel originates at the remote site and terminates at the corporate site in New York. The NetScreen device at the corporate site routes traffic exiting one tunnel into the other tunnel.

By disabling intrazone blocking, the NetScreen device at the corporate site only needs to do a route lookup—not a policy lookup—when conducting traffic from tunnel to tunnel because both remote endpoints are in the same zone (the Untrust Zone)[24].

You bind the tunnels to the tunnel interfaces—tunnel.1 and tunnel.2—which are both unnumbered. The tunnels use AutoKey IKE, with the preshared keys. You select the security level predefined as “Compatible” for both Phase 1 and Phase 2 proposals. You bind the Untrust zone to the untrust-vr. The Untrust zone interface is ethernet3.

Note: *The following configuration is for route-based VPNs. If you configure policy-based hub-and-spoke VPNs, you must use the Trust and Untrust zones in the policies; you cannot use user-defined security zones.*

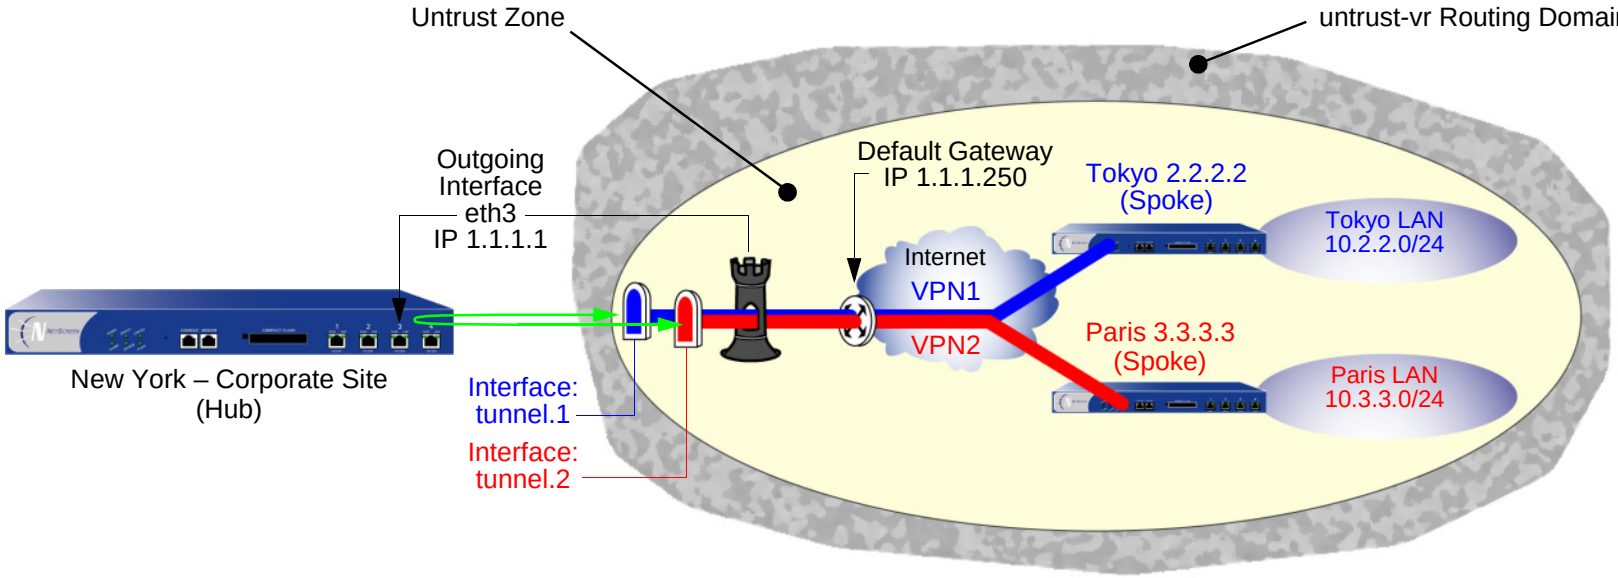

24. Optionally, you can leave intrazone blocking enabled and define an intrazone policy permitting traffic between the two tunnel interfaces.

WebUI (New York)

1. Security Zones and Virtual Routers

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

IP Address/Netmask: 0.0.0.0/0

Manage IP: 0.0.0.0

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Null

Network > Zones > Edit (for Untrust): Enter the following, and then click **OK**:

Virtual Router Name: untrust-vr

Block Intra-Zone Traffic: (clear)

2. Interfaces

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 1.1.1.1/24

Network > Interfaces > New Tunnel IF: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.1

Zone (VR): Untrust (untrust-vr)

Unnumbered: (select)

Interface: ethernet3 (untrust-vr)

Network > Interfaces > New Tunnel IF: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.2

Zone (VR): Untrust (untrust-vr)

Unnumbered: (select)

Interface: ethernet3 (untrust-vr)

3. VPN for Tokyo Office

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: VPN1

Security Level: Compatible

Remote Gateway: Create a Simple Gateway: (select)

Gateway Name: Tokyo

Type: Static IP: (select), Address/Hostname: 2.2.2.2

Preshared Key: netscreen1

Security Level: Compatible

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Proxy-ID: (select)

Local IP / Netmask: 0.0.0.0/0

Remote IP / Netmask: 0.0.0.0/0

Service: ANY

4. VPN for Paris Office

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: VPN2

Security Level: Compatible

Remote Gateway: Create a Simple Gateway: (select)

Gateway Name: Paris

Type: Static IP: (select), Address/Hostname: 3.3.3.3

Preshared Key: netscreen2

Security Level: Compatible

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Proxy-ID: (select)

Local IP / Netmask: 0.0.0.0/0

Remote IP / Netmask: 0.0.0.0/0

Service: ANY

5. Routes

Network > Routing > Routing Entries > untrust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 10.2.2.0/24

Gateway: (select)

Interface: tunnel.1

Gateway IP Address: 0.0.0.0

Network > Routing > Routing Entries > untrust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 10.3.3.0/24

Gateway: (select)

Interface: tunnel.2

Gateway IP Address: 0.0.0.0

Network > Routing > Routing Entries > untrust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 1.1.1.250

WebUI (Tokyo)

1. Security Zones and Virtual Routers

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

IP Address/Netmask: 0.0.0.0/0

Manage IP: 0.0.0.0

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Null

Network > Zones > Edit (for Untrust): Enter the following, and then click **OK**:

Virtual Router Name: untrust-vr

Block Intra-Zone Traffic: (select)

2. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.2.2.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 2.2.2.2/24

Network > Interfaces > New Tunnel IF: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.1

Zone (VR): Untrust (untrust-vr)

Unnumbered: (select)

Interface: ethernet3 (untrust-vr)

3. Address

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Paris

IP Address/Domain Name:

IP/Netmask: (select), 10.3.3.0/24

Zone: Untrust

4. VPN

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: VPN1

Security Level: Compatible

Remote Gateway: Create a Simple Gateway: (select)

Gateway Name: New York

Type: Static IP: (select), Address/Hostname: 1.1.1.1

Preshared Key: netscreen1

Security Level: Compatible

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Proxy-ID: (select)

Local IP / Netmask: 0.0.0.0/0

Remote IP / Netmask: 0.0.0.0/0

Service: ANY

5. Routes

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 0.0.0.0/0

Next Hop Virtual Router Name: (select); untrust-vr

Network > Routing > Routing Entries > untrust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 2.2.2.250

Network > Routing > Routing Entries > untrust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 10.3.3.0/24

Gateway: (select)

Interface: tunnel.1

Gateway IP Address: 0.0.0.0

6. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Paris

Service: ANY

Action: Permit

WebUI (Paris)

1. Security Zones and Virtual Routers

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

IP Address/Netmask: 0.0.0.0/0

Manage IP: 0.0.0.0

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Null

Network > Zones > Edit (for Untrust): Enter the following, and then click **OK**:

Virtual Router Name: untrust-vr

Block Intra-Zone Traffic: (select)

2. Interfaces

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.3.3.1/24

Select the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 3.3.3.3/24

Network > Interfaces > New Tunnel IF: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.1

Zone (VR): Untrust (untrust-vr)

Unnumbered: (select)

Interface: ethernet3 (untrust-vr)

3. Address

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: Tokyo

IP Address/Domain Name:

IP/Netmask: (select), 10.2.2.0/24

Zone: Untrust

4. VPN

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: VPN2

Security Level: Compatible

Remote Gateway: Create a Simple Gateway: (select)

Gateway Name: New York

Type: Static IP: (select), Address/Hostname: 1.1.1.1

Preshared Key: netscreen2

Security Level: Compatible

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Proxy-ID: (select)

Local IP / Netmask: 0.0.0.0/0

Remote IP / Netmask: 0.0.0.0/0

Service: ANY

5. Routes

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 0.0.0.0/0

Next Hop Virtual Router Name: (select); untrust-vr

Network > Routing > Routing Entries > untrust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 3.3.3.250

Network > Routing > Routing Entries > untrust-vr New: Enter the following, and then click **OK**:

Network Address / Netmask: 10.2.2.0/24

Gateway: (select)

Interface: tunnel.1

Gateway IP Address: 0.0.0.0

6. Policies

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), Any

Destination Address:

Address Book Entry: (select), Tokyo

Service: ANY

Action: Permit

CLI (New York)

1. Security Zones and Virtual Routers

```
unset interface ethernet3 ip
unset interface ethernet3 zone
set zone untrust vrouter untrust-vr
unset zone untrust block
```

2. Interfaces

```
set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
set interface tunnel.1 zone untrust
set interface tunnel.1 ip unnumbered interface ethernet3
set interface tunnel.2 zone untrust
set interface tunnel.2 ip unnumbered interface ethernet3
```

3. VPN for Tokyo Office

```
set ike gateway Tokyo address 2.2.2.2 outgoing-interface ethernet3 preshare
    netscreen1 sec-level compatible
set vpn VPN1 gateway Tokyo sec-level compatible
set vpn VPN1 bind interface tunnel.1
set vpn VPN1 proxy-id local-ip 0.0.0.0/0 remote-ip 0.0.0.0/0 any
```

4. VPN for Paris Office

```
set ike gateway Paris address 3.3.3.3 outgoing-interface ethernet3 preshare
    netscreen2 sec-level compatible
set vpn VPN2 gateway Paris sec-level compatible
set vpn VPN2 bind interface tunnel.2
set vpn VPN2 proxy-id local-ip 0.0.0.0/0 remote-ip 0.0.0.0/0 any
```

5. Routes

```
set vrouter untrust-vr route 10.2.2.0/24 interface tunnel.1
set vrouter untrust-vr route 10.3.3.0/24 interface tunnel.2
set vrouter untrust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
save
```

CLI (Tokyo)

1. Security Zones and Virtual Routers

```
unset interface ethernet3 ip
unset interface ethernet3 zone
set zone untrust vrouter untrust-vr
```

2. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.2.2.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 2.2.2.2/24

set interface tunnel.1 zone untrust
set interface tunnel.1 ip unnumbered interface ethernet3
```

3. Address

```
set address untrust Paris 10.3.3.0/24
```

4. VPN

```
set ike gateway “New York” address 1.1.1.1 outgoing-interface ethernet3
    preshare netscreen1 sec-level compatible
set vpn VPN1 gateway “New York” sec-level compatible
set vpn VPN1 bind interface tunnel.1
set vpn VPN1 proxy-id local-ip 0.0.0.0/0 remote-ip 0.0.0.0/0 any
```

5. Routes

```
set vrouter trust-vr route 0.0.0.0/0 vrouter untrust-vr
set vrouter untrust-vr route 0.0.0.0/0 interface ethernet3 gateway 2.2.2.250
set vrouter untrust-vr route 10.3.3.0/24 interface tunnel.1
```

6. Policies

```
set policy from trust to untrust any Paris any permit
set policy from untrust to trust Paris any any permit
save
```

CLI (Paris)

1. Security Zones and Virtual Routers

```
unset interface ethernet3 ip
unset interface ethernet3 zone
set zone untrust vrouter untrust-vr
```

2. Interfaces

```
set interface ethernet1 zone trust
set interface ethernet1 ip 10.3.3.1/24
set interface ethernet1 nat

set interface ethernet3 zone untrust
set interface ethernet3 ip 3.3.3.3/24

set interface tunnel.1 zone untrust
set interface tunnel.1 ip unnumbered interface ethernet3
```

3. Address

```
set address untrust Tokyo 10.2.2.0/24
```

4. VPN

```
set ike gateway “New York” address 1.1.1.1 outgoing-interface ethernet3
    preshare netscreen2 sec-level compatible
set vpn VPN2 gateway “New York” sec-level compatible
set vpn VPN2 bind interface tunnel.1
set vpn VPN2 proxy-id local-ip 0.0.0.0/0 remote-ip 0.0.0.0/0 an
```

5. Routes

```
set vrouter trust-vr route 0.0.0.0/0 vrouter untrust-vr
set vrouter untrust-vr route 0.0.0.0/0 interface ethernet3 gateway 3.3.3.250
set vrouter untrust-vr route 10.2.2.0/24 interface tunnel.1
```

6. Policies

```
set policy from trust to untrust any Tokyo any permit
set policy from untrust to trust Tokyo any any permit
save
```

Index

Symbols

3DES 8

A

AES (Advanced Encryption Standard) 8
Aggressive Mode 12
AH 3, 7
anti-replay checking 46, 54
attacks
  Replay 14
authentication
  algorithms 7, 44, 49, 53, 57
Authentication Header
  See AH
AutoKey IKE VPN 9
  management 9

C

CA certificates 18, 22
certificates 10
  CA 18, 22
  loading 26
  local 22
  requesting 23
  revocation 21, 36
  via e-mail 22
Challenge Handshake Authentication Protocol
  See CHAP
CHAP 273, 276
character types, ScreenOS supported x
CLI
  conventions vi
container 242
conventions
  CLI vi
  illustration ix
  names x
  WebUI vii
CRL (Certificate Revocation List) 20, 36
  loading 20
cryptographic options 40–57
  anti-replay checking 46, 54
  authentication algorithms 44, 49, 53, 57
  authentication types 42, 51
  certificate bit lengths 43, 51
  dialup 50–57
  dialup VPN recommendations 57
  Diffie-Hellman groups 43, 46, 52, 55
  encryption algorithms 44, 48, 52, 57
  ESP 48, 56
  IKE ID 44–46, 53–54
  IPSec protocols 47, 56
  key methods 42
  PFS 46, 55
  Phase 1 modes 42, 51
  site-to-site 41–49
  site-to-site VPN recommendations 49
  Transport mode 56
  Tunnel mode 56

D

Data Encryption Standard
  See DES
DES 8
Diffie-Hellman exchange 13
Diffie-Hellman groups 13, 43, 46, 52, 55
digital signature 16
DIP pools
  extended interfaces 168
  NAT for VPNs 168
DN (distinguished name) 237
DNS
  L2TP settings 276

E

Encapsulating Security Payload
  See ESP
encryption
  algorithms 8, 44, 48, 52, 57
ESP 3, 7, 8
  authenticate only 48
  encrypt and authenticate 48, 56
  encrypt only 48

G

group IKE ID
  certificates 238–249
  preshared key 250–258
group IKE ID user 237–258
  certificates 238
  preshared key 250

H

hash-based message authentication code
  See HMAC
HMAC 7

I

IKE 9, 77, 91, 201
  group IKE ID user 237–258
  group IKE ID, container 242
  group IKE ID, wildcard 241
  heartbeats 384
  hello messages 384
  IKE ID 44–46, 53–54
  IKE ID recommendations 68
  IKE ID, Windows 200 288
  local ID, ASN1-DN 240
  Phase 1 proposals, predefined 11
  Phase 2 proposals, predefined 14
  proxy IDs 14
  redundant gateways 382–400
  remote ID, ASN1-DN 240
  shared IKE ID user 259–267
illustration
  conventions ix

interfaces
extended 168
Internet Key Exchange
See IKE
IP addresses
extended 168
IP Security
See IPSec
IPSec 3
AH 2, 47, 56
digital signature 16
ESP 2, 47, 56
SAs 2, 10, 11, 13
SPI 2
transport mode 4, 273, 279, 286
tunnel 2
tunnel mode 5
tunnel negotiation 11

K

keepalive
frequency, NAT-T 303
L2TP 283

L

L2TP 269–298
access concentrator, *See* LAC
compulsory configuration 270
decapsulation 275
default parameters 276
encapsulation 274
hello signal 284
Keep Alive 283, 284
L2TP-only on Windows 2000 273
network server, *See* LNS
operational mode 273
RADIUS server 276
ScreenOS support 273
SecurID server 276
tunnel 279
voluntary configuration 270
Windows 2000 291
Windows 2000 tunnel authentication 283
L2TP-over-IPSec 4, 279, 286
tunnel 279
LAC 270
NetScreen-Remote 5.0 270
Windows 2000 270
Layer 2 Tunneling Protocol
See L2TP
LNS 270
local certificate 22

M

Main Mode 12
Manual Key 131, 142
management 9
MD5 7
Message Digest version 5
See MD5
MIB files, importing 325
MIP
VPNs 168
modulus 13

N

names
conventions x
NAT
IPSec and NAT 301
NAT servers 301
NAT-dst
VPNs 168
NAT-src
VPNs 171
NAT-T 301
enabling 305
keepalive frequency 303
NAT-Traversal
See NAT-T
NetScreen-Remote
AutoKey IKE VPN 201
dynamic peer 209, 220
NAT-T option 301
NHTB table 326–331
addressing scheme 328
automatic entries 331
manual entries 330
mapping routes to tunnels 327

O

OCSP (Online Certificate Status Protocol) 36
client 36
responder 36

P

packet flow
inbound VPN 63–64
outbound VPN 61–62
policy-based VPN 65–66
route-based VPN 60–64
PAP 273, 276
Password Authentication Protocol
See PAP
Perfect Forward Secrecy
See PFS
PFS 14, 46, 55
Phase 1 11
proposals 11
proposals, predefined 11
Phase 2 13
proposals 13
proposals, predefined 14
PKI 18
Point-to-Point Protocol
See PPP
policies
bidirectional VPNs 143
policy-based VPNs 58
PPP 271
preshared key 9, 201
proposals
Phase 1 11, 67
Phase 2 13, 67
protocols
CHAP 273
PAP 273
PPP 271
proxy IDs 14
matching 67
VPNs and NAT 168–169

Public key infrastructure
See PKI
Public/private key pair 19

R

RADIUS
L2TP 276
redundant gateways 382–400
recovery procedure 385
TCP SYN flag checking 388
rekey option, VPN monitoring 308
replay protection 14
route-based VPNs 58

S

SAs 10, 11, 13
check in packet flow 62
SCEP (Simple Certificate Enrollment Protocol) 30
Secure Hash Algorithm-1
See SHA-1
SecurID
L2TP 276
security association
See SAs
SHA-1 7
SNMP
MIB files, importing 325
VPN monitoring 325

T

TCP
SYN flag checking 388
transport mode 4, 273, 279, 286
Triple DES
See 3DES
tunnel mode 5

U

UDP
checksum 303
NAT-T encapsulation 301
users
group IKE ID 237–258
shared IKE ID 259–267

V

Valicert 36
Verisign 36
VPN monitoring 307–322
destination address 308–311
destination address, XAuth 309
ICMP echo requests 325
outgoing interface 308–311
policies 310
rekey option 308, 331
routing design 323
SNMP 325
status changes 307, 310
VPNs
Aggressive mode 12
AutoKey IKE 9
configuration tips 67–68
cryptographic options 40–57
Diffie-Hellman exchange 13
Diffie-Hellman groups 13
FQDN aliases 152
FQDN for gateway 151–167
Main mode 12
MIP 168
multiple tunnels per tunnel interface 326–381
NAT for overlapping addresses 168–185
NAT-dst 168
NAT-src 171
packet flow 60–66
Phase 1 11
Phase 2 13
proxy IDs, matching 67
redundant gateways 382–400
redundant groups, recovery procedure 385
replay protection 14
route- vs policy-based 58
SAs 10
tunnel always up 308
VPN groups 382
VPN monitoring and rekey 308

W

WebUI
conventions vii
wildcard 241
WINS
L2TP settings 276

X

XAuth
VPN monitoring 309

NetScreen Concepts & Examples

ScreenOS Reference Guide

Volume 6: Dynamic Routing

ScreenOS 5.0.0

P/N 093-0929-000

Rev. B

Copyright Notice

NetScreen Technologies, Inc.
Building #3
805 11th Avenue
Sunnyvale, CA 94089
www.netscreen.com

FCC Statement

The following information is for FCC compliance of Class A devices: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. The equipment generates, uses, and can radiate radio-frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case users will be required to correct the interference at their own expense.

The following information is for FCC compliance of Class B devices: The equipment described in this manual generates and may radiate radio-frequency energy. If it is not installed in accordance with NetScreen's installation instructions, it may cause interference with radio and television reception. This equipment has been tested and found to comply with the limits for a Class B digital device in accordance with the specifications in part 15 of the FCC rules. These specifications are designed to provide reasonable protection against such interference in a residential installation. However, there is no guarantee that interference will not occur in a particular installation.

If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Consult the dealer or an experienced radio/TV technician for help.
- Connect the equipment to an outlet on a circuit different from that to which the receiver is connected.

Caution: Changes or modifications to this product could void the user's warranty and authority to operate this device.

Disclaimer

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR NETSCREEN REPRESENTATIVE FOR A COPY.

Contents

Preface v
Conventions vi
CLI Conventions vi
WebUI Conventions vii
Illustration Conventions ix
Naming Conventions and Character Types x
NetScreen Documentation xi
Chapter 1 Virtual Routers 1
Virtual Routers on NetScreen Devices 3
Using Two VRs 3
Forwarding Traffic between VRs 4
Configuring Two Virtual Routers 4
Example: Binding a Zone to the untrust-vr 5
Custom Virtual Routers 7
Example: Creating a Custom Virtual Router 7
Example: Removing a Custom Virtual Router 8
Virtual Routers and Virtual Systems 9
Example: Creating a Custom Virtual Router in a vsys 10
Example: Defining a Route with a Shared Virtual Router as the Next-Hop 11
Modifying Virtual Routers 12
Virtual Router ID 12
Example: Assigning a Virtual Router ID 13
Maximum Number of Routing Table Entries 14
Example: Limiting the Maximum Number of Routing Table Entries 14
Route Selection 15
Route Preference 15
Example: Setting a Route Preference 16
Route Metric 17
Source-Based Routing 17
Example: Source-Based Routing 19
Route Redistribution 21
Configuring a Route Map 22
Route Filtering 24
Access Lists 24
Example: Configuring an Access List 25
Example: Redistributing BGP Routes into OSPF 26
Exporting and Importing Routes between VRs 28
Example: Configuring a Route Export Rule 29
Chapter 2 Open Shortest Path First (OSPF) 33
Overview of OSPF 34
Areas 34
Router Classification 35
Hello Protocol 35
Network Types 36
Broadcast Networks 36
Point-to-Point Networks 36
Link State Advertisements 37
Basic OSPF Configuration 38
Creating an OSPF Routing Instance in a Virtual Router 39

Example: Creating an OSPF Routing Instance 39
Example: Removing an OSPF Routing Instance 40
Defining an OSPF Area 41
Example: Creating an OSPF Area 41
Assigning Interfaces to an OSPF Area 42
Example: Assigning Interfaces to OSPF Areas 42
Example: Configuring an Area Range 43
Enabling OSPF on Interfaces 44
Example: Enabling OSPF on Interfaces 44
Example: Disabling OSPF on an Interface 45
Verifying the Configuration 46
Redistributing Routes 49
Example: Redistributing a BGP Route into OSPF 49
Summarizing Redistributed Routes 50
Example: Summarizing Redistributed Routes 50
Global OSPF Parameters 51
Example: Advertising the Default Route 52
Virtual Links 53
Example: Creating a Virtual Link 54
Example: Creating an Automatic Virtual Link 56
OSPF Interface Parameters 57
Example: Setting OSPF Interface Parameters 59
Security Configuration 60
Authenticating Neighbors 60
Example: Configuring the Clear-Text Password Authentication Method 60
Example: Configuring the MD5 Password Authentication Method 61
Filtering OSPF Neighbors 62
Example: Configuring a Neighbor List 62
Rejecting Default Routes 63
Example: Removing the Default Route from the Route Table 63
Protecting against Flooding 64
Example: Configuring the Hello Threshold 64
Example: Configuring the LSA Threshold 65
Chapter 3 Routing Information Protocol (RIP) 67
Overview of RIP 68
Basic RIP Configuration 69
Creating a RIP Routing Instance in a Virtual Router 70
Example: Creating a RIP Routing Instance 70
Example: Removing a RIP Routing Instance 71
Enabling RIP on Interfaces 72
Example: Enabling RIP on Interfaces 72
Example: Disabling RIP on an Interface 73
Redistributing Routes 73
Example: Redistributing Routes into RIP 74
Global RIP Parameters 76
Example: Advertising the Default Route to RIP Neighbors 77
RIP Interface Parameters 78
Example: Setting RIP Interface Parameters 79
Security Configuration 80
Authenticating Neighbors 80
Example: Configuring the MD5 Password Authentication Method 81
Filtering RIP Neighbors 82
Example: Configuring Trusted Neighbors 82

Rejecting Default Routes 83
Example: Rejecting Default Routes 83
Protecting Against Flooding 84
Example: Configuring an Update Threshold 84
Example: RIP on Tunnel Interfaces 85
Chapter 4 Border Gateway Protocol (BGP) 87
Overview of BGP 88
Types of BGP Messages 89
Path Attributes 89
External and Internal BGP 90
Basic BGP Configuration 91
Creating and Enabling a BGP Routing Instance in a Virtual Router 92
Example: Creating a BGP Routing Instance 92
Example: Removing a BGP Routing Instance 93
Enabling BGP on Interfaces 94
Example: Enabling BGP on Interfaces 94
Example: Disabling BGP on Interfaces 94
Configuring a BGP Peer 95
Example: Configuring a BGP Peer 97
Example: Configuring an IBGP Peer-Group 98
Verifying the BGP Configuration 100
Security Configuration 102
Authenticating Neighbors 102
Example: Configuring MD5 Authentication for BGP Peers 102
Rejecting Default Routes 103
Example: Rejecting Default Routes 103
Optional BGP Configurations 104
Redistributing Routes 105
Example: Redistributing an OSPF Route into BGP 105
AS-Path Access List 106
Example: Configuring an AS-Path Access List 106
Route Reflection 107
Example: Configuring the Virtual Router as a Route Reflector 108
Confederations 110
Example: Configuring a Confederation 111
BGP Communities 113
Index ..... IX-I

Preface

Routing is an essential part of security devices such as NetScreen appliances and systems. Dynamic routing allows NetScreen devices to exchange routing information with routers and other network devices using commonly-implemented protocols and automatically build and update routing tables. Dynamic routing protocols greatly reduce the time lag between network topology changes and routing table adjustments because the adjustments occur automatically.

Volume 6, “Dynamic Routing” describes how to configure virtual routers on NetScreen devices and how to redistribute routing table entries between protocols or between virtual routers, and how to configure the following dynamic routing protocols on NetScreen devices: Open Shortest Path First (OSPF), Routing Information Protocol (RIP), and Border Gateway Protocol (BGP).

CONVENTIONS

This document contains several types of conventions, which are introduced in the following sections:

- "CLI Conventions"
- "WebUI Conventions" on page vii
- "Illustration Conventions" on page ix
- "Naming Conventions and Character Types" on page x

CLI Conventions

The following conventions are used when presenting the syntax of a command line interface (CLI) command:

- Anything inside square brackets [ ] is optional.
- Anything inside braces { } is required.
- If there is more than one choice, each choice is separated by a pipe (|). For example,

  ```
  set interface { ethernet1 | ethernet2 | ethernet3 } manage
  ```

  means "set the management options for the ethernet1, ethernet2, or ethernet3 interface".
- Variables appear in *italic.* For example:

  ```
  set admin user name password
  ```

When a CLI command appears within the context of a sentence, it is in **bold** (except for variables, which are always in *italic*). For example: "Use the **get system** command to display the serial number of a NetScreen device."

Note: *When typing a keyword, you only have to type enough letters to identify the word uniquely. For example, typing* ***set adm u joe j12fmt54*** *is enough to enter the command* ***set admin user joe j12fmt54****. Although you can use this shortcut when entering commands, all the commands documented here are presented in their entirety.*

WebUI Conventions

Throughout this book, a chevron (>) is used to indicate navigation through the WebUI by clicking menu options and links. For example, the path to the address configuration dialog box is presented as **Objects > Addresses > List > New**. This navigational sequence is shown below.

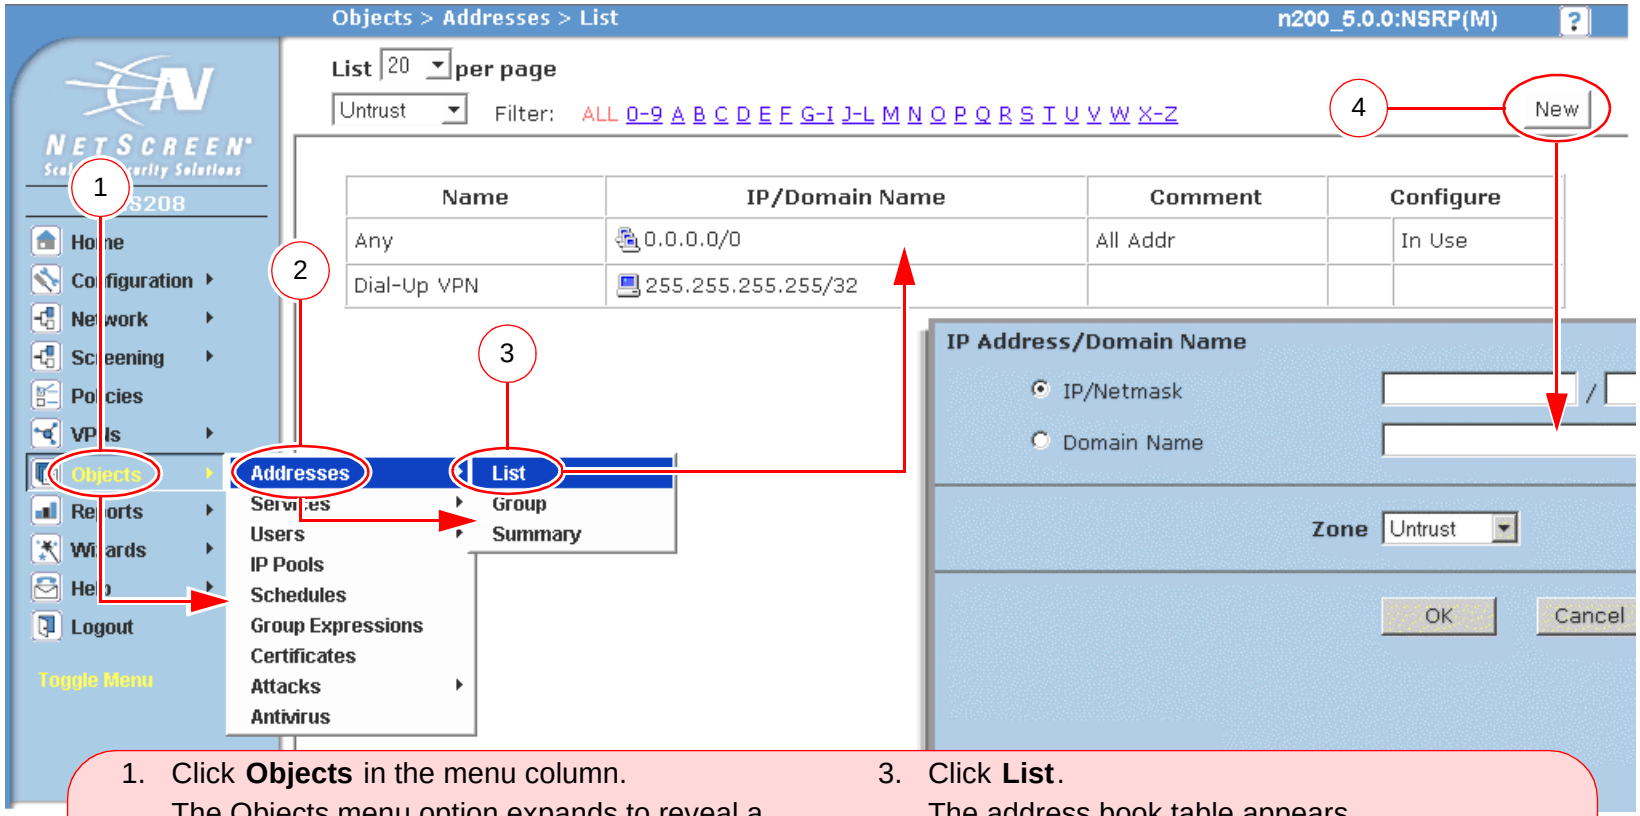

1. Click **Objects** in the menu column.
   The Objects menu option expands to reveal a subset of options for Objects.
2. (Applet menu) Hover the mouse over **Addresses**.
   (DHTML menu) Click **Addresses**.
   The Addresses option expands to reveal a subset of options for Addresses.
3. Click **List**.
   The address book table appears.
4. Click the **New** link.
   The new address configuration dialog box appears.

To perform a task with the WebUI, you must first navigate to the appropriate dialog box where you can then define objects and set parameters. The set of instructions for each task is divided into two parts: a navigational path and configuration details. For example, the following set of instructions includes the path to the address configuration dialog box and the settings for you to configure:

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: addr_1

IP Address/Domain Name:

IP/Netmask: (select), 10.2.2.5/32

Zone: Untrust

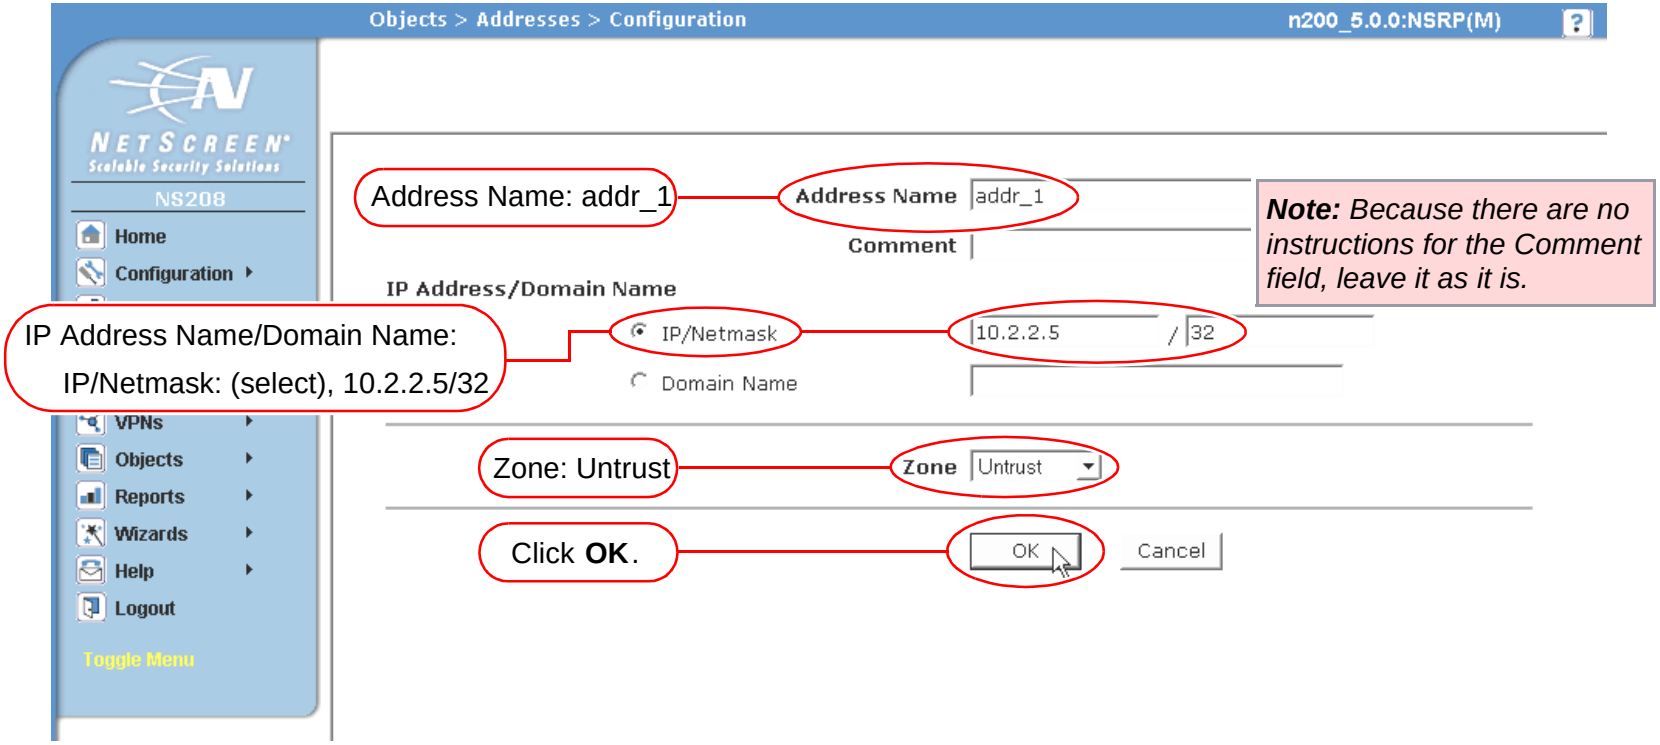

Illustration Conventions

The following graphics make up the basic set of images used in illustrations throughout this book:

Generic NetScreen Device

Virtual Routing Domain

Security Zone

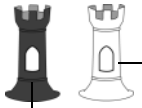

Security Zone Interfaces

White = Protected Zone Interface (example: Trust Zone)

Black = Outside Zone Interface (example: Untrust Zone)

Tunnel Interface

VPN Tunnel

Router Icon

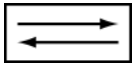

Switch Icon

Local Area Network (LAN) with a Single Subnet

(example: 10.1.1.0/24)

Internet

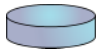

Dynamic IP (DIP) Pool

Desktop Computer

Laptop Computer

Generic Network Device

(examples: NAT server, Access Concentrator)

Server

Naming Conventions and Character Types

ScreenOS employs the following conventions regarding the names of objects—such as addresses, admin users, auth servers, IKE gateways, virtual systems, VPN tunnels, and zones—defined in ScreenOS configurations.

- If a name string includes one or more spaces, the entire string must be enclosed within double quotes ("); for example, **set address trust "local LAN" 10.1.1.0/24**.
- NetScreen trims any spaces leading or trailing text within a set of double quotes; for example, **" local LAN "** becomes **"local LAN"**.
- NetScreen treats multiple consecutive spaces as a single space.
- Name strings are case sensitive, although many CLI key words are case insensitive. For example, **"local LAN"** is different from **"local lan"**.

ScreenOS supports the following character types:

- Single-byte character sets (SBCS) and multiple-byte character sets (MBCS). Examples of SBCS are ASCII, European, and Hebrew. Examples of MBCS—also referred to as double-byte character sets (DBCS)—are Chinese, Korean, and Japanese.

  Note: *A console connection only supports SBCS. The WebUI supports both SBCS and MBCS, depending on the character sets that your Web browser supports.*

- ASCII characters from 32 (0x20 in hexidecimals) to 255 (0xff), except double quotes ("), which have special significance as an indicator of the beginning or end of a name string that includes spaces.

NETSCREEN DOCUMENTATION

To obtain technical documentation for any NetScreen product, visit www.netscreen.com/resources/manuals/.

To obtain the latest software version, visit www.netscreen.com/services/download_soft. Select a category of software product from the dropdown list, then follow the displayed instructions. (You must be a registered user to download NetScreen software.)

If you find any errors or omissions in the following content, please contact us at the e-mail address below:

techpubs@netscreen.com

1

Virtual Routers

Routing is an essential part of security devices such as NetScreen appliances and systems. Without routing, security devices cannot effectively forward secure traffic to targeted destinations. You can configure a NetScreen device to only use static routes, but you must manually add, remove, or modify routing table entries whenever changes occur on the network. (See the chapter on “Routing Tables and Static Routing” in Volume 2 for more information about configuring static routes.) Dynamic routing allows NetScreen devices to exchange routing information with routers and other network devices using commonly-implemented protocols and automatically build and update routing tables. Dynamic routing protocols greatly reduce the time lag between network topology changes and routing table adjustments because the adjustments occur automatically.

This chapter explains how to configure virtual routers on NetScreen devices and how to redistribute routing table entries between protocols or between virtual routers. The remaining chapters in this volume describe how to configure specific dynamic routing protocols on NetScreen devices.

This chapter contains the following sections:

- “Virtual Routers on NetScreen Devices” on page 3
  - “Using Two VRs” on page 3
  - “Forwarding Traffic between VRs” on page 4
  - “Configuring Two Virtual Routers” on page 4
  - “Custom Virtual Routers” on page 7
  - “Virtual Routers and Virtual Systems” on page 9
- “Modifying Virtual Routers” on page 12
  - “Virtual Router ID” on page 12
  - “Maximum Number of Routing Table Entries” on page 14

- "Route Selection" on page 15
  - "Route Preference" on page 15
  - "Route Metric" on page 17
  - "Source-Based Routing" on page 17
- "Route Redistribution" on page 21
  - "Configuring a Route Map" on page 22
  - "Route Filtering" on page 24
  - "Access Lists" on page 24
- "Exporting and Importing Routes between VRs" on page 28

VIRTUAL ROUTERS ON NETSCREEN DEVICES

ScreenOS can divide its routing component into two or more virtual routers. A virtual router (VR) supports static routing and dynamic routing protocols, which you can enable simultaneously in one virtual router. There are two predefined virtual routers on NetScreen devices:

- trust-vr, which by default contains all predefined security zones and any user-defined zones
- untrust-vr, which by default does not contain any security zones

You cannot delete the trust-vr or untrust-vr virtual routers. On some NetScreen devices, you can create and configure additional virtual routers (see "Custom Virtual Routers" on page 7 for more information on creating custom virtual routers). You can configure certain parameters for the predefined and custom virtual routers (see "Modifying Virtual Routers" on page 12).

Using Two VRs

By separating routing information into two virtual routers, you can control the information in a given routing domain that is visible to other routing domains. For example, you can keep the routing information for all the security zones inside a corporate network on the predefined virtual router trust-vr, and all the routing information for all the zones outside the corporate network on the other predefined virtual router, untrust-vr. Because the information in the routing table of one virtual router is not visible to the other virtual router, you can keep internal network routing information separate from untrusted sources outside the company.

Forwarding Traffic between VRs

When there are two virtual routers on a NetScreen device, traffic is not automatically forwarded between zones that reside in different VRs, even if there are policies that permit the traffic. To enable traffic to pass from one virtual router to another, you need to make sure that there are appropriate entries in the routing table. To do this, you can:

- Configure a static route in one virtual router that defines another VR as the next hop for the route. This route can even be the default route for the virtual router. For example, you can configure a default route for the trust-vr with the untrust-vr as the next hop. If the destination in an outbound packet does not match any other entries in the trust-vr routing table, it is forwarded to the untrust-vr. See the chapter on “Routing Tables and Static Routing” in Volume 2 for more information about configuring static routes.
- Export routes from the routing table in one virtual router into the routing table of another VR. You can export and import specific routes. You can also export all routes in the trust-vr routing table to the untrust-vr. This enables packets received in the untrust-vr to be forwarded to destinations in the trust-vr. See “Exporting and Importing Routes between VRs” on page 28 for more information.

Configuring Two Virtual Routers

As mentioned previously, you can configure multiple virtual routers in a NetScreen device with each virtual router maintaining a separate routing table. By default, all predefined and user-defined security zones are bound to the trust-vr virtual router. This also means that all interfaces that are bound to those security zones also belong to the trust-vr virtual router. This section discusses how to bind a security zone (and its interfaces) to the untrust-vr virtual router.

You can bind a security zone to only one virtual router. You can bind multiple security zones to a single virtual router when there is no address overlap between zones. That is, all interfaces in the zones must be in route mode. Once a zone is bound to a virtual router, all the interfaces in that zone belong to the virtual router. You can change the binding of a security zone from one virtual router to another, however, you must first remove all interfaces from the zone. (For more information on binding and unbinding an interface to a security zone, see “Interfaces” on page **2**-65.)

The following are the basic steps in binding a security zone to the untrust-vr virtual router:

1. Remove all interfaces from the zone that you want to bind to the untrust-vr. You cannot modify a zone-to-virtual router binding if there is an interface assigned to the zone. If you have assigned an IP address to an interface, you need to remove the address assignment before removing the interface from the zone.
2. Assign the zone to the untrust-vr virtual router.
3. Assign interface(s) back to the zone.

Example: Binding a Zone to the untrust-vr

In the following example, the untrust security zone is bound by default to the trust-vr virtual router and the interface ethernet3 is bound to the untrust security zone. (There are no other interfaces bound to the untrust security zone.) You must first set the IP address and netmask of the ethernet3 interface to 0.0.0.0, then change the bindings so that the untrust security zone is bound to the untrust-vr virtual router.

WebUI

1. Unbind Interface from untrust Zone

Network > Interfaces(ethernet3) > Edit: Enter the following, and then click **OK**:

Zone Name: Null

IP Address/Netmask: 0.0.0.0/0

2. Bind untrust Zone to untrust-vr

Network > Zones (untrust) > Edit: Select **untrust-vr** from the Virtual Router Name drop-down list, and then click **OK**.

3. Bind Interface to untrust zone

Network > Interfaces(ethernet3) > Edit: Select **Untrust** from the Zone Name drop-down list, and then click **OK**.

CLI

1. **Unbind Interface from Untrust Zone**

```
set interface ethernet3 0.0.0.0/0
unset interface ethernet3 zone
```

2. **Bind untrust Zone to untrust-vr**

```
set zone untrust vr untrust-vr
```

3. **Bind Interface to untrust zone**

```
set interface eth3 zone untrust
save
```

In the following displays, the **get zone** output on the left shows the default interface, zone, and virtual router bindings. In the default bindings, the untrust zone is bound to the trust-vr. The **get zone** output on the right shows the interface, zone, and virtual router bindings after you have reconfigured the bindings; the untrust zone is now bound to the untrust-vr.

Untrust zone bound to trust-vr (default bindings)

```
ns-> get zone
Total of 12 zones in vsys root. 7 policy configurable zone(s)
------------------------------------------------------------
  ID Name          Type    Attr    VR          Default-IF   VSYS
   0 Null          Null    Shared untrust-vr   null         Root
   1 Untrust       Sec(L3) Shared trust-vr     ethernet3    Root
   2 Trust         Sec(L3)        trust-vr     ethernet1    Root
   3 DMZ           Sec(L3)        trust-vr     ethernet2    Root
   4 Self          Func           trust-vr     self         Root
   5 MGT           Func           trust-vr     vlan1        Root
   6 HA            Func           trust-vr     null         Root
  10 Global        Sec(L3)        trust-vr     null         Root
  11 V1-Untrust    Sec(L2)        trust-vr     v1-untrust   Root
  12 V1-Trust      Sec(L2)        trust-vr     v1-trust     Root
  13 V1-DMZ        Sec(L2)        trust-vr     v1-dmz       Root
  16 Untrust-Tun Tun              trust-vr     null         Root
```

Untrust zone bound to untrust-vr

```
ns-> get zone
Total of 12 zones in vsys root. 7 policy configurable zone(s)
------------------------------------------------------------
  ID Name          Type    Attr    VR          Default-IF   VSYS
   0 Null          Null    Shared untrust-vr   null         Root
   1 Untrust       Sec(L3) Shared untrust-vr   ethernet3    Root
   2 Trust         Sec(L3)        trust-vr     ethernet1    Root
   3 DMZ           Sec(L3)        trust-vr     ethernet2    Root
   4 Self          Func           trust-vr     self         Root
   5 MGT           Func           trust-vr     vlan1        Root
   6 HA            Func           trust-vr     null         Root
  10 Global        Sec(L3)        trust-vr     null         Root
  11 V1-Untrust    Sec(L2)        trust-vr     v1-untrust   Root
  12 V1-Trust      Sec(L2)        trust-vr     v1-trust     Root
  13 V1-DMZ        Sec(L2)        trust-vr     v1-dmz       Root
  16 Untrust-Tun Tun              trust-vr     null         Root
------------------------------------------------------------
```

Custom Virtual Routers

Some NetScreen devices[1] allow you to create custom virtual routers in addition to the two predefined virtual routers. You can modify all aspects of a user-defined virtual router, including the virtual router ID, the maximum number of entries allowed in the routing table, and the preference value for routes from specific protocols.

Example: Creating a Custom Virtual Router

In this example, you create a custom virtual router called trust2-vr and you enable automatic route exporting from the trust2-vr VR to the untrust-vr.

WebUI

Network > Routing > Virtual Routers > New: Enter the following, and then click **OK**:

Virtual Router Name: trust2-vr

Auto Export Route to Untrust-VR: (select)

CLI

```
set vrouter name trust2-vr
set vrouter trust2-vr auto-route-export
save
```

1. Only certain NetScreen devices support custom virtual routers. To create custom virtual routers, you need a software license key.

Example: Removing a Custom Virtual Router

In this example, you delete an existing user-defined virtual router named “trust2-vr”.

WebUI

Network > Routing > Virtual Routers: Click **Remove** for the trust2-vr.

When the prompt appears asking you to confirm the removal, click **OK**.

CLI

```
unset vrouter trust2-vr
```

When the prompt appears asking you to confirm the removal (vrouter unset, are you sure? y/[n]), type **Y**.

```
save
```

Note: *You cannot delete the predefined untrust-vr and trust-vr virtual routers, but you can delete any user-defined virtual router. To modify the name of a user-defined virtual router or change the virtual router ID, you must first delete the virtual router, and then recreate it with the new name or virtual router ID.*

Virtual Routers and Virtual Systems

When a root-level admin creates a vsys on virtual system-enabled[2] systems, the vsys automatically has the following virtual routers available for its use:

- Any root-level virtual routers that have been defined as sharable. The untrust-vr is, by default, a shared virtual router that is accessible by any vsys. You can configure other root-level virtual routers to be sharable.
- A vsys-level virtual router. When you create a vsys, a vsys-level virtual router is automatically created that maintains the routing table for the Trust-*vsysname* zone. You can choose to name the virtual router *vsysname*-vr or a user-defined name. A vsys-level virtual router cannot be shared by other vsys.

You can define one or more custom virtual routers for a vsys. For more information about virtual systems, see “Virtual Systems” on page **7** -1. In the following illustration, each of the three vsys has two virtual routers associated with it: a vsys-level virtual router named *vsysname*-vr, and the untrust-vr.

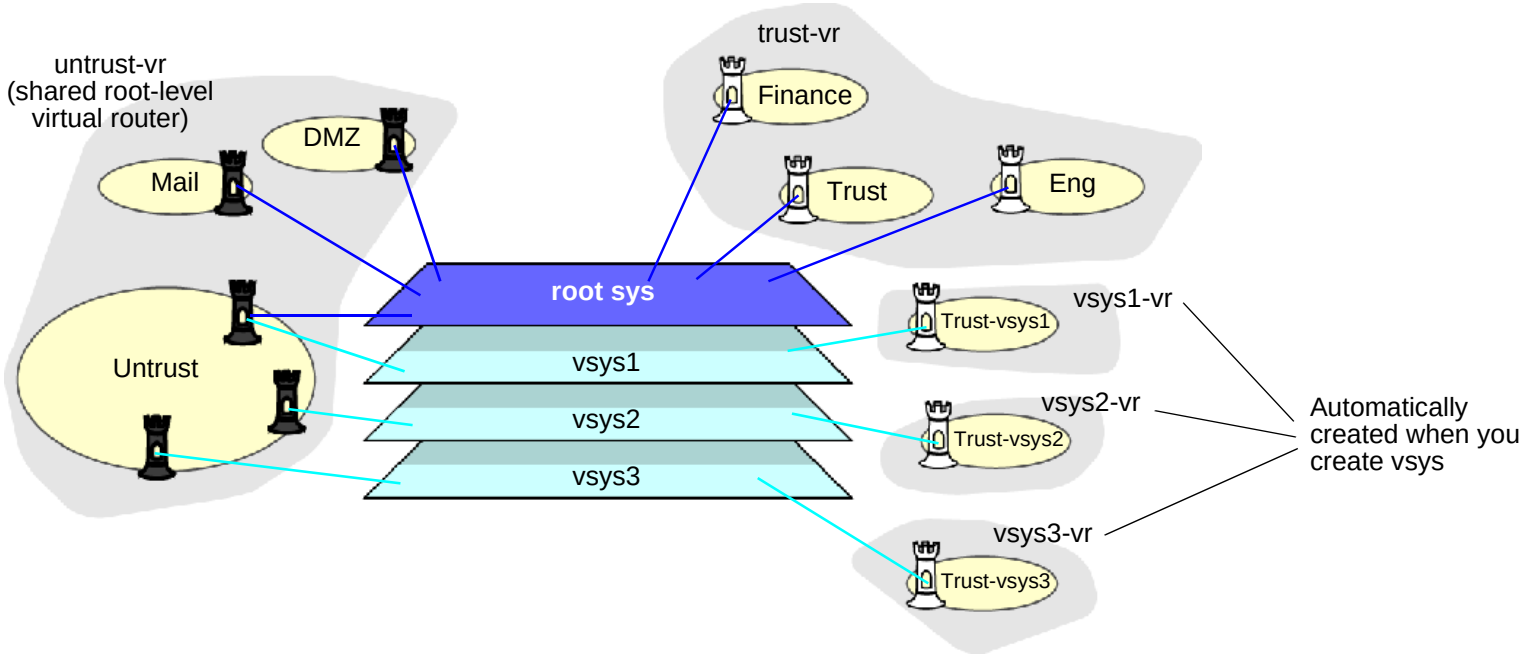

2. Only NetScreen systems (NetScreen-500, -5200, -5400) support virtual systems. To create vsys objects, you need a software license key.

Example: Creating a Custom Virtual Router in a vsys

In this example, you define a custom virtual router vr-1a with the router ID 10.1.1.9 for the vsys my-vsys1.

WebUI

Vsys > Enter (for my-vsys1) > Network > Routing > Virtual Routers > New: Enter the following, and then click **Apply**:

Virtual Router Name: vr-1a

Virtual Router ID: Custom (select)

In the text box, enter 10.1.1.9

CLI

```
set vsys my-vsys1
(my-vsys1) set vrouter name vr-1a
(my-vsys1/vr-1a) set router-id 10.1.1.9
(my-vsys1/vr-1a) exit
(my-vsys1) exit
```

Enter **y** at the following prompt:

```
Configuration modified, save? [y]/n
```

The vsys-level virtual router that is created when you create the vsys is the default virtual router for a vsys. The predefined Trust-*vsysname* security zone is bound by default to the default virtual router. You can bind a user-defined vsys-level security zone to any virtual router available to the vsys.

The untrust-vr is shared by default across all vsys. While vsys-level virtual routers are not sharable, you can define any root-level virtual router to be shared by the vsys. This allows you to define routes in a vsys-level virtual router that use a shared root-level virtual router as the next-hop. You can also configure route redistribution between a vsys-level virtual router and a shared root-level virtual router.

Example: Defining a Route with a Shared Virtual Router as the Next-Hop

In this example, the root-level virtual router my-router contains route table entries for the 4.0.0.0/8 network. If you configure the root-level virtual router my-router to be sharable by the vsys, then you can define a route in a vsys-level virtual router for the 4.0.0.0/8 destination with my-router as the next hop. In this example, the vsys is my-vsys1 and the vsys-level virtual router is my-vsys1-vr.

WebUI

Network > Routing > Virtual Routers > New: Enter the following, and then click **OK**:

Virtual Router Name: my-router

Shared and accessible by other vsys (select)

Vsys > Enter (for my-vsys1) > Network > Routing > Routing Entries > New (for my-vsys1-vr): Enter the following, and then click **OK**:

Network Address/Netmask: 40.0.0.0 255.0.0.0

Next Hop Virtual Router Name: (select) my-router

CLI

```
set vrouter name my-router sharable
set vsys my-vsys1
(my-vsys1) set vrouter my-vsys1-vr route 40.0.0.0/8 vrouter my-router
(my-vsys1) exit
```

Enter **y** at the following prompt:

```
Configuration modified, save? [y]/n
```

MODIFYING VIRTUAL ROUTERS

You can modify the following parameters for virtual routers:

- Virtual router ID
- Maximum number of entries allowed in the routing table
- Preference value for routes, based on protocol
- (For the trust-vr only) Enable or disable automatic route exporting to the untrust-vr for interfaces configured in Route mode.

Virtual Router ID

With dynamic routing protocols, each routing device uses a *unique* router identifier to communicate with other routing devices. The identifier can be in the form of a dotted decimal notation, like an IP address., or an integer value. If you do not define a specific virtual router ID before enabling a dynamic routing protocol, ScreenOS automatically selects the highest IP address of the active interfaces in the VR for the router identifier.

Note: *By default all NetScreen devices have IP address 192.168.1.1 assigned to the VLAN1 interface. If you do not specify a router ID before enabling a dynamic routing protocol on a NetScreen device, the IP address chosen for the router ID will likely be the default 192.168.1.1 address. This can cause a problem with routing as there cannot be multiple NetScreen virtual routers with the same router ID in a routing domain. Therefore, NetScreen recommends that you always explicitly assign a virtual router ID that is unique in the network. You can set the virtual router ID to the loopback interface address, as long as the loopback interface is not a Virtual Security Interface (VSI) in an NetScreen Redundancy Protocol (NSRP) cluster. (See Volume 8, "High Availability" for more information about configuring an NSRP cluster.)*

Example: Assigning a Virtual Router ID

In this example, you assign 0.0.0.10 as the router ID for the trust-vr.

Note: *In the WebUI, you must enter the router ID in dotted decimal notation. In the CLI, you can enter the router ID either in dotted decimal notation (0.0.0.10) or simply enter 10 (this is converted by the CLI to 0.0.0.10).*

WebUI

Network > Routing > Virtual Router (trust-vr) > Edit: Enter the following, and then click **OK**:

Virtual Router ID: Custom (select)

In the text box, enter 0.0.0.10

CLI

```
set vrouter trust-vr router-id 10
save
```

Note: *You cannot assign or change a router ID if you have already enabled a dynamic routing protocol in the VR. If you need to change the router ID, you must first disable the dynamic routing protocol(s) in the VR. For information on disabling a dynamic routing protocol in the VR, see the appropriate chapter in this volume.*

Maximum Number of Routing Table Entries

Each virtual router is allocated the routing table entries it needs from a system-wide pool. The maximum number of entries available depends upon the NetScreen device[3] and the number of virtual routers configured on the device. You can limit the maximum number of routing table entries that can be allocated for a specific virtual router. This helps prevent one virtual router from using up all the entries in the system.

Example: Limiting the Maximum Number of Routing Table Entries

In this example, you set the maximum number of routing table entries for the trust-vr to 20.

WebUI

Network > Routing > Virtual Routers > Edit (for trust-vr): Enter the following, and then click **OK**:

Maximum Route Entry:

Set limit at: (select), 20

CLI

```
set vrouter trust-vr max-routes 20
save
```

3. See the relevant product data sheet to determine the maximum number of routing table entries available on your NetScreen device.

ROUTE SELECTION

Multiple routes with the same prefix (IP address and mask) can exist in the routing table. Where the routing table contains multiple routes to the same destination, the preference values of each router are compared. The route that has the lowest preference value is selected. If the preference values are the same, the metric values are then compared. The route that has the lowest metric value is then selected.[4]

Route Preference

A route preference is a weight added to the route that influences the determination of the best path for traffic to reach its destination. When importing or adding a route to the routing table, the virtual router adds a preference value — determined by the protocol by which the route is learned — to the route. A low preference value (a number closer to 0) is preferable to a high preference value (a number further from 0).

In a virtual router, you can set the preference value for routes according to protocol. The following table shows the default preference values for routes of each protocol.

Protocol	Default Preference
Connected	0
Static	20
Auto-Exported	30
EBGP	40
OSPF	60
RIP	100
Imported	140
OSPF External Type 2	200
IBGP	250

4. If there are multiple routes to the same destination with the *same* preference values and the *same* metric values, then any one of those routes can be selected. In this case, selection of one specific route over another is not guaranteed or predictable.

You can also adjust the route preference value to direct traffic along preferred paths.

Note: *If the route preference changes for any type of route (for example, OSPF type 1 routes), the new preference displays in the route table, but the new preference does not take effect until the route is relearned (which can be achieved by disabling, then enabling the dynamic routing protocol), or, in the case of static routes, deleted and added again.*

Example: Setting a Route Preference

In this example, you specify a value of 4 as the preference for any "connected[5]" routes added to the route table for the untrust-vr.

WebUI

Network > Routing > Virtual Routers > Edit (for untrust-vr): Enter the following, and then click **OK**:

Route Preference:

Connected: 4

CLI

```
set vrouter untrust-vr preference connected 4
save
```

5. A route is connected when the router has an interface with an IP address in the destination network.

Route Metric

Route metrics determine the best path a packet can take to reach a given destination. Routers use route metrics to weigh two routes to the same destination and determine the use of one route over the other. When there are multiple routes to the same destination network with the same preference value, the route with the lowest metric prevails.

A route metric can be based on the number of routers a packet must traverse to reach a destination, the relative speed and bandwidth of the path, the dollar cost of the links making up the path, or a combination of these (and other) elements. When routes are learned dynamically, the neighboring router from which the route originates provides the metric. The default metric for connected routes is always 0. The default metric for static routes is 1, although you can specify a different metric value when configuring a static route.

Source-Based Routing

You can direct a ScreenOS virtual router to forward traffic based on the source IP address of a data packet instead of just the destination IP address. For example, this feature allows traffic from users on a specific subnet to be forwarded on one path while traffic from users on a different subnet are forwarded on another path.

By default, ScreenOS uses only destination IP addresses to find the best route in the virtual router's routing table. When source-based routing is enabled in a virtual router, ScreenOS first performs routing table lookup based on the source IP address. If ScreenOS does not find a route based on the source IP address, then the destination IP address is used for route lookup.

You define source-based routes as statically configured routes on specified virtual routers. Source-based routes only apply to the virtual router in which you configure them. For example, you cannot specify another virtual router as the next-hop for a source-based route. You also cannot redistribute source-based routes into another virtual router or into a routing protocol.

To use this feature:

1. Create one or more source-based routes for a specific virtual router, specifying the following information:
   - The name of the virtual router in which source-based routing applies
   - The source IP address on which ScreenOS performs a routing table lookup (This address appears as an entry in the routing table.)
   - The name of the outgoing interface on which the packet is forwarded
   - The next-hop for the source-based route (Note that if you have already specified a default gateway for the interface with the CLI **set interface** *interface* **gateway** *ip_addr* command, you do not need to specify the gateway parameter; the interface's default gateway is used as the next hop for the source-based route. You cannot specify another virtual router as the next-hop for the source-based route.)
   - The metric for the source-based route (If there are multiple source-based routes with the same prefix, only the route with the lowest metric is used for route lookup and other routes with the same prefix are marked as "inactive.")
2. Enable source-based routing for the virtual router. For any routing table lookup in the specified virtual router, ScreenOS uses the source IP address of the packet first. If no route is found for the source IP address, the destination IP address is used for the routing table lookup.

Example: Source-Based Routing

In the following example, traffic from users on the 10.1.1.0/24 subnetwork is forwarded to ISP 1, while traffic from users on the 10.1.2.0/24 subnetwork is forwarded to ISP 2. You need to configure two entries in the default trust-vr virtual router routing table and enable source-based routing:

- The subnetwork 10.1.1.0/24, with ethernet1 as the forwarding interface, and ISP 1's router (1.1.1.1) as the next-hop
- The subnetwork 10.1.2.0/24, with ethernet2 as the forwarding interface, and ISP 2's router (2.2.2.2) as the next-hop

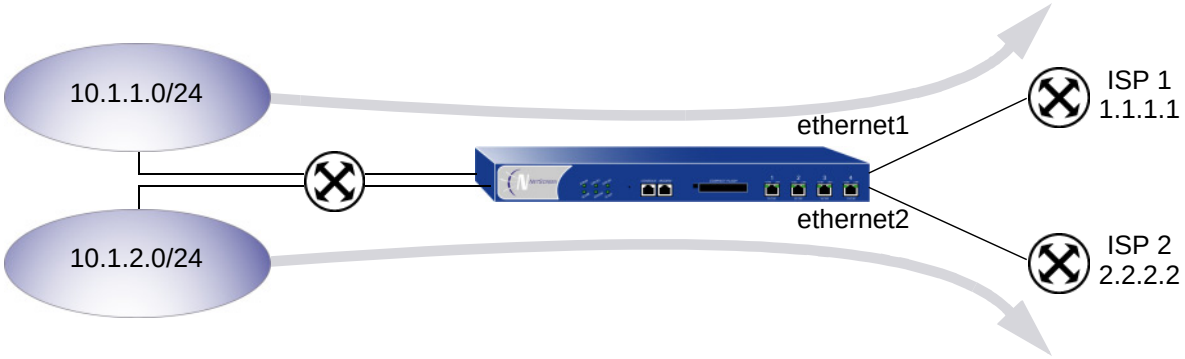

WebUI

Network > Routing > Source Routing > New (for trust-vr): Enter the following, and then click **OK**:

Network Address/Netmask: 10.1.1.0 255.255.255.0

Interface: ethernet1 (select)

Gateway IP Address: 1.1.1.1

Network > Routing > Source Routing > New (for trust-vr): Enter the following, and then click **OK**:

Network Address/Netmask: 10.1.2.0 255.255.255.0

Interface: ethernet2 (select)

Gateway IP Address: 2.2.2.2

Note: *In the WebUI, the default metric value is 1.*

Network > Routing > Source Routing: Select **Source Routing** (for trust-vr).

CLI

```
set vrouter trust-vr route source 10.1.1.0/24 interface ethernet1 gateway
    1.1.1.1 metric 1
set vrouter trust-vr route source 10.1.2.0/24 interface ethernet2 gateway
    2.2.2.2 metric 1
set vrouter trust-vr enable-source-routing
save
```

ROUTE REDISTRIBUTION

The routing table in a virtual router contains routes gathered by all dynamic routing protocols running in the VR, as well as static routes and directly-connected routes. By default, a dynamic routing protocol (such as OSPF, RIP, or BGP) advertises to its neighbors or peers only the routes that meet the following conditions:

- The routes must be active in the routing table.
- The routes must be learned by the dynamic routing protocol[6].

To allow a dynamic routing protocol to advertise routes that were learned by another protocol, including statically-configured routes, you need to *redistribute* routes from the source protocol into the advertising protocol.

You can redistribute routes learned from a routing protocol (including statically configured routes) into a different routing protocol in the same VR. This allows the receiving routing protocol to advertise the redistributed routes. When importing a route, the current domain has to translate all the information, particularly known routes, from the other protocol to its own protocol. For example, if a routing domain uses the OSPF protocol and it connects to a routing domain using the BGP protocol, the OSPF domain has to import all the routes from the BGP domain to inform all of its OSPF neighbors about how to reach devices in the BGP domain.

Routes are redistributed between protocols according to a *redistribution rule*[7] that is defined by the system or network administrator. When a route is added to a virtual router's routing table, all redistribution rules defined in the VR are applied one-by-one to the route to determine whether the route is to be redistributed. When a route is deleted from a virtual router's routing table, all redistribution rules defined in the VR are applied one-by-one to the route to determine whether the route is to be deleted from another routing protocol within the VR. Note that all redistribution rules are applied to the added or deleted route. There is no concept of rule order or "first matching rule" for redistribution rules.

On the NetScreen device, you configure a *route map* to specify which routes are to be redistributed and the attributes of the redistributed routes.

6. OSPF, RIP, and BGP also advertise connected routes for the ScreenOS interfaces on which these protocols are enabled.
7. You can only define one redistribution rule between any two protocols.

Configuring a Route Map

A *route map* consists of a set of statements that are applied in sequential order to a route. Each statement in the route map defines a condition that is compared to the route. A route is compared to each statement in a specified route map in order of increasing sequence number until there is a match, then the action specified by the statement is applied. If the route matches the condition in the route map statement, the route is either permitted or rejected. A route map statement can also modify certain attributes of a matching route. There is an implicit deny at the end of every route map; that is, if a route does not match any entry in the route map, the route is rejected.

The following are match conditions that you can configure in a route map statement:

Match Condition	Description
BGP AS Path	Matches a specified AS path access list. See “Route Filtering” on page 24.
BGP Community	Matches a specified community list. See “Route Filtering” on page 24.
OSPF route type	Matches either OSPF internal, external type 1, or external type 2.
Interface	Matches a specified interface.
IP address	Matches a specified access list. See “Route Filtering” on page 24.
Metric	Matches a specified route metric value.
Next-hop	Matches a specified access list. See “Route Filtering” on page 24.
Tag	Matches a specified route tag value or IP address.

For each match condition, you specify whether a route that matches the condition is accepted (permitted) or rejected (denied). If a route matches a condition and is permitted, you can optionally set attribute values for the route. The following are attributes that you can set in a route map statement:

Set Attributes	Description
BGP AS Path	Prepends a specified AS path access list to the path list attribute of the matching route.
BGP Community	Sets the community attribute of the matching route to the specified community list.
BGP local preference	Sets the local-pref attribute of the matching route to the specified value.
BGP Weight	Sets the weight of the matching route.
OSPF metric type	Sets the OSPF metric type of the matching route to either external type 1 or external type 2.
RIP offset metric	Sets the offset of the matching route between 1-16. This increases the metric on a less desirable path.
Metric	Sets the metric of the matching route to the specified value.
Next-hop of route	Sets the next-hop of the matching route to the specified IP address.
Tag	Sets the tag of the matching route to the specified tag value or IP address.

Route Filtering

Route filtering allows you to control which routes are allowed into a virtual router, which routes are advertised to peers, and which routes are redistributed from one routing protocol to another. You can apply filters to incoming routes sent by a routing peer or to outgoing routes sent by the NetScreen virtual router to peer routers. You can use the following filtering mechanisms:

- Access list—An access list is a set of specified IP address prefixes. You can use an access list to filter routes based on network prefixes. See Access Lists for information on configuring an access list.
- BGP AS-path access list—An AS-path attribute is a list of autonomous systems through which a route advertisement has passed and is part of the route information. An AS-path access list is a set of regular expressions that represent specific ASs. You can use an AS-path access list to filter routes based on the AS through which the route has traversed. See “AS-Path Access List” on page 106 for information on configuring an AS-path access list.
- BGP community list—A community attribute contains identifiers for the communities to which a BGP route belongs. A BGP community list is a set of BGP communities that you can use to filter routes based on the communities to which a route belongs. See “BGP Communities” on page 113 for information on configuring a BGP community list.

Access Lists

An access list is a sequential list of statements against which a route is compared. Each statement specifies the IP address/netmask of a network prefix and the forwarding status (permit or deny the route). For example, a statement in an access list can allow routes for the 1.1.1.0/24 subnet. Another statement in the same access list can deny routes for the 2.2.2.0/24 subnet. If a route matches a statement in the access list, the specified forwarding status is applied.

Note that the sequence of statements in an access list is important, as a route is compared to the first statement in the access list and then to subsequent statements until there is a match. If there is a match, all subsequent statements in the access list are ignored. Therefore, you should sequence the more specific statements before less specific statements. For example, place the statement that denies routes for the 1.1.1.1/30 subnet before the statement that permits routes for the 1.1.1.0/24 subnet.

Example: Configuring an Access List

In this example, you create an access list on the trust-vr. The access list has the following characteristics:

- Identifier: 2 (you must specify an access list identifier when configuring the access list)
- Forwarding Status: permit
- IP Address/Netmask Filtering: 1.1.1.1/24
- Sequence Number: 10 (positions this statement relative to other statements in the access list)

WebUI

Network > Routing > Virtual Routers > Access List: > New (for trust-vr): Enter the following and then click **OK**:

Access List ID: 2

Sequence No: 10

IP/Netmask: 1.1.1.1/24

Action: Permit

CLI

```
set vrouter trust-vr access-list 2 permit ip 1.1.1.1/24 10
save
```

Example: Redistributing BGP Routes into OSPF

In this example, you redistribute specified BGP routes that have passed through the autonomous system 65000 into OSPF. You first configure an AS-path access list that allows routes that have passed through AS 65000. (For more information about configuring an AS-path access list, see “AS-Path Access List” on page 106.) Next, you configure a route map “rtmap1” to match routes in the AS path access list. Finally, in OSPF you specify a redistribution rule that uses the route map ‘rtmap1” and specify BGP as the source protocol for the routes.

WebUI

1. BGP AS-Path Access List

Network > Routing > Virtual Routers > Edit (for trust-vr) > Edit BGP Instance > AS Path: Enter the following, and then click **Add**:

AS Path Access List ID: 1

Permit: (select)

AS Path String: _65000_

2. Route Map

Network > Routing > Virtual Routers > Route Map > New (for trust-vr): Enter the following, and then click **OK**:

Map Name: rtmap1

Sequence No.: 10

Action: permit (select)

Match Properties:

AS Path: (select), 1

3. Redistribution Rule

Network > Routing > Virtual Router > Edit (for trust-vr) > Edit OSPF Instance > Redistributable Rules: Select the following, and then click **Add**:

Route Map: rtmap1

Protocol: BGP

CLI

1. BGP AS-Path Access List

```
set vrouter trust-vr protocol bgp as-path-access-list 1 permit _65000_
```

2. Route Map

```
set vrouter trust-vr
ns(trust-vr)-> set route-map name rtmap1 permit 10
ns(trust-vr/rtmap1-10)-> set match as-path 1
ns(trust-vr/rtmap1-10)-> exit
ns(trust-vr)-> exit
```

3. Redistribution Rule

```
set vrouter trust-vr protocol ospf redistribute route-map rtmap1 protocol bgp
save
```

EXPORTING AND IMPORTING ROUTES BETWEEN VRS

If you have two virtual routers configured on a NetScreen device, you can allow specified routes in one VR to be learned by the other VR. To do this, you must define *export rules* on the source VR that will export routes to the destination VR. When exporting routes, a virtual router allows other VRs to learn about its network. On the destination VR, you can optionally configure *import rules* to control the routes that are allowed to be imported from the source VR. If there are no import rules on the destination VR, all exported routes are accepted.

To export and import routes between virtual routers:

1. On the source VR, define an export rule.
2. (Optional) On the destination VR, define an import rule. While this step is optional, an import rule allows you to further control the routes that the destination virtual router accepts from the source virtual router.

On the NetScreen device, you configure an export or import rule by specifying the following:

- The destination virtual router (for export rules) or source virtual router (for import rules)
- The protocol of the routes to be exported/imported
- Which routes are to be exported/imported
- (Optional) New or modified attributes of the exported/imported routes

Configuring an export or import rule is similar to configuring a redistribution rule. You configure a *route map* to specify which routes are to be exported/imported and the attributes of the routes.

You can configure the trust-vr to automatically export all its route table entries to the untrust-vr. You can also configure a user-defined virtual router to automatically export routes to other virtual routers. Routes in networks directly connected to interfaces in NAT mode cannot be exported.

Example: Configuring a Route Export Rule

In this example, OSPF routes for the 1.1.1.1/24 network in the trust-vr virtual router are exported to the untrust-vr routing domain. You first create an access list for the network prefix 1.1.1.1/24, which is then used in the route map "rtmap1"to filter for matches of routes for the 1.1.1.1/24 network. You then create a route export rule to export matching OSPF routes from the trust-vr to the untrust-vr virtual router.

WebUI

trust-vr

1. Access List

   Network > Routing > Virtual Routers > Access List: > New (for trust-vr): Enter the following and then click **OK**:

   Access List ID: 2

   Sequence No: 10

   IP/Netmask: 1.1.1.1/24

   Action: Permit

2. Route Map

   Network > Routing > Virtual Routers > Route Map > New (for trust-vr): Enter the following, and then click **OK**:

   Map Name: rtmap1

   Sequence No.: 10

   Action: permit (select)

   Match Properties:

   Access List: (select), 2

3. Export Rule

Network > Routing > Virtual Routers > Export Rules > New (for trust-vr): Enter the following, and then click **OK**:

Destination Virtual Router: untrust-vr

Route Map: rtmap1

Protocol: OSPF

CLI

trust-vr

1. Access List

```
set vrouter trust-vr
ns(trust-vr)-> set access-list 2 permit ip 1.1.1.1/24 10
```

2. Route Map

```
ns(trust-vr)-> set route-map name rtmap1 permit 10
ns(trust-vr/rtmap1-10)-> set match ip 2
ns(trust-vr/rtmap1-10)-> exit
```

3. Export Rule

```
ns(trust-vr)-> set export-to vrouter untrust-vr route-map rtmap1 protocol ospf
ns(trust-vr)-> exit
save
```

Example: Configuring the trust-vr to Automatically Export Routes to the untrust-vr

You can configure the trust-vr to automatically export all its routes to the untrust-vr. However, this does not necessarily mean that the untrust-vr imports all the routes exported by the trust-vr. If you define import rules for the untrust-vr, only routes that match the import rules are imported. In this example, the trust-vr automatically exports all routes to the untrust-vr, but an import rule on the untrust-vr allows only internal OSPF routes.

WebUI

trust-vr

Network > Routing > Virtual Router > Edit (for trust-vr): Select **Auto Export Route to Untrust-VR**, and then click **OK**.

untrust-vr

Network > Routing > Virtual Router > Route Map (for untrust-vr) > New: Enter the following, and then click **OK**:

Map Name: from-ospf-trust

Sequence No.: 10

Action: permit (select)

Route Type: internal-ospf (select)

CLI

trust-vr

```
set vrouter trust-vr auto-route-export
```

untrust-vr

```
set vrouter untrust-vr
ns(untrust-vr)-> set route-map name from-ospf-trust permit 10
ns(untrust-vr/from-ospf-trust-10)-> set match route-type internal-ospf
ns(untrust-vr/from-ospf-trust-10)-> exit
ns(untrust-vr)-> set import-from vrouter trust-vr route-map from-ospf-trust
    protocol ospf
ns(untrust-vr)-> exit
save
```

2

Open Shortest Path First (OSPF)

This chapter describes the Open Shortest Path First (OSPF) routing protocol on NetScreen devices. The following topics are covered:

- "Overview of OSPF" on page 34
  - "Areas" on page 34
  - "Router Classification" on page 35
  - "Hello Protocol" on page 35
  - "Network Types" on page 36
  - "Link State Advertisements" on page 37
- "Basic OSPF Configuration" on page 38
  - "Creating an OSPF Routing Instance in a Virtual Router" on page 39
  - "Defining an OSPF Area" on page 41
  - "Assigning Interfaces to an OSPF Area" on page 42
  - "Enabling OSPF on Interfaces" on page 44
  - "Verifying the Configuration" on page 46
- "Redistributing Routes" on page 49
  - "Summarizing Redistributed Routes" on page 50
- "Global OSPF Parameters" on page 51
  - "Virtual Links" on page 53
- "OSPF Interface Parameters" on page 57
- "Security Configuration" on page 60
  - "Authenticating Neighbors" on page 60
  - "Filtering OSPF Neighbors" on page 62
  - "Rejecting Default Routes" on page 63
  - "Protecting against Flooding" on page 64

OVERVIEW OF OSPF

The Open Shortest Path First (OSPF) routing protocol is an Interior Gateway Protocol (IGP) intended to operating within a single Autonomous System (AS). A router running OSPF distributes its state information (such as usable interfaces and neighbor reachability) by periodically flooding *link-state advertisements* (LSAs) throughout the AS.

Each OSPF router uses LSAs from neighboring routers to maintain a *link-state database*. The link-state database is a listing of topology and state information for the surrounding networks. The constant distribution of LSAs throughout the routing domain enables all routers in an AS to maintain identical link-state databases.

OSPF uses the link-state database to determine the best path to any network within the AS. This is done by generating a *shortest-path tree*, which is a graphical representation of the shortest path to any network within the AS. While all routers have the same link state database, they all have unique shortest-path trees because routers always generate the tree with themselves at the top of the tree.

Areas

By default, all routers are grouped into a single “backbone” area called area 0 (usually denoted as area 0.0.0.0). However, large geographically dispersed networks are typically segmented into multiple areas. This is because as networks grow, link-state databases grow and dividing the link-state database into smaller groups allows for better scalability.

Areas reduce the amount of routing information passed throughout the network because a router only maintains a link-state database for the area in which it resides. No link-state information is maintained for networks or routers outside the area. A router connected to multiple areas maintains a link-state database for each area to which it is connected. It is important to note that all areas must be directly connected to area 0, with one exception (to be covered later).

AS external advertisements describe routes to destinations in other autonomous systems and are flooded throughout an AS. Certain OSPF areas can be configured as *stub areas*; AS external advertisements are not flooded into these areas. There are two common types of stub areas used in OSPF:

- **Stub area** - An area that receives route summaries from the backbone area but does not receive link-state advertisements from other areas for routes learned through non-OSPF sources (BGP, for example). A stub area can be considered a *totally stubby area* if no summary routes are allowed in the stub area.
- **Not So Stubby Area (NSSA)** - Like a normal stub area, NSSAs cannot receive routes from non-OSPF sources outside the current area. However, external routes learned within the area can be learned and passed to other areas.

Router Classification

Routers that participate in OSPF routing are classified according to their function or location in the network:

- **Internal Router** - A router with all interfaces belonging to the same area.
- **Backbone Router** - A router that has an interface in the backbone area.
- **Area Border Router** - A router that attaches to multiple areas is called an area border router (ABR). An ABR summarizes routes from non-backbone areas for distribution to the backbone area. On NetScreen devices running OSPF, the backbone area is created by default. If you create a second area in ScreenOS, the device functions as an ABR.
- **AS Boundary Router** - When an OSPF area borders another AS, the router between the two autonomous systems is called an autonomous system boundary router (ASBR). An ASBR is responsible for advertising external AS routing information throughout an AS.

Hello Protocol

Two routers with interfaces on the same subnet are considered *neighbors*. Routers use the hello protocol to establish and maintain these neighbor relationships. When two routers establish bidirectional communication, they are said to have established an *adjacency*. If two routers do not establish an adjacency, they cannot exchange routing information.

In cases were there are multiple routers on a network, it is necessary to establish one router as the *designated router* (DR) and another as the *backup designated router* (BDR). The DR is responsible for flooding the network with LSAs that contain a list of all OSPF-enabled routers attached to the network. The DR is the only router that can form adjacencies with other routers on the network. Therefore, the DR is the only router on a network that can provide routing information to other routers. The BDR is responsible for becoming the designated router if the DR should fail.

Network Types

ScreenOS supports the following network types:

- Broadcast Networks
- Point-to-Point Networks

Broadcast Networks

A *broadcast network* is a network that connects many routers together and can send, or broadcast, a single physical message to all the attached routers. Pairs of routers on a broadcast network are assumed to be able to communicate with each other. Ethernet is an example of a broadcast network.

On broadcast networks, the OSPF router dynamically detects its neighbor routers by sending Hello packets to the multicast address 224.0.0.5. For broadcast networks, the Hello protocol elects a Designated Router and Backup Designated Router for the network.

A *non-broadcast network* is a network that connects many routers together but cannot broadcast messages to attached routers. On non-broadcast networks, OSPF protocol packets that are normally multicast need to be sent to each neighboring router. ScreenOS does not support OSPF on non-broadcast networks.

Point-to-Point Networks

A *point-to-point* network typically joins two routers over a Wide Area Network (WAN). An example of a point-to-point network is two NetScreen devices connected via an IPSec VPN tunnel. On point-to-point networks, the OSPF router dynamically detects neighbor routers by sending Hello packets to the multicast address 224.0.0.5.

Link State Advertisements

Each OSPF router sends out LSAs that define the router's local state information. Additionally, there are other types of LSAs that a router can send out, depending upon the router's OSPF function. The following table summarizes the LSA types:

LSA Type	Sent By	Flooded Throughout	Information Sent in LSA
Router LSA	All OSPF routers	Area	Describes the state of all router interfaces throughout the area.
Network LSA	Designated Router on broadcast and NBMA networks	Area	Contains a list of all routers connected to the network.
Summary LSA	Area Border Routers	Area	Describes a route to a destination outside the area but still inside the AS. There are two types: Type 3 summary-LSAs describe routes to networks. Type 4 summary-LSAs describe routes to AS boundary routers.
AS-External	Autonomous System Boundary Router	Autonomous System	Routes to networks in another AS. Often, this is the default route (0.0.0.0/0).

BASIC OSPF CONFIGURATION

Like RIP and BGP, you create OSPF on a per-virtual router basis on a NetScreen device. If you have multiple virtual routers (VRs) in a system, you can enable multiple instances of OSPF, one instance for each VR.

Note: *Before you configure a dynamic routing protocol on the NetScreen device, you should assign a virtual router ID, as described in Chapter 1, "Virtual Routers".*

This section describes the following basic steps to configure OSPF in a VR on a NetScreen device:

1. Create and enable the OSPF routing instance in a VR. This step also automatically creates an OSPF backbone area, with an area ID of 0.0.0.0, which cannot be deleted.
2. (Optional) Unless all OSPF interfaces will be connected to the backbone area, you need to define a new OSPF area with its own area ID. For example, if the NetScreen device is to act as an ABR, you need to create a new OSPF area in addition to the backbone area. You can configure the new area as a normal, stub, or not-so-stubby area.
3. Assign one or more interfaces to each OSPF area. You must explicitly add interfaces to an OSPF area, including the backbone area.
4. Enable OSPF on each interface.
5. Verify that OSPF is properly configured and operating.

This section describes how to perform each of these tasks for the example shown below using either the CLI or the WebUI. In the example, you configure the NetScreen device as an ABR connecting to area 0 through the ethernet3 interface and connecting to area 10 through the ethernet1 interface.

You can optionally configure other OSPF parameters, such as the following:

- Global parameters, such as virtual links, that are set at the VR level for the OSPF protocol (see "Global OSPF Parameters" on page 51)
- Interface parameters, such as authentication, that are set on a per-interface basis for the OSPF protocol (see "OSPF Interface Parameters" on page 57)
- Security-related OSPF parameters that are set at either the VR level or on a per-interface basis (see "Security Configuration" on page 60)

Creating an OSPF Routing Instance in a Virtual Router

You create and enable an OSPF routing instance on a specific virtual router on a NetScreen device. Creating the OSPF routing instance also automatically creates an OSPF backbone area. When you create and enable an OSPF routing instance on a VR, OSPF can transmit and receive packets on all OSPF-enabled interfaces in the VR.

Example: Creating an OSPF Routing Instance

In the following example, you first assign 0.0.0.10 as the router ID for the trust-vr virtual router. You then create an OSPF routing instance on the trust-vr. (For more information about virtual routers and configuring a virtual router on NetScreen devices, see Chapter 1, "Virtual Routers".)

WebUI

1. Router ID

Network > Routing > Virtual Router (trust-vr) > Edit: Enter the following, and then click **OK**:

Virtual Router ID: Custom (select)

In the text box, enter 0.0.0.10

2. OSPF Routing Instance

Network > Routing > Virtual Router (trust-vr) > Edit > Create OSPF Instance: Select **OSPF Enabled**, and then click **OK**.

CLI

1. Router ID

```
set vrouter trust-vr router-id 10
```

2. OSPF Routing Instance

```
set vrouter trust-vr protocol ospf
set vrouter trust-vr protocol ospf enable
save
```

> ***Note:*** *In the CLI, you must first create the OSPF routing instance before you can enable it. Thus, you must issue two separate CLI commands to enable an OSPF routing instance.*

Example: Removing an OSPF Routing Instance

In this example, you disable the OSPF routing instance in the trust-vr. OSPF stops transmitting and processing OSPF packets on all OSPF-enabled interfaces in the trust-vr.

WebUI

Network > Routing > Virtual Routers (trust-vr) > Edit > Edit OSPF Instance: Deselect OSPF Enabled, and then click **OK**.

Network > Routing > Virtual Routers (trust-vr) > Edit > Delete OSPF Instance, and then click **OK** at the confirmation prompt.

CLI

```
unset vrouter trust-vr protocol ospf enable
unset vrouter trust-vr protocol ospf
save
```

> ***Note:*** *In the CLI, you must first disable the OSPF routing instance before you can delete it. Thus, you must issue two separate CLI commands to remove an OSPF routing instance.*

Defining an OSPF Area

Areas reduce the amount of routing information that needs to be passed through the network because an OSPF router maintains a link-state database only for the area it resides in. No link-state information is maintained for networks or routers outside the area.

All areas must be connected to area 0, which is defined by default on the NetScreen virtual router when you create the OSPF routing instance on the virtual router. If you need to create an additional OSPF area, you can optionally define the area as a stub area or not-so-stubby area. See "Areas" on page 34 for more explanations of these types of areas.

You can optionally configure the following parameters for areas:

Area Parameter	Description	Default Value
Metric for default route	(NSSA and stub areas only) Specifies the metric for the default route advertisement.	1
Metric type for the default route.	(NSSA area only) Specifies the external metric type (1 or 2) for the default route.	1
No summary	(NSSA and stub areas only) Specifies that summary LSAs are *not* advertised into the area.	Summary LSAs are advertised into the area
Range	(All areas) Specifies a range of IP addresses to be advertised in summary LSAs, and whether they are advertised or not.	

Example: Creating an OSPF Area

In the following example, you create an OSPF area with an area ID of 10.

WebUI

Network > Routing > Virtual Routers > Edit (trust-vr) > Edit OSPF Instance > Area: Enter the following, and then click **OK**:

Area ID: 10

Type: normal (select)

Action: Add

CLI

```
set vrouter trust-vr protocol ospf area 10
save
```

Assigning Interfaces to an OSPF Area

Once an area is created, you can assign one or more interfaces to the area, using either the WebUI or the CLI **set interface** command.

Example: Assigning Interfaces to OSPF Areas

In the following example, you assign the ethernet1 interface to OSPF area 10 and assign the ethernet3 interface to OSPF area 0.

WebUI

Network > Routing > Virtual Routers > Edit (trust-vr) > Edit OSPF Instance > Area > Configure (Area 10): Use the **Add** button to move the ethernet1 interface from the Available Interface(s) column to the Selected Interfaces column. Click **OK**.

Network > Routing > Virtual Routers > Edit (for trust-vr) > Edit OSPF Instance > Area > Configure (Area 0): Use the **Add** button to move the ethernet3 interface from the Available Interface(s) column to the Selected Interfaces column. Click **OK**.

CLI

```
set interface ethernet1 protocol ospf area 10
set interface ethernet3 protocol ospf area 0
save
```

Example: Configuring an Area Range

By default, an ABR does not aggregate routes sent from one area to another area. Configuring an area range allows a group of subnets in an area to be consolidated into a single network address to be advertised in a single summary link advertisement to other areas. When you configure an area range, you specify whether to advertise or withhold the defined area range in advertisements.

In the following example, you define the following area ranges for area 10:

- 10.1.1.0/24 to be advertised
- 10.1.2.0/24 not to be advertised

WebUI

Network > Routing > Virtual Routers > Edit (for trust-vr) > Edit OSPF Instance > Area > Configure (0.0.0.10): Enter the following in the Area Range section, and then click **Add**:

IP / Netmask: 10.1.1.0/24

Type: (select) Advertise

Enter the following in the Area Range section, and then click **Add**:

IP / Netmask: 10.1.2.0/24

Type: (select) No Advertise

CLI

```
set vrouter trust-vr protocol ospf area 10 range 10.1.1.0/24 advertise
set vrouter trust-vr protocol ospf area 10 range 10.1.2.0/24 no-advertise
save
```

Enabling OSPF on Interfaces

By default, OSPF is disabled on all interfaces in the VR. You must explicitly enable OSPF on an interface after you assign the interface to an area. When you disable OSPF on an interface, OSPF does not transmit or receive packets on the specified interface, but interface configuration parameters are preserved.

Note: If you disable the OSPF routing instance in the VR (see "Example: Removing an OSPF Routing Instance" on page 40), OSPF stops transmitting and processing packets on all OSPF-enabled interfaces in the VR.

Example: Enabling OSPF on Interfaces

In this example, you enable the OSPF routing instance on the ethernet1 interface (which was previously assigned to area 10) and on the ethernet3 interface (which was previously assigned to area 0).

WebUI

Network > Interfaces > Edit (for ethernet1) > OSPF: Select **Enable Protocol OSPF**, and then click **Apply**.

Network > Interfaces > Edit (for ethernet3) > OSPF: Select **Enable Protocol OSPF**, and then click **Apply**.

CLI

```
set interface ethernet1 protocol ospf enable
set interface ethernet3 protocol ospf enable
save
```

Example: Disabling OSPF on an Interface

In this example, you disable the OSPF routing instance only on the ethernet1 interface. Note that any other interfaces in the trust-vr virtual router on which you have enabled OSPF are still able to transmit and process OSPF packets.

WebUI

Network > Interfaces > Edit (for ethernet1) > OSPF: Clear **Enable Protocol OSPF**, and then click **Apply**.

CLI

```
unset interface ethernet1 protocol ospf enable
save
```

Note: If you disable the OSPF routing instance in the VR (see "Example: Removing an OSPF Routing Instance" on page 40), OSPF stops transmitting and processing packets on all OSPF-enabled interfaces in the VR.

Verifying the Configuration

You can review the configuration you entered through the WebUI or the CLI by executing the following CLI command:

```
ns-> get vrouter trust-vr protocol ospf config
VR: trust-vr RouterId: 10.1.1.250
----------------------------------
set protocol ospf
set enable
set area 0.0.0.10 range 10.1.1.0 255.255.255.0 advertise
set area 0.0.0.10 range 10.1.2.0 255.255.255.0 no-advertise
set area 0.0.0.10
set vlink area-id 0.0.0.10 router-id 10.1.1.250
exit
set interface ethernet1 protocol ospf area 0.0.0.10
set interface ethernet1 protocol ospf enable
set interface ethernet3 protocol ospf area 0.0.0.0
set interface ethernet3 protocol ospf enable
```

You can verify that OSPF is running on the VR by executing the following CLI command:

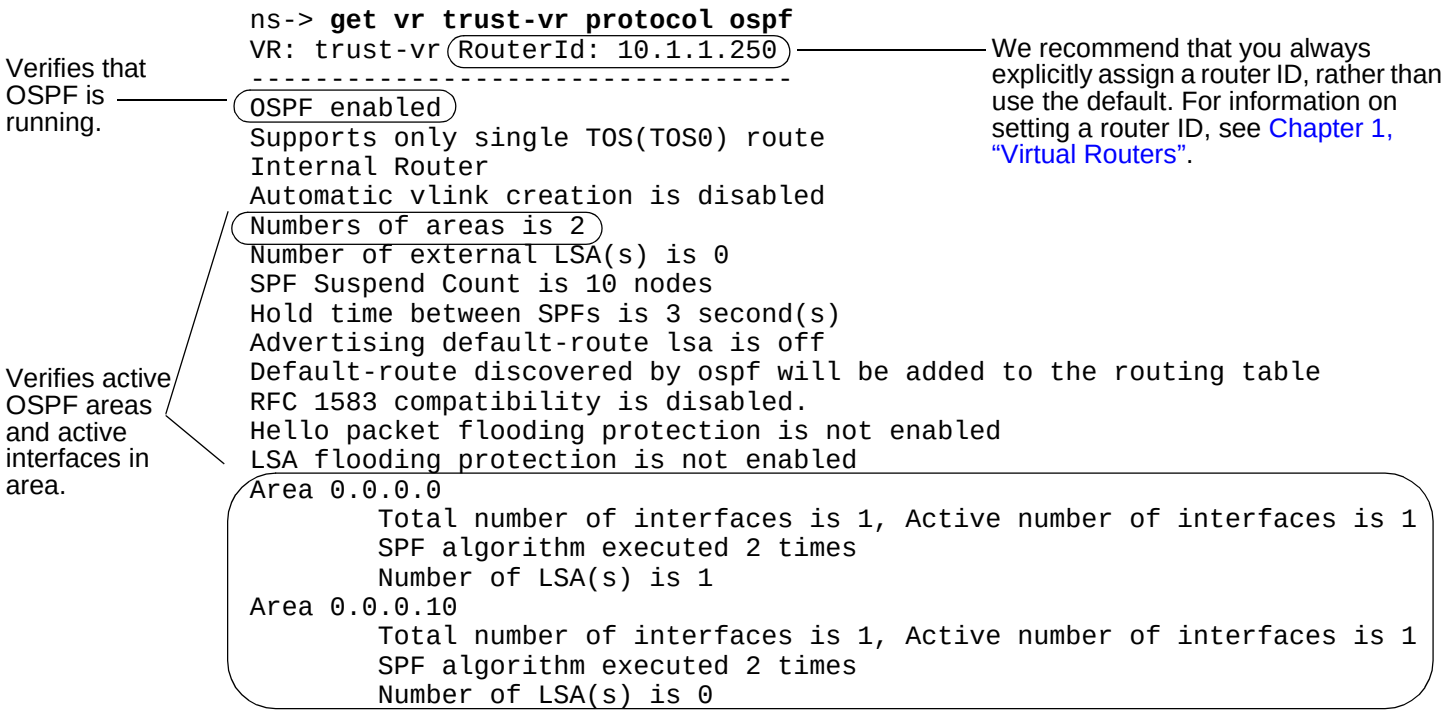

You can verify that OSPF is enabled on the interfaces and see the state of the interfaces by executing the following CLI command:

```
ns-> get vr trust-vr protocol ospf interface
VR: trust-vr RouterId: 10.1.1.250
---------------------------------
Interface   IpAddr        NetMask       AreaId        Status   State
-------------------------------------------------------------------------------

ethernet3   2.2.2.2       255.255.255.0 0.0.0.0       enabled Designated Router

ethernet1   10.1.1.1      255.255.255.0 0.0.0.10      enabled Up
```

You can configure the priority of the virtual router to be elected the DR or the BDR. See "OSPF Interface Parameters" on page 57.

You can verify that the OSPF routing instance on the NetScreen device has established adjacencies with OSPF neighbors by executing the following CLI command:

```
ns-> get vrouter trust-vr protocol ospf neighbor
VR: trust-vr RouterId: 10.1.1.250
---------------------------------
                Neighbor(s) on interface ethernet3 (Area 0.0.0.0)
IpAddr/If Index RouterId        Priority State    Options
-------------------------------------------------------------------------------
2.2.2.2         2.2.2.250             1 Full     E

                Neighbor(s) on interface ethernet1 (Area 0.0.0.10)
IpAddr/If Index RouterId        Priority State    Options
-------------------------------------------------------------------------------
10.1.1.1        10.1.1.252            1 Full     E
```

Indicates full OSPF adjacencies with neighbors on these interfaces.

REDISTRIBUTING ROUTES

Route redistribution is the exchange of route information between routing protocols. For example, you can redistribute the following types of routes into the OSPF routing instance in the same VR:

- Routes learned from BGP
- Directly connected routes
- Imported routes
- Statically configured routes

When you configure route redistribution, you must first specify a route map to filter the routes that are redistributed. For more information about creating route maps for route redistribution, refer to Chapter 1, “Virtual Routers”.

Example: Redistributing a BGP Route into OSPF

In the following example, you redistribute a route that originated from a BGP routing domain into the current OSPF routing domain. Both the CLI and WebUI examples assume that you previously created a route map called add-bgp.

WebUI

Network > Routing > Virtual Routers > Edit (for trust-vr) > Edit OSPF Instance > Redistributable Rules: Enter the following, and then click **Add**:

Route Map: add-bgp

Protocol: BGP

CLI

```
set vrouter trust-vr protocol ospf redistribute route-map add-bgp protocol bgp
save
```

Summarizing Redistributed Routes

In large internetworks where hundreds or even thousands of network addresses can exist, some routers may become overly congested with route information. Once you have redistributed a series of routes from an external protocol to the current OSPF routing instance, you can bundle the routes into one generalized or *summarized* network route. By summarizing multiple addresses, you enable a series of routes to be recognized as one route, simplifying the process.

An advantage to using route summarization in a large, complex network is that it can isolate topology changes from other routers. That is, if a specific link in a given domain is intermittently failing, the summary route would not change, so no router external to the domain would need to keep modifying its routing table due to the link failure.

In addition to creating fewer entries in the routing tables on the backbone routers, route summarization prevents the propagation of LSAs to other areas when one of the summarized networks goes down or comes up. You can summarize inter-area routes or external routes.

Example: Summarizing Redistributed Routes

In the following example, you redistribute BGP routes defined by the route-map add-bgp into the current OSPF routing instance. You then summarize the set of imported routes under the network address 2.1.1.0/24.

WebUI

Network > Routing > Virtual Routers > Edit (for trust-vr) > Edit OSPF Instance > Redistributable Rules: Enter the following, and then click **Add**:

Route Map: add-bgp

Protocol: BGP

Network > Routing > Virtual Routers > Edit (for trust-vr) > Edit OSPF Instance > Summary Import: Enter the following, and then click **Add**:

IP/Netmask: 2.1.1.0/24

CLI

```
set vrouter trust-vr protocol ospf redistribute route-map add-bgp protocol bgp
set vrouter trust-vr protocol ospf summary-import ip 2.1.1.0/24
save
```

GLOBAL OSPF PARAMETERS

This section describes optional OSPF global parameters that you can configure at the VR level. When you configure an OSPF parameter at the VR level, the parameter setting affects operations on all OSPF-enabled interfaces. You can modify global parameter settings through the OSPF routing protocol context in the CLI or by using the WebUI.

The following table describes the OSPF global parameters and their default values.

OSPF Global Parameter	Description	Default Value
Advertise default route	Specifies that an active default route (0.0.0.0/0) in the VR route table is advertised into all OSPF areas. You must also specify the metric and metric type (ASE type 1 or type 2) for the default route. You can also specify that the default route is always advertised.	Default route is not advertised.
Reject default route	Specifies that any default route learned in OSPF is not added to the route table.	Default route learned in OSPF is added to the route table.
Automatic virtual link	Specifies that the virtual router is to automatically create a virtual link when it cannot reach the OSPF backbone.	Disabled
Maximum hello packets	Specifies the maximum number of OSPF hello packets that the virtual router can receive in a hello interval.	10
Maximum LSA packets	Specifies the maximum number of OSPF LSA packets that the virtual router can receive within the specified number of seconds.	No default
RFC 1583 compatibility	Specifies that the ScreenOS OSPF routing instance is compatible with RFC 1583, an earlier version of OSPF.	ScreenOS supports OSPF version 2, as defined by RFC 2328.
Virtual link configuration	Configures the OSPF area and router ID for the virtual link. You can optionally configure the authentication method, hello interval, retransmit interval, transmit delay, or neighbor dead interval for the virtual link.	No virtual link configured.

Example: Advertising the Default Route

The default route, 0.0.0.0/0, matches every destination network in a routing table, although a more specific prefix overrides the default route.

In the following example, you advertise the current OSPF routing instance's default route.

WebUI

Network > Routing > Virtual Routers > Edit (for trust-vr) > Edit OSPF Instance: Select **Advertising Default Route Enable**, and then click **OK**.

Note: *In the WebUI, the default metric is 1 and the default metric-type is ASE type 1.*

CLI

```
set vrouter trust-vr protocol ospf advertise-default-route metric 1 metric-type 1
save
```

Virtual Links

All areas in an OSPF internetwork must connect directly to the backbone area. Sometimes, you need to create a new area that cannot be physically connected to the backbone area. To solve this problem, you configure a virtual link. The virtual link provides a remote area with a logical path to the backbone through another area.

You must configure the virtual link on the routers on both ends of the link. To configure a virtual link on the NetScreen device, you need to define:

- The ID of the OSPF area through which the virtual link will pass. You cannot create a virtual link that passes through the backbone area or a stub area.
- The ID of the router at the other end of the virtual link.

You can optionally configure the following parameters for virtual links:

Virtual Link Parameter	Description	Default Value
Authentication	Specifies either clear text password or MD5 authentication.	No authentication used.
Dead interval	Specifies the number of seconds that elapses with no response from an OSPF neighbor before the neighbor is determined to be not running.	40 seconds
Hello interval	Specifies the number of seconds between OSPF hellos.	10 seconds
Retransmit interval	Specifies the number of seconds that elapses before the interface resends an LSA to a neighbor that did not respond to the original LSA.	5 seconds
Transmit delay	Specifies the number of seconds between transmissions of link-state update packets sent on an interface.	1 second

Example: Creating a Virtual Link

In the following example, you create a virtual link through OSPF area 10 from NetScreen-A with router ID 10.10.1.250 to NetScreen-B with router ID 10.1.1.250. (See Chapter 1, "Virtual Routers" for information on how to configure router IDs on NetScreen devices.) You also configure the virtual link for a transmit delay of 10 seconds. On each NetScreen device, you need to identify the router ID of the device at the other end of the virtual link.

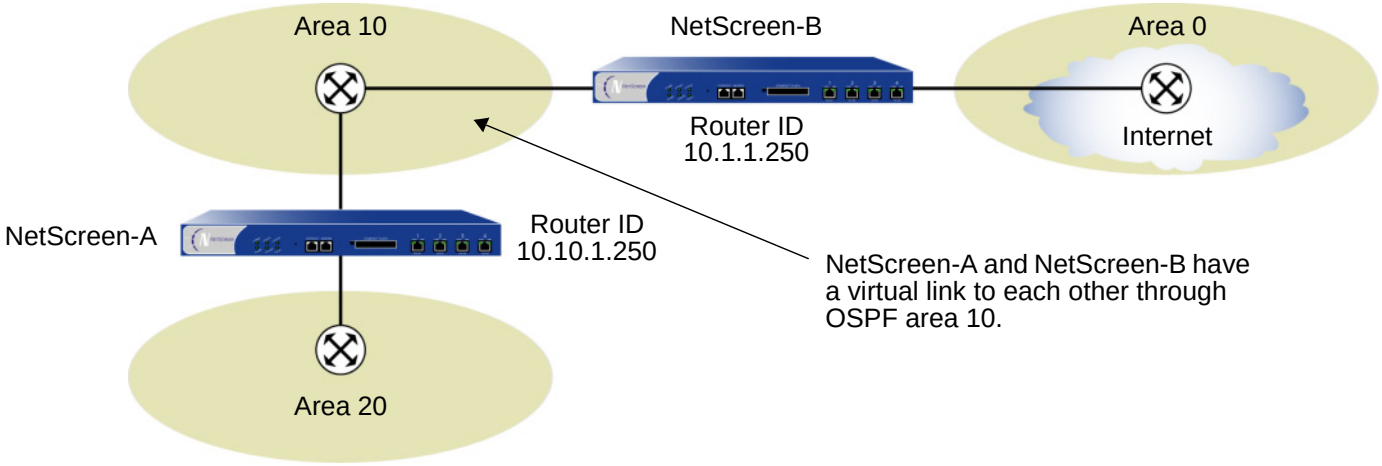

WebUI (NetScreen-A)

Network > Routing > Virtual Routers > Edit (for trust-vr) > Edit OSPF Instance > Virtual Link: Enter the following, and then click **Add**:

Area ID: 10 (select)

Router ID: 10.1.1.250

> Configure: In the Transmit Delay field, type **10**, and then click **OK**.

CLI (NetScreen-A)

```
set vrouter trust-vr protocol ospf vlink area-id 10 router-id 10.1.1.250
set vrouter trust-vr protocol ospf vlink area-id 10 router-id 10.1.1.250
    transit-delay 10
save
```

Note: *In the CLI, you must first create the virtual link before you can configure any optional parameters for the virtual link. Thus, in the CLI example above, you must issue two separate commands to create and then configure the virtual link.*

WebUI (NetScreen-B)

Network > Routing > Virtual Routers > Edit (for trust-vr) > Edit OSPF Instance > Virtual Link: Enter the following, and then click **Add**:

Area ID: 10

Router ID: 10.10.1.250

> Configure: In the Transmit Delay field, type **10**, and then click **OK**.

CLI (NetScreen-B)

```
set vrouter trust-vr protocol ospf vlink area-id 10 router-id 10.10.1.250
set vrouter trust-vr protocol ospf vlink area-id 10 router-id 10.10.1.250
    transit-delay 10
save
```

Example: Creating an Automatic Virtual Link

You can direct a VR to automatically create a virtual link for instances when it cannot reach the network backbone. Having the VR automatically create virtual links replaces the more time-consuming process of creating each virtual link manually. In the following example, you configure automatic virtual link creation.

WebUI

Network > Routing > Virtual Routers > Edit (for trust-vr) > Edit OSPF Instance: Select **Automatically Generate Virtual Links**, and then click **OK**.

CLI

```
set vrouter trust-vr protocol ospf auto-vlink
save
```

OSPF Interface Parameters

This section describes OSPF parameters that you configure at the interface level. When you configure an OSPF parameter at the interface level, the parameter setting affects the OSPF operation only on the specific interface. You can modify interface parameter settings with **interface** commands in the CLI or by using the WebUI.

The following table describes optional OSPF interface parameters and their default values.

OSPF Interface Parameter	Description	Default Value
Authentication	Specifies either clear text password or message digest 5 (MD5) authentication to verify OSPF communication on the interface. A clear text password requires password string of up to 8 digits, and an MD5 authentication password requires a password string of up to 16 digits. The MD5 password also requires that you configure key strings.	No authentication used
Cost	Specifies the metric for the interface. The cost associated with an interface depends upon the bandwidth of the link to which the interface is connected. The higher the bandwidth, the lower (more desirable) the cost value.	1 for a 100MB or more link 10 for a 10MB link 100 for a 1MB link
Dead interval	Specifies the number of seconds that elapses with no response from an OSPF neighbor before OSPF determines the neighbor is not running.	40 seconds
Hello interval	Specifies the interval, in seconds, at which OSPF sends out hello packets to the network.	10 seconds
Link type	Specifies the interface as a point-to-point link.	Ethernet interfaces are treated as broadcast interfaces
Neighbor list	Specifies subnets, in the form of one or more access lists, on which OSPF neighbors reside that are eligible to form adjacencies.	None (adjacencies are formed with all neighbors on the interface)

OSPF Interface Parameter	Description	Default Value
Passive Interface	Specifies that the IP address of the interface is advertised into the OSPF domain as an OSPF route and not as an external route, but the interface does not transmit or receive OSPF packets. This option is useful when BGP is also enabled on the interface.	OSPF-enabled interfaces transmit and receive OSPF packets
Priority	Specifies the priority for the virtual router to be elected the Designated Router or Backup Designated Router. The router with the larger priority value has the best chance (although not guaranteed) chance of being elected.	1
Retransmit interval	Specifies the number of seconds that elapses before the interface resends an LSA to a neighbor that did not respond to the original LSA.	5 seconds
Transit delay	Specifies the number of seconds between transmissions of link-state update packets sent on the interface.	1 second

Note: *To form adjacencies, all OSPF routers in an area must use the same hello, dead, and retransmit interval values.*

Example: Setting OSPF Interface Parameters

In this example, you configure the following OSPF parameters for the ethernet1 interface:

- Increase the interval between OSPF hello messages to 15 seconds.
- Increase the interval between OSPF retransmissions to 7 seconds.
- Increase the interval between LSA transmissions to 2 seconds.

WebUI

Network > Interfaces > Edit (for ethernet1) > OSPF: Enter the following, and then click **Apply**:

Hello Interval: 15

Retransmit Interval: 7

Transit Delay: 2

CLI

```
set interface ethernet1 protocol ospf hello-interval 15
set interface ethernet1 protocol ospf retransmit-interval 7
set interface ethernet1 protocol ospf transit-delay 2
save
```

SECURITY CONFIGURATION

This section describes possible security problems in the OSPF routing domain and methods of preventing attacks.

Note: *To make OSPF more secure, you should configure all routers in the OSPF domain to be at the same security level. Otherwise, a compromised OSPF router can bring down the entire OSPF routing domain.*

Authenticating Neighbors

An OSPF router can be easily spoofed, since LSAs are not encrypted and most protocol analyzers provide decapsulation of OSPF packets. Authenticating OSPF neighbors is the best way to fend off these types of attacks.

OSPF provides both simple password and MD5 authentication to validate OSPF packets received from neighbors. All OSPF packets received on the interface that are not authenticated are discarded. By default, there is no authentication enabled on any OSPF interface.

MD5 authentication requires that the same key be used for both the sending and receiving OSPF routers. You can specify more than one MD5 key on the NetScreen device; each key is paired with a key identifier. If you configure multiple MD5 keys on the NetScreen device, you can then select the key identifier of the key that is to be used for authenticating communications with the neighbor router. This allows MD5 keys on pairs of routers to be changed periodically with minimal risk of packets being dropped.

Example: Configuring the Clear-Text Password Authentication Method

In this example, you set a clear-text password 12345678 for OSPF on interface ethernet1.

WebUI

Network > Interfaces > Edit (for ethernet1) > OSPF: Enter the following, and then click **Apply**:

Password: (select), 12345678

CLI

```
set interface ethernet1 protocol ospf authentication password 12345678
save
```

Example: Configuring the MD5 Password Authentication Method

In the following example, you set the two different MD5 keys on interface ethernet1 and select one of the keys to be the active key. Note that the default key-id is 0 so you do not have to specify the key-id for the first MD5 key you enter.

WebUI

Network > Interfaces > Edit (for ethernet1) > OSPF: Enter the following, and then click **Apply**:

Authentication:

MD5 Keys: (select)

1234567890123456

9876543210987654

Key ID: 1

Preferred: (select)

CLI

```
set interface ethernet1 protocol ospf authentication md5 1234567890123456
set interface ethernet1 protocol ospf authentication md5 9876543210987654
    key-id 1
set interface ethernet1 protocol ospf authentication md5 active-md5-key-id 1
save
```

Filtering OSPF Neighbors

Multi-access environments can allow devices, including routers, to be connected into a network relatively easily. This can cause stability or performance issues if the connected device is not reliable.

By default, the OSPF routing instance on the NetScreen virtual router forms adjacencies with all OSPF neighbors communicating on an OSPF-enabled interface. You can limit the devices on an interface that can form adjacencies with the OSPF routing instance by defining a list of subnets that contain eligible OSPF neighbors. Only hosts or routers that reside in the specified subnets can form adjacencies with the OSPF routing instance. To specify the subnets that contain eligible OSPF neighbors, define an access list for the subnets at the virtual router level.

Example: Configuring a Neighbor List

In this example, you configure an access list that permits the hosts on subnet 10.10.10.130/27. You then specify the access list to configure eligible OSPF neighbors.

WebUI

Network > Routing > Virtual Router (trust-vr) > Access List > New: Enter the following, and then click **OK**:

Access List ID: 4

Sequence No.: 10

IP/Netmask: 10.10.10.130/27

Action: Permit (select)

Network > Interfaces > Edit (for ethernet1) > OSPF: Enter the following, and then click **Apply**:

Neighbor List: 4

CLI

```
set vrouter trust-vr access-list 4
set vrouter trust-vr access-list 4 permit ip 10.10.10.130/27 10
set interface ethernet1 protocol ospf neighbor-list 4
save
```

Rejecting Default Routes

In a Route Detour Attack, a router injects a default route (0.0.0.0/0) into the routing domain in order to detour packets to itself. The router can then either drop the packets, causing service disruption, or it can obtain sensitive information in the packets before forwarding them. On NetScreen devices, OSPF by default accepts any default routes that are learned in OSPF and adds the default route to the routing table.

Example: Removing the Default Route from the Route Table

In the following example, you specify that a default route not be learned from OSPF.

WebUI

Network > Routing > Virtual Routers > Edit (for trust-vr) > Edit OSPF Instance: Select the **Do Not Add Default-route Learned in OSPF** check box and then click **OK**.

CLI

```
set vrouter trust-vr protocol ospf reject-default-route
save
```

Protecting against Flooding

A malfunctioning or compromised router can flood its neighbors with OSPF hello packets or with LSAs. LSAs enable OSPF routers to provide device, network, and routing information for the link state database. Each router retrieves information from the LSAs sent by other routers on the network to distill path information for the routing table. LSA flood protection enables you to manage the number of LSAs entering the virtual router. If the virtual router receives too many LSAs, the router fails because of LSA flooding. An LSA attack happens when a router generates an excessive number of LSAs in a short period of time, thus keeping other OSPF routers in the network busy running the SPF algorithm.

On NetScreen virtual routers, you can configure both the maximum number of hello packets per hello interval and the maximum number of LSAs that can be received on an OSPF interface within a certain interval. Packets that exceed a configured threshold are dropped. By default, the OSPF hello packet threshold is 10 packets per hello interval (the default hello interval for an OSPF interface is 10 seconds). There is no default LSA threshold; if you do not set an LSA threshold, all LSAs are accepted.

Example: Configuring the Hello Threshold

In the following example, you configure a threshold of 20 packets per hello interval. The hello interval, which is configurable on each OSPF interface, is not changed from its default of 10 seconds.

WebUI

Network > Routing > Virtual Routers > Edit (for trust-vr) > Edit OSPF Instance: Enter the following, and then click **OK**:

Prevent Hello Packet Flooding Attack: On

Max Hello Packet: 20

CLI

```
set vrouter trust-vr protocol ospf hello-threshold 20
save
```

Example: Configuring the LSA Threshold

In this example, you create an OSPF LSA flood attack threshold of 10 packets per 10 seconds.

WebUI

Network > Routing > Virtual Routers > Edit (for trust-vr) > Edit OSPF Instance: Enter the following, and then click **OK**:

LSA Packet Threshold Time: 10

Maximum LSAs: 10

CLI

```
set vrouter trust-vr protocol ospf lsa-threshold 10 10
save
```

3

Routing Information Protocol (RIP)

This chapter describes the Routing Information Protocol (RIP) version 2 routing protocol on NetScreen devices. The following topics are covered:

- "Overview of RIP" on page 68
- "Basic RIP Configuration" on page 69
  - "Creating a RIP Routing Instance in a Virtual Router" on page 70
  - "Enabling RIP on Interfaces" on page 72
  - "Redistributing Routes" on page 73
- "Global RIP Parameters" on page 76
- "RIP Interface Parameters" on page 78
- "Security Configuration" on page 80
  - "Authenticating Neighbors" on page 80
  - "Filtering RIP Neighbors" on page 82
  - "Rejecting Default Routes" on page 83
  - "Protecting Against Flooding" on page 84

OVERVIEW OF RIP

Routing information protocol (RIP) is a distance vector protocol used as an Interior Gateway Protocol (IGP) in moderate-sized autonomous systems (AS). ScreenOS supports RIP version 2 (RIPv2), as defined by RFC 2453. While RIPv2 supports only simple password (plain text) authentication, NetScreen's RIP implementation also supports MD5 authentication extensions, as defined by RFC 2082.

As mentioned previously, RIP is intended for moderate-sized networks. It can also be used to manage route information within a small, homogeneous, network such as a corporate LAN. The longest path allowed in a RIP network is 15 hops. A metric value of 16 indicates an invalid or unreachable destination (this value is also referred to as "infinity" since it is larger than the 15-hop maximum allowed in RIP networks).

RIP is not intended for large networks or networks where routes are chosen based on real-time parameters such as measured delay, reliability, or load. RIP supports both point-to-point networks (used with VPNs) and broadcast/multicast Ethernet networks. RIP does not support point-to-multipoint interfaces.

RIP sends out messages that contain the complete routing table to every neighboring router every 30 seconds. These messages are normally sent as multicasts to address 224.0.0.9 from the RIP port.

The RIP routing database contains one entry for every destination that is reachable through the RIP routing instance. The RIP routing database includes the following information:

- IPv4 address of a destination. Note that RIP does not distinguish between networks and hosts.
- IP address of the first router along the route to the destination (the next hop).
- Network interface used to reach the first router.
- Metric that indicates the distance, or cost, of getting to the destination. Most RIP implementations use a metric of 1 for each network.
- A timer that indicates the time that has elapsed since the database entry was last updated.

BASIC RIP CONFIGURATION

Like OSPF and BGP, you create RIP on a per-Virtual Router basis on a NetScreen device. If you have multiple virtual routers (VRs) in a system, you can enable multiple instances of RIP, one instance for each VR.

Note: *Before you configure a dynamic routing protocol on the NetScreen device, you should assign a virtual router ID, as described in Chapter 1, "Virtual Routers".*

This section describes the following basic steps to configure RIP on a NetScreen device:

1. Create the RIP routing instance in a Virtual Router.
2. Enable the RIP instance.
3. Enable RIP on interfaces that connect to other RIP routers.
4. Redistribute routes learned from different routing protocols (such as OSPF, BGP, or statically configured routes) into the RIP instance.

This section describes how to perform each of these tasks using either the CLI or the WebUI.

You can also optionally configure other RIP parameters such as the following:

- Global parameters, such as timers and trusted RIP neighbors, that are set at the VR level for the RIP protocol (see "Global RIP Parameters" on page 76)
- Interface parameters, such as neighbor authentication, that are set on a per-interface basis for the RIP protocol (see "RIP Interface Parameters" on page 78)
- Security-related RIP parameters, that are set at either the VR level or on a per-interface basis (see "Security Configuration" on page 80)

Creating a RIP Routing Instance in a Virtual Router

You create and enable a RIP routing instance on a specific virtual router on a NetScreen device. When you create and enable a RIP routing instance on a VR, RIP can transmit and receive packets on all RIP-enabled interfaces in the VR.

Deleting a RIP routing instance in a VR removes the corresponding RIP configurations for all interfaces that are in the VR. For more information about virtual routers and configuring a virtual router on NetScreen devices, see Chapter 1, “Virtual Routers”.

Example: Creating a RIP Routing Instance

In the following example, you first assign 0.0.0.10 as the router ID for the trust-vr virtual router. You then create create a RIP routing instance on the trust-vr.

WebUI

1. Router ID

   Network > Routing > Virtual Router (trust-vr) > Edit: Enter the following, and then click **OK**:

   Virtual Router ID: Custom (select)

   In the text box, enter 0.0.0.10

2. RIP Routing Instance

   Network > Routing > Virtual Router (trust-vr) > Edit: Select **Create RIP Instance**.

   Select Enable RIP, and then click **OK**.

CLI

1. Router ID

```
set vrouter trust-vr router-id 10
```

2. RIP Routing Instance

```
set vrouter trust-vr protocol rip
set vrouter trust-vr protocol rip enable
save
```

> ***Note:*** *In the CLI, you must first create the RIP routing instance before you can enable it. Thus, you must issue two separate CLI commands to enable a RIP routing instance.*

Example: Removing a RIP Routing Instance

In this example, you disable the RIP routing instance in the trust-vr. RIP stops transmitting and processing packets on all RIP-enabled interfaces in the trust-vr.

WebUI

Network > Routing > Virtual Router (trust-vr) > Edit > Edit RIP Instance: Deselect Enable RIP, and then click **OK**.

Network > Routing > Virtual Router (trust-vr) > Edit > Delete RIP Instance, and enter **OK** at the confirmation prompt.

CLI

```
unset vrouter trust-vr protocol rip enable
unset vrouter trust-vr protocol rip
save
```

> ***Note:*** *In the CLI, you must first disable the RIP routing instance before you can delete it. Thus, you must issue two separate CLI commands to remove a RIP routing instance.*

Enabling RIP on Interfaces

By default, RIP is disabled on all interfaces in the VR and you must explicitly enable it on an interface. When you disable RIP at the interface level, RIP does not transmit or receive packets on the specified interface. Interface configuration parameters are preserved when you disable RIP on an interface.

Note: If you disable the RIP routing instance in the VR (see "Example: Removing a RIP Routing Instance" on page 71), RIP stops transmitting and processing packets on all RIP-enabled interfaces in the VR.

Example: Enabling RIP on Interfaces

In this example, you enable RIP on the Trust interface.

WebUI

Network > Interface > Edit (for Trust) > RIP: Select Protocol RIP **Enable**, and then click **Apply**.

CLI

```
set interface trust protocol rip enable
save
```

Example: Disabling RIP on an Interface

In this example, you disable RIP on the Trust interface.

WebUI

Network > Interface (for Trust) > RIP: Clear Protocol RIP **Enable**, and then click **Apply**.

CLI

```
unset interface trust protocol rip
save
```

Redistributing Routes

Route redistribution is the exchange of route information between routing protocols. For example, you can redistribute the following types of routes into the RIP routing instance in the same virtual router:

- Routes learned from BGP
- Routes learned from OSPF
- Directly connected routes
- Imported routes
- Statically configured routes

You need to configure a route map to filter the routes that are redistributed. For more information about creating route maps for route redistribution, see Chapter 1, “Virtual Routers”.

Routes imported into RIP from other protocols have a default metric of 1. You can change the default metric (see “Global RIP Parameters” on page 76).

Example: Redistributing Routes into RIP

In this example, you redistribute static routes that are in the subnetwork 20.1.0.0/16 to RIP neighbors in the trust-vr virtual router. To do this, you first create an access list to permit addresses in the 20.1.0.0/16 subnetwork. Then, configure a route map that permits addresses that match the access list you configured. Use the route map to specify the redistribution of static routes into the RIP routing instance.

WebUI

Network > Routing > Virtual Router (trust-vr) > Access List > New: Enter the following, and then click **OK**:

Access List ID: 20

Sequence No.: 1

IP/Netmask: 20.1.0.0/16

Action: Permit (select)

Network > Routing > Virtual Router (trust-vr) > Route Map > New: Enter the following, and then click **OK**:

Map Name: rtmap1

Sequence No.: 1

Action: Permit (select)

Match Properties:

Access List: (select), 20 (select)

Network > Routing > Virtual Router (trust-vr) > Edit > Edit RIP Instance > Redistributable Rules: Enter the following, and then click **Add**:

Route Map: rtmap1 (select)

Protocol: Static (select)

CLI

```
set vrouter trust-vr acc-list 20 permit ip 20.1.0.0/16 1
set vrouter trust-vr route-map name rtmap1 permit 1
set vrouter trust-vr route-map rtmap1 1 match ip 20
set vrouter trust-vr protocol rip redistribute route-map rtmap1 protocol static
save
```

GLOBAL RIP PARAMETERS

This section describes RIP global parameters that you can configure at the VR level. When you configure a RIP parameter at the VR level, the parameter setting affects operations on all RIP-enabled interfaces. You can modify global parameter settings through the RIP routing protocol context in the CLI or by using the WebUI.

The following table describes the RIP global parameters and their default values.

RIP Global Parameter	Description	Default Value
Default metric	Default metric value for routes imported into RIP from other protocols, such as OSPF and BGP.	10
Update timer	Specifies, in seconds, when to issue updates of RIP routes to neighbors.	30 seconds
Maximum packets per update	Specifies the maximum number of packets received per update.	No maximum
Invalid timer	Specifies, in seconds, when a route becomes invalid from the time a neighbor stops advertising the route.	180 seconds
Flush timer	Specifies, in seconds, when a route is removed from the time the route is invalidated.	120 seconds
Maximum neighbors	The maximum number of RIP neighbors allowed.	16
Trusted neighbors	Specifies an access list that defines RIP neighbors. If no neighbors are specified, RIP uses multicasting or broadcasting to detect neighbors on an interface.	All neighbors are trusted
Allow neighbors on different subnet.	Specifies that RIP neighbors on different subnets are allowed.	Disabled
Advertise default route	Specifies whether the default route (0.0.0.0/0) is advertised.	Disabled
Reject default route	Specifies whether RIP rejects a default route learned from another protocol.	Disabled
Incoming route map	Specifies the filter for routes to be learned by RIP.	None
Outgoing route map	Specifies the filter for routes to be advertised by RIP.	None

Example: Advertising the Default Route to RIP Neighbors

By default, the default route (0.0.0.0/0) is not advertised to RIP neighbors. The following command advertises the default route to RIP neighbors in the trust-vr virtual router with a metric of 5 (you must enter a metric value). The default route must exist in the routing table.

WebUI

Network > Routing > Virtual Router (trust-vr) > Edit > Edit RIP Instance: Enter the following, and then click **OK**:

Advertising Default Route: (select)

Metric: 5

CLI

```
set vrouter trust-vr protocol rip adv-default-route metric number 5
save
```

Note: *See the* NetScreen CLI Reference Guide *for more information about global parameters that you can configure in the RIP routing protocol context.*

RIP INTERFACE PARAMETERS

This section describes RIP parameters that you configure at the interface level. When you configure a RIP parameter at the interface level, the parameter setting affects the RIP operation only on the specific interface. You can modify interface parameter settings with **interface** commands in the CLI or by using the WebUI.

The following table describes the RIP interface parameters and their default values.

RIP Interface Parameter	Description	Default Value
Split-horizon	Specifies whether to enable split-horizon (do not advertise routes learned from an interface in updates sent to the same interface). If split horizon is enabled with the poison-reverse option, routes that are learned from an interface are advertised with a metric of 16 in updates sent to the same interface.	Disabled
RIP metric	Specifies the RIP metric for the interface.	1
Authentication	Specifies either clear text password or MD5 authentication.	No authentication used.
Passive mode	Specifies that the interface is to receive but not transmit RIP packets.	No
Incoming route map	Specifies the filter for routes to be learned by RIP.	None
Outgoing route map	Specifies the filter for routes to be advertised by RIP.	None

You can define incoming and outgoing route map filters at the VR level or at the interface level. A route map filter you define at the interface level takes precedence over a route map filter defined at the VR level. For example, if you define an incoming route map at the VR level and a different incoming route map at the interface level, the incoming route map defined at the interface level takes precedence.

Example: Setting RIP Interface Parameters

In this example, you configure the following RIP parameters for the trust interface:

- Set MD5 authentication, with the key 1234567898765432 and the key ID 215.
- Enable split horizon with poison reverse for the interface.

WebUI

Network > Interfaces > Edit (for Trust) > RIP: Enter the following, and then click **OK**:

Authentication: MD5 (select)

Key: 1234567898765432

Key ID: 215

Split Horizon: Enabled with poison reverse (select)

CLI

```
set interface trust protocol rip authentication md5 1234567898765432 key-id 215
set interface trust protocol rip split-horizon poison-reverse
save
```

SECURITY CONFIGURATION

This section describes possible security problems in the RIP routing domain and methods of preventing attacks.

Note: *To make RIP more secure, you should configure all routers in the RIP domain to be at the same security level. Otherwise, a compromised RIP router can bring down the entire RIP routing domain.*

Authenticating Neighbors

A RIP router can be easily spoofed, since RIP packets are not encrypted and most protocol analyzers provide decapsulation of RIP packets. Authenticating RIP neighbors is the best way to fend off these types of attacks.

RIP provides both simple password and MD5 authentication to validate RIP packets received from neighbors. All RIP packets received on the interface that are not authenticated are discarded. By default, there is no authentication enabled on any RIP interface.

MD5 authentication requires that the same key be used for both the sending and receiving RIP routers. You can specify more than one MD5 key on the NetScreen device; each key is paired with a key identifier. If you configure multiple MD5 keys on the NetScreen device, you can then select the key identifier of the key that is to be used for authenticating communications with the neighbor router. This allows MD5 keys on pairs of routers to be changed periodically with minimal risk of packets being dropped.

Example: Configuring the MD5 Password Authentication Method

In the following example, you set the two different MD5 keys on interface ethernet1 and select one of the keys to be the active key. Note that the default key-id is 0 so you do not have to specify the key-id for the first MD5 key you enter.

WebUI

Network > Interfaces > Edit (for ethernet1) > RIP: Enter the following, and then click **Apply**:

MD5 Keys: (select)

1234567890123456 (first key field)

9876543210987654 (second key field)

Key ID: 1

Preferred: (select)

CLI

```
set interface ethernet1 protocol rip authentication md5 1234567890123456
set interface ethernet1 protocol rip authentication md5 9876543210987654 key-id 1
set interface ethernet1 protocol rip authentication md5 active-md5-key-id 1
save
```

Filtering RIP Neighbors

Multi-access environments can allow devices, including routers, to be connected into a network relatively easily. This can cause stability or performance issues if the connected device is not reliable. To prevent this problem, you can use an access list to filter the devices that are allowed to become RIP neighbors. By default, RIP neighbors are limited to devices that are on the same subnet as the NetScreen virtual router.

Example: Configuring Trusted Neighbors

In this example, you configure the following global parameters for the RIP routing instance running in the trust-vr virtual router:

- Maximum number of RIP neighbors is 1.
- The IP address of the trusted neighbor, 10.1.1.1, is specified in an access-list.

WebUI

Network > Routing > Virtual Router (trust-vr) > Access List > New: Enter the following, and then click **OK**:

Access List ID: 10

Sequence No.: 1

IP/Netmask: 10.1.1.1/32

Action: Permit (select)

Network > Routing > Virtual Router (trust-vr) > Edit > Edit RIP Instance: Enter the following, and then click **OK**:

Trusted Neighbors: (select), 10

Maximum Neighbors: 1

CLI

```
set vrouter trust-vr
ns(trust-vr)-> set access-list 10 permit ip 10.1.1.1/32 1
ns(trust-vr)-> set protocol rip
ns(trust-vr/rip)-> set max-neighbor-count 1
ns(trust-vr/rip)-> set trusted-neighbors 10
ns(trust-vr/rip)-> exit
ns(trust-vr)-> exit
save
```

Rejecting Default Routes

In a Route Detour Attack, a router injects a default route (0.0.0.0/0) into the routing domain in order to detour packets to itself. The router can then either drop the packets, causing service disruption, or it can obtain sensitive information in the packets before forwarding them. On NetScreen devices, RIP by default accepts any default routes that are learned in RIP and adds the default route to the routing table.

Example: Rejecting Default Routes

In this example, you configure the RIP routing instance running in the trust-vr virtual router to reject any default routes that are learned in RIP.

WebUI

Network > Routing > Virtual Router (trust-vr) > Edit > Edit RIP Instance: Enter the following, and then click **OK**:

Reject Default Route Learnt by RIP: (select)

CLI

```
set vrouter trust-vr protocol rip reject-default-route
save
```

Protecting Against Flooding

A malfunctioning or compromised router can flood its neighbors with RIP routing update packets. On NetScreen virtual routers, you can configure the maximum number of update packets that can be received on a RIP interface within a certain interval to avoid flooding of update packets. All update packets that exceed the configured update threshold are dropped. If you do not set an update threshold, all update packets are accepted.

You need to exercise care when configuring an update threshold when neighbors have large routing tables, as the number of routing updates can be quite high within a given duration because of flash updates. Update packets that exceed the threshold are dropped and valid routes may not be learned.

Example: Configuring an Update Threshold

In this example, you set the maximum number of routing update packets that RIP can receive on an interface to 4.

WebUI

Network > Routing > Virtual Router (trust-vr) > Edit > Edit RIP Instance: Enter the following, and then click **OK**:

Maximum Number Packets per Update Time: (select), 4

CLI

```
set vrouter trust-vr protocol rip threshold-update 4
save
```

Example: RIP on Tunnel Interfaces

The following example creates and enables a RIP routing instance in the Trust-VR virtual router on the NetScreen-A device. You enable RIP on both the VPN tunnel interface and the Trust zone interface. Only routes that are in the subnet 10.10.0.0/16 are advertised to the RIP neighbor on NetScreen-B. This is done by first configuring an access list that permits only addresses in the subnet 10.10.0.0/16, then specifying a route map *abcd* that permits routes that match the access list. You then specify the route map to filter the routes that are advertised to RIP neighbors.

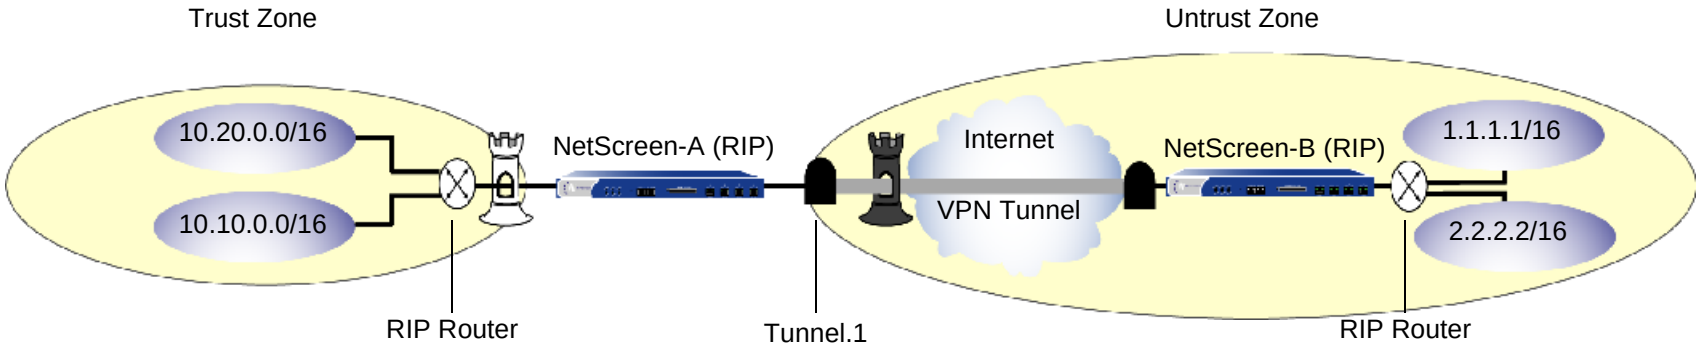

WebUI

Network > Routing > Virtual Router > Edit (for trust-vr) > Create RIP Instance: Select **Enable RIP**, and then click **OK**.

Network > Routing > Virtual Router > Access List (for trust-vr) > New: Enter the following, and then click **OK**:

Access List ID: 10

Sequence No.: 10

IP/Netmask: 10.10.0.0/16

Action: Permit

Network > Routing > Virtual Router > Route Map (for trust-vr) > New: Enter the following, and then click **OK**:

Map Name: abcd

Sequence No.: 10

Action: Permit

Match Properties:

Access List: (select), 10

Network > Routing > Virtual Router > Edit (for trust-vr) > Edit RIP Instance: Select the following, and then click **OK**:

Outgoing Route Map Filter: abcd

Network > Interfaces > Edit (for tunnel.1) > RIP: Enter the following, and then click **Apply**:

Enable RIP: (select)

Network > Interfaces > Edit (for trust) > RIP: Enter the following, and then click **Apply**:

Enable RIP: (select)

CLI

```
set vrouter trust-vr protocol rip
set vrouter trust-vr protocol rip enable
set interface tunnel.1 protocol rip enable
set interface trust protocol rip enable
set vrouter trust-vr access-list 10 permit ip 10.10.0.0/16 10
set vrouter trust-vr route-map name abcd permit 10
set vrouter trust-vr route-map abcd 10 match ip 10
set vrouter trust-vr protocol rip route-map abcd out
save
```

4

Border Gateway Protocol (BGP)

This chapter describes the Border Gateway Protocol (BGP) on NetScreen devices. The following topics are covered:

- "Overview of BGP" on page 88
  - "Types of BGP Messages" on page 89
  - "Path Attributes" on page 89
  - "External and Internal BGP" on page 90
- "Basic BGP Configuration" on page 91
  - "Creating and Enabling a BGP Routing Instance in a Virtual Router" on page 92
  - "Enabling BGP on Interfaces" on page 94
  - "Configuring a BGP Peer" on page 95
  - "Verifying the BGP Configuration" on page 100
- "Security Configuration" on page 102
  - "Authenticating Neighbors" on page 102
  - "Rejecting Default Routes" on page 103
- "Optional BGP Configurations" on page 104
  - "Redistributing Routes" on page 105
  - "AS-Path Access List" on page 106
  - "Route Reflection" on page 107
  - "Confederations" on page 110
  - "BGP Communities" on page 113

OVERVIEW OF BGP

The Border Gateway Protocol (BGP) is a path vector protocol that is used to carry routing information between Autonomous Systems[1] (ASs). The BGP routing information includes the sequence of AS numbers that a network prefix (a route) has traversed. The path information that is associated with the prefix is used to enable loop prevention and enforce routing policies. ScreenOS supports BGP version 4 (BGP-4), as defined in RFC 1771.

Two *BGP peers* establish a *BGP session* in order to exchange routing information. A BGP router can participate in BGP sessions with different peers. BGP peers must first establish a TCP connection between themselves to open a BGP session. Upon forming the initial connection, peers exchange entire routing tables. As routing table changes occur, BGP routers exchange update messages with peers. A BGP router maintains current versions of the routing tables of all the peers with which it has sessions, periodically sending keepalive messages to peers to verify the connections.

A BGP peer only advertises those routes that it is actively using. When a BGP peer advertises a route to its neighbor, it also includes path attributes that describe the characteristics of the route. A BGP router compares the path attributes and prefix to select the best route from all paths that are available to a given destination.

1. An autonomous system is a set of routers that are in the same administrative domain.

Types of BGP Messages

BGP uses four different types of messages to communicate with peers:

- **Open** messages identify BGP peers to each other to initiate the BGP session. These messages are sent after the peers establish a TCP session. During the exchange of open messages, BGP peers specify their protocol version, AS number, hold time and BGP identifier.
- **Update** messages announce routes to the peer and withdraw previously-advertised routes.
- **Notification** messages indicate errors. The BGP session is terminated and then the TCP session is closed.

  Note: *The NetScreen device does not send a Notification message to a peer if, during the exchange of open messages, the peer indicates that it supports protocol capabilities that the NetScreen device does not support.*

- **Keepalive** messages are used to maintain the BGP session. By default, the NetScreen device sends keepalive messages to peers at 60-second intervals. This interval is configurable.

Path Attributes

BGP path attributes are a set of parameters that describe the characteristics of a route. BGP couples the attributes with the route they describe, then compares all paths available to a destination to select the best route to use to reach the destination. The path attributes are:

- **Origin** describes where the route was learned—it can be IGP, EGP, or incomplete.
- **AS-Path** contains a list of autonomous systems through which the route advertisement has passed.
- **Next-Hop** is the IP address of the router to which traffic for the route is sent.
- **Multi-Exit Discriminator** (MED) is a metric for a path where there are multiple links between ASs (the MED is set by one AS and used by another AS to choose a path).
- **Local-Pref** is a metric used to inform BGP peers of the local router's preference for the route.
- **Atomic-Aggregate** informs BGP peers that the local router selected a less-specific route from a set of overlapping routes received from a peer.

- **Aggregator** specifies the AS and router that performed aggregation of the route.
- **Communities** specifies one or more communities to which this route belongs
- **Cluster List** contains a list of the reflector clusters through which the route has passed

Many of the path attribute values are required or optional settings that you configure for the BGP routing instance. For example, for routes that are advertised by BGP on the NetScreen device, the AS number that you specify when you create the BGP routing instance is appended to the AS-Path attribute. A BGP router can choose to add or modify the path attributes before advertising the route to peers.

External and Internal BGP

External BGP (EBGP) is used between autonomous systems, as when different ISP networks connect to each other or an enterprise network connects to an ISP network. Internal BGP (IBGP) is used within an AS, such as within an enterprise network. The main goal of IBGP is to distribute the routes learned from EBGP to routers in the AS. Therefore, an IBGP router can readvertise routes that it learns from its EBGP peers to its IBGP peers, but it cannot advertise routes learned from IBGP peers to other IBGP peers. This restriction prevents route advertisement loops within the network, but means that an IBGP network must be fully-meshed (that is, every BGP router in the network must have a session with every other router in the network).

Some path attributes are only applicable to EBGP or IBGP. For example, the MED attribute is only used for EBGP messages, while the LOCAL-PREF attribute is only present in IBGP messages.

BASIC BGP CONFIGURATION

Like RIP and OSPF, you create a BGP instance on a per-virtual router basis on a NetScreen device. If you have multiple virtual routers on a device, you can enable multiple instances of BGP — one instance for each VR.

Note: *Before you configure a dynamic routing protocol on the NetScreen device, you should assign a virtual router ID, as described in Chapter 1, "Virtual Routers".*

This section describes the following basic steps to configure BGP in a virtual router on a NetScreen device:

1. Create and enable the BGP routing instance in a virtual router by first assigning an autonomous system number to the BGP instance, then enabling the instance.
2. Enable BGP on the interface that is connected to the peer.
3. Enable each BGP peer.
4. Configure one or more remote BGP peers.
5. Verify that BGP is properly configured and operating.

This section describes how to perform each of these tasks using either the CLI or the WebUI for the example shown below. In this example, you configure the NetScreen device as a BGP peer in AS 65000. The NetScreen device will establish a BGP session with the peer in AS 65500.

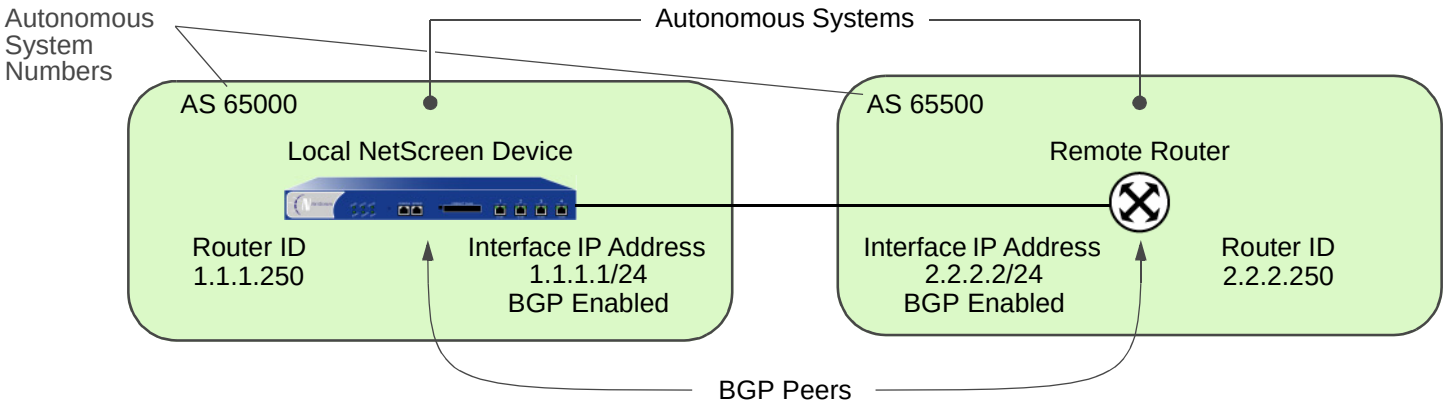

Creating and Enabling a BGP Routing Instance in a Virtual Router

You create and enable a BGP routing instance on a specific virtual router on a NetScreen device. To create a BGP routing instance, you need to first specify the autonomous system number[2] in which the virtual router resides. If the virtual router is an IBGP router, the autonomous system number must be the same as other IBGP routers in the network. When you enable the BGP routing instance on a VR, the BGP routing instance will be able to contact and establish a session with the BGP peers that you configure.

Example: Creating a BGP Routing Instance

In the following example, you first assign 0.0.0.10 as the router ID for the trust-vr virtual router. You then create and enable a BGP routing instance on trust-vr, which resides on the NetScreen device in AS 65000. (For more information about virtual routers and configuring a virtual router on NetScreen devices, see Chapter 1, “Virtual Routers”.)

WebUI

1. Router ID

Network > Routing > Virtual Router (trust-vr) > Edit: Enter the following, and then click **OK**:

Virtual Router ID: Custom (select)

In the text box, enter 0.0.0.10

2. BGP Routing Instance

Network > Routing > Virtual Routers > Edit (for trust-vr) > Create BGP Instance: Enter the following, and then click **OK**:

AS Number (required): 65000

BGP Enabled: (select)

2. Autonomous system numbers are globally unique numbers that are used to exchange EBGP routing information and to identify the AS. The following entities allocate AS numbers: the American Registry for Internet Numbers (ARIN), Reseaux IP Europeens (RIPE), and Asia Pacific Network Information Center (APNIC). The numbers 64512 through 65535 are for private use and not to be advertised on the global Internet.

CLI

1. Router ID

```
set vrouter trust-vr router-id 10
```

2. BGP Routing Instance

```
set vrouter trust-vr protocol bgp 65000
set vrouter trust-vr protocol bgp enable
save
```

Example: Removing a BGP Routing Instance

In this example, you disable and remove the BGP routing instance in the trust-vr. BGP stops sessions with all peers.

WebUI

Network > Routing > Virtual Routers (trust-vr) > Edit > Edit BGP Instance: Deselect BGP Enabled, and then click **OK**.

Network > Routing > Virtual Routers (trust-vr) > Edit: Select **Delete BGP Instance**, and then click **OK** at the confirmation prompt.

CLI

```
unset vrouter trust-vr protocol bgp enable
unset vrouter trust-vr protocol bgp 65000
save
```

Enabling BGP on Interfaces

You must enable BGP on the interface on which the peer resides. (By default, interfaces on the NetScreen device are not bound to any routing protocol.)

Example: Enabling BGP on Interfaces

In this example, you enable BGP on the interface ethernet4.

WebUI

Network > Interfaces > Configure (for ethernet4): Select **Protocol BGP**, and then click **OK**.

CLI

```
set interface ethernet4 protocol bgp
save
```

Example: Disabling BGP on Interfaces

In this example, you disable BGP on the interface ethernet4. Note that any other interfaces on which you have enabled BGP are still able to transmit and process BGP packets.

WebUI

Network > Interfaces > Configure (for ethernet4): Clear **Protocol BGP**, and then click **OK**.

CLI

```
unset interface ethernet4 protocol bgp
save
```

Configuring a BGP Peer

Before two BGP devices can communicate and exchange routes, they need to identify each other so they can start a BGP session. You need to specify the IP addresses of the BGP peers and, optionally, configure parameters for establishing and maintaining the session. Peers can be either internal (IBGP) or external (EBGP) peers. For an EBGP peer, you need to specify the autonomous system in which the peer resides.

All BGP sessions are authenticated by checking the BGP peer identifier and the AS number advertised by the peers. A successful connection with a peer is logged. If anything goes wrong with the peer connection, a BGP notification message will either be sent to or received from the peer, which causes the connection to fail or close.

You can configure parameters for individual peer addresses. You can also assign peers to a *peer-group*, which then allows you configure parameters for the peer-group as a whole. Note that you cannot assign IBGP and EBGP peers to the same peer-group.

The following table describes parameters you can configure for BGP peers and the default values. An “X” in the Peer column indicates a parameter you can configure for an individual peer IP address while an “X” in the Peer Group column indicates a parameter you can configure for a peer-group.

BGP Parameter	Peer	Peer Group	Description	Default Value
Advertise default route	X		Advertises the default route in the virtual route to BGP peers.	Default route is not advertised
EBGP multihop	X	X	Number of nodes between local BGP and neighbor.	0 (disabled)
Force connect	X	X	Causes peer to drop existing BGP connection and accept a new connection.	N/A
Hold time	X	X	Time elapsed without a message from a peer before the peer is considered down.	180 seconds
Keepalive	X	X	Time between keepalive transmissions.	1/3 of hold-time
MD5 authentication	X	X	Configures MD-5 authentication.	Only peer identifier and AS number checked
MED	X		Configures MED attribute value.	0

BGP Parameter	Peer	Peer Group	Description	Default Value
Next-hop self	X	X	For routes sent to the peer, the next hop path attribute is set to the IP address of the interface of the local VR.	Next hop attribute unchanged
Reflector client	X	X	Peer is a reflector client when the local BGP is set as the route reflector.	None
Reject default route	X		Ignores default route advertisements from BGP peers.	Default routes from peers are added to routing table
Retry time	X	X	Time after a failed session attempt that the BGP session is reattempted.	120 seconds
Send community	X	X	Transmits community attribute to peer.	Community attribute not sent to peers
Weight	X	X	Priority of path between local BGP and peer.	100

You can configure some parameters at both the peer level and the protocol level (see “Optional BGP Configurations” on page 104). For example, you can configure the hold-time value for a specific peer at 210 seconds, while the default hold-time value at the protocol level is 180 seconds; the peer configuration takes precedence. You can set different MED values at the protocol level and at the peer level; the MED value you set at the peer level applies only to routes that are advertised to those peers.

Example: Configuring a BGP Peer

In the following example, you configure and enable a BGP peer. This peer has the following attributes:

- IP address 1.1.1.250
- Resides in AS 65500

Note: *You must enable each peer connection that you configure.*

WebUI

Network > Routing > Virtual Routers > Edit (for trust-vr) > Edit BGP Instance > Neighbors: Enter the following, and then click **Add**:

AS Number: 65500

Remote IP: 1.1.1.250

Network > Routing > Virtual Routers > Edit (for trust-vr) > Edit BGP Instance > Neighbors > Configure (for the peer you just added): Select **Peer Enabled**, and then click **OK**.

CLI

```
set vrouter trust-vr protocol bgp neighbor 1.1.1.250 remote-as 65500
set vrouter trust-vr protocol bgp neighbor 1.1.1.250 enable
save
```

Example: Configuring an IBGP Peer-Group

In the following example, you configure an IBGP peer group called **ibgp** that contains the following IP addresses: 10.1.2.250 and 10.1.3.250. Once you have defined a peer group, you can configure parameters (such as MD5 authentication) that apply to all members of the peer group.

Note: *You must enable each peer connection that you configure. If you configure peers as part of a peer group, you still need to enable the peer connections one by one*

WebUI

Network > Routing > Virtual Routers > Edit (for trust-vr) > Edit BGP Instance > Peer Group: Enter **ibgp** for Group Name, and then click **Add**.

> Configure (for ibgp): In the Peer authentication field, enter **verify03**, and then click **OK**.

Network > Routing > Virtual Routers > Edit (for trust-vr) > Edit BGP Instance > Neighbors: Enter the following, and then click **Add**:

AS Number: 65000

Remote IP: 10.1.2.250

Peer Group: ibgp (select)

Enter the following, and then click **Add**:

AS Number: 65000

Remote IP: 10.1.3.250

Peer Group: ibgp (select)

Network > Routing > Virtual Routers > Edit (for trust-vr) > Edit BGP Instance > Neighbors > Configure (for 10.1.2.250): Select **Peer Enabled**, and then click **OK**.

Network > Routing > Virtual Routers > Edit (for trust-vr) > Edit BGP Instance > Neighbors > Configure (for 10.1.3.250): Select **Peer Enabled**, and then click **OK**.

CLI

```
set vrouter trust-vr protocol bgp neighbor peer-group ibgp remote-as 65000
set vrouter trust-vr protocol bgp neighbor 10.1.2.250 peer-group ibgp
set vrouter trust-vr protocol bgp neighbor 10.1.3.250 peer-group ibgp
set vrouter trust-vr protocol bgp neighbor 10.1.2.250 enable
set vrouter trust-vr protocol bgp neighbor 10.1.3.250 enable
set vrouter trust-vr protocol bgp neighbor peer-group ibgp md5-authentication
    verify03
save
```

Verifying the BGP Configuration

You can review the configuration you entered through the WebUI or the CLI by executing the following CLI command:

```
ns-> get vrouter trust-vr protocol bgp config
set protocol bgp 65000
set enable
unset synchronization
set neighbor 1.1.1.250 remote-as 65500
set neighbor 1.1.1.250 enable
exit
```

You can verify that BGP is running on the VR by executing the following CLI command:

Verifies that BGP is running.

```
ns-> get vr trust-vr protocol bgp
Admin State:          enable
Local Router ID:      10.1.1.250
Local AS number:      65000
Hold time:            180
Keepalive interval:   60 = 1/3 hold time, default
Local MED is:         0
Always compare MED:   disable
Local preference:     100
Route Flap Damping:   disable
IGP synchronization:  disable
Route reflector:      disable
Cluster ID:           not set (ID = 0)
Confederation based on RFC 1965
Confederation:        disable (confederation ID = 0)
Member AS:            none
Origin default route: disable
Ignore default route: disable
```

We recommend that you always explicitly assign a router ID, rather than use the default. For information on setting a router ID, see Chapter 1, “Virtual Routers”.

You can verify that a BGP peer or peer group is enabled and see the state of the BGP session by executing the following CLI command:

```
ns-> get vrouter trust-vr protocol bgp neighbor
Peer AS Remote IP       Local IP          Wt ConnID Status   State     Flag
 65500 1.1.1.250       0.0.0.0          100      0 Enabled ACTIVE    0000

total 1 BGP peers shown
```

Indicates peer is enabled and session is active.

The state can be one of the following:

- Idle—the first state of the connection
- Connect—BGP is waiting for successful TCP transport connection
- Active—BGP is initiating a transport connection[3]
- OpenSent—BGP is waiting for an OPEN message from the peer
- OpenConfirm—BGP is waiting for a KEEPALIVE or NOTIFICATION message from the peer
- Established—BGP is exchanging UPDATE packets with the peer

3. A session state that continually changes between the Active and Connect may indicate a problem with the connection between the peers.

SECURITY CONFIGURATION

This section describes possible security problems in the BGP routing domain and methods of preventing attacks.

Note: *To make BGP more secure, you should configure all routers in the BGP domain to be at the same security level. Otherwise, a compromised BGP router can bring down the entire BGP routing domain.*

Authenticating Neighbors

A BGP router can be easily spoofed, since BGP packets are not encrypted and most protocol analyzers provide decapsulation of BGP packets. Authenticating BGP peers is the best way to fend off these types of attacks.

BGP provides MD5 authentication to validate BGP packets received from peer. MD5 authentication requires that the same key be used for both the sending and receiving BGP routers. All BGP packets received from the specified peer that are not authenticated are discarded. By default, only the peer identifier and AS number are checked for a BGP peer.

Example: Configuring MD5 Authentication for BGP Peers

In the following example, you first configure a BGP peer with the remote IP address 1.1.1.250 in AS 65500. You then configure the peer for MD5 authentication using the key 1234567890123456.

WebUI

Network > Routing > Virtual Routers > Edit (for trust-vr) > Edit BGP Instance > Neighbors: Enter the following, and then click **Add**:

AS Number: 65500

Remote IP: 1.1.1.250

> Configure (for Remote IP 1.1.1.250): Enter the following, and then click **OK**:

Peer Authentication: Enable (select)

MD5 password: 1234567890123456

Peer Enabled: (select)

CLI

```
set vrouter trust-vr
(trust-vr)-> set protocol bgp
(trust-vr/bgp)-> set neighbor 1.1.1.250 remote-as 65500
(trust-vr/bgp)-> set neighbor 1.1.1.250 md5-authentication 1234567890123456
(trust-vr/bgp)-> set neighbor 1.1.1.250 enable
(trust-vr/bgp)-> exit
(trust-vr)-> exit
save
```

Rejecting Default Routes

In a Route Detour Attack, a router injects a default route (0.0.0.0/0) into the routing domain in order to detour packets to itself. The router can then either drop the packets, causing service disruption, or it can obtain sensitive information in the packets before forwarding them. On NetScreen devices, BGP by default accepts any default routes that are sent from BGP peers and adds the default route to the routing table.

Example: Rejecting Default Routes

In this example, you configure the BGP routing instance running in the trust-vr virtual router to ignore any default routes that are sent from BGP peers.

WebUI

Network > Routing > Virtual Router (trust-vr) > Edit > Edit BGP Instance: Enter the following, and then click **OK**:

Ignore default route from peer: (select)

CLI

```
set vrouter trust-vr protocol bgp reject-default-route
save
```

OPTIONAL BGP CONFIGURATIONS

This section describes the parameters you can configure for the BGP routing protocol in the virtual router. You can configure these parameters with either the CLI BGP context commands or the WebUI. This section explains some of the more complex parameter configurations.

The following table describes BGP parameters and their default values.

BGP Protocol Parameter	Description	Default Value
Advertise default route	Advertises the default route in the virtual router to BGP peers.	Default route not advertised
Aggregate	Create aggregated address.	Disabled
Always compare MED	Compares MED values in routes.	Disabled
AS path access list	Create an AS path access list to permit or deny routes.	
Community list	Create community lists. See “BGP Communities” on page 113.	
AS confederation	Create confederations. See “Confederations” on page 110.	
Flap damping	Block advertisement of a route until it becomes stable.	Disabled
Hold time	Time elapsed without a message from a peer before the peer is considered down.	180 seconds
Keepalive	Time between keepalive transmissions.	1/3 of hold-time
Local preference	Configures LOCAL_PREF metric	100
MED	Configures MED attribute value.	0
Network	Create network and subnetwork entries.	
Route redistribution	Import routes into BGP from other routing protocols. See “Redistributing Routes” on page 105.	
Reflector	Configures the local BGP instance as a route reflector to clients. See “Route Reflection” on page 107.	Disabled
Reject default route	Ignores default route advertisements from BGP peers.	Default routes from peers are added to routing table

BGP Protocol Parameter	Description	Default Value
Retry time	Time after an unsuccessful BGP session establishment with a peer that session establishment is retried.	120 seconds
Synchronization	Enables synchronization with an IGP, such as OSPF or RIP.	Disabled

Redistributing Routes

Route redistribution is the exchange of route information between routing protocols. For example, you can redistribute the following types of routes into the BGP routing instance in the same VR:

- Routes learned from OSPF or RIP
- Directly connected routes
- Imported routes
- Statically configured routes

When you configure route redistribution, you must first specify a route map to filter the routes that are redistributed. For more information about creating route maps for route redistribution, refer to Chapter 1, “Virtual Routers”.

Example: Redistributing an OSPF Route into BGP

In the following example, you redistribute a route that originated from an OSPF routing domain into the current BGP routing domain. Both the CLI and WebUI examples assume that you previously created a route map called add-ospf.

WebUI

Network > Routing > Virtual Routers > Edit (for trust-vr) > Edit BGP Instance > Redist. Rules: Enter the following, and then click **Add**:

Route Map: add-ospf

Protocol: OSPF

CLI

```
set vrouter trust-vr protocol bgp redistribute route-map add-ospf protocol ospf
save
```

AS-Path Access List

The AS-path attribute contains a list of the ASs through which a route has traversed. BGP prepends the local AS number to the AS-path attribute when a route passes through the AS. You can use an *AS-path access list* to filter routes based on the AS-path information. An AS-path access list consists of a set of regular expressions that define AS-path information and whether the routes that match the information are permitted or denied. For example, you can use an AS-path access list to filter routes that have passed through a particular AS or routes that originated in a particular AS.

Regular expressions are a way to define a search for specific pattern in the AS-path attribute. You can use special symbols and characters in constructing a regular expression. For example, to match routes that have passed through AS 65000, use the regular expression _**65000**_ (the underscores match any characters before or after 65000). You can use the regular expression “65000$” to match routes that originated in AS 65000 (the dollar sign matches the end of the AS-path attribute, which would be the AS where the route originated).

Example: Configuring an AS-Path Access List

The following example configures an AS-path access list for the trust-vr virtual router that allows routes that have passed through AS 65000 but does not allow routes that originated in AS 65000.

WebUI

Network > Routing > Virtual Routers > Edit (for trust-vr) > Edit BGP Instance > AS Path: Enter the following, and then click **Add**:

AS Path Access List ID: 2

Deny: (select)

AS Path String: 65000$

Network > Routing > Virtual Routers > Edit (for trust-vr) > Edit BGP Instance > AS Path: Enter the following, and then click **Add**:

AS Path Access List ID: 2

Permit: (select)

AS Path String: _65000_

CLI

```
set vrouter trust-vr protocol bgp as-path-access-list 2 deny 65000$
set vrouter trust-vr protocol bgp as-path-access-list 2 permit _65000_
save
```

Route Reflection

Because an IBGP router cannot readvertise routes learned from one IBGP peer to another IBGP peer (see "External and Internal BGP" on page 90), a *full mesh* of IBGP sessions is required where each router in a BGP AS is a peer to every other router in the AS. Note that having a full mesh does not mean each pair of routers needs to be directly connected, but each router needs to be able to establish and maintain an IBGP session with every other router. For example, in an AS with 8 routers, each of the 8 routers would need to peer with the 7 other routers.

A full mesh configuration of IBGP sessions does not scale very well. The number of full mesh IBGP sessions required for an AS is calculated by the following formula:

(x * (x-1))/2

For an AS containing 8 routers, the number of full mesh IBGP sessions would be 28. If there were 20 routers in the AS, the number of full mesh IBGP sessions needed in the network would be 190.

Route reflection is a method for solving the IBGP scalability problem and is described in RFC 1966. A *route reflector* is a router that passes IBGP learned routes to specified IBGP neighbors (*clients*), thus eliminating the need for full mesh sessions. The route reflector and its clients make up a *cluster*, which you can further identify with a cluster ID. Routers outside of the cluster treat the entire cluster as a single entity, instead of interfacing with each individual router in full mesh. This arrangement greatly reduces overhead. The clients exchange routes with the route reflector, while the route reflector reflects routes between clients.

On the NetScreen device, you can specify the local virtual router to act as a route reflector. You can optionally specify a cluster ID for the route reflector. If you specify a cluster ID, the BGP routing instance appends the cluster ID to the Cluster-List attribute of a route. The cluster ID helps prevent routing loops as the local BGP routing instance drops a route when its cluster ID appears in the route's cluster list.

Note: *Before you can configure a cluster ID, the BGP routing instance must be disabled.*

After you set up a route reflector on the local virtual router, you then define the route reflector's clients. You can specify individual IP addresses or a peer-group for the clients. You do not need to configure anything on the clients.

Example: Configuring the Virtual Router as a Route Reflector

In the following example, the EBGP router advertises the 5.5.5.0/24 prefix to Client 3. Without route reflection, Client 3 advertises the route to NetScreen-A, but NetScreen-A does not readvertise that route to Clients 1 and 2. If you configure NetScreen-A as the route reflector with Clients 1, 2, and 3 as its clients, NetScreen-A readvertises routes received from Client 3 to Clients 1 and 2.

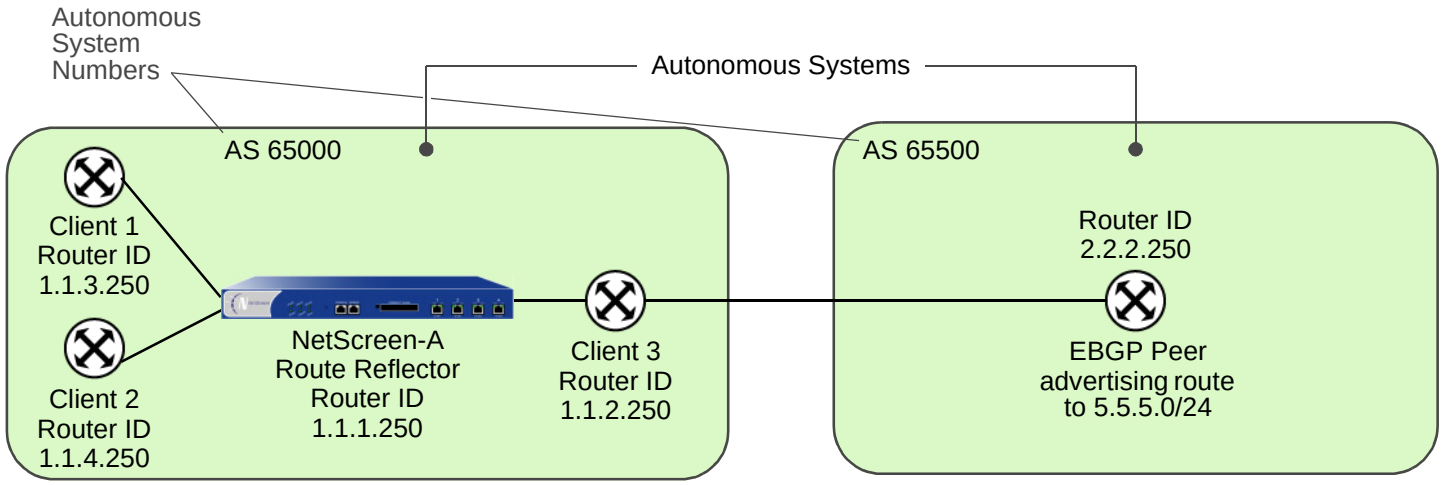

WebUI

Network > Routing > Virtual Routers > Edit (for trust-vr) > Edit BGP Instance: Enter the following, and then click **Apply**:

Route reflector: Enable

Cluster ID: 99

> Neighbors: Enter the following, and then click **Add**:

AS Number: 65000

Remote IP: 1.1.2.250

Enter the following, and then click **Add**:

AS Number: 65000

Remote IP: 1.1.3.250

Enter the following, and then click **Add**:

AS Number: 65000

Remote IP: 1.1.4.250

> Configure (for Remote IP 1.1.2.250): Select **Reflector Client**, and then click **OK**.

> Configure (for Remote IP 1.1.3.250): Select **Reflector Client**, and then click **OK**.

> Configure (for Remote IP 1.1.4.250): Select **Reflector Client**, and then click **OK**.

CLI

```
set vrouter trust-vr protocol bgp reflector
set vrouter trust-vr protocol bgp reflector cluster-id 99
set vrouter trust-vr protocol bgp neighbor 1.1.2.250 reflector-client
set vrouter trust-vr protocol bgp neighbor 1.1.3.250 reflector-client
set vrouter trust-vr protocol bgp neighbor 1.1.4.250 reflector-client
save
```

Confederations

Like route reflection (see “Route Reflection” on page 107), *confederations* are another approach to solving the problem of full mesh scaling in an IBGP environment and are described in RFC 1965. A confederation splits an autonomous system into several smaller ASs, with each sub-AS a fully-meshed IBGP network. A router outside the confederation sees the entire confederation as a single autonomous system with a single identifier; the sub-AS networks are not visible outside the confederation. Sessions between routers in two different sub-ASs in the same confederation, known as EIBGP sessions, are essentially EBGP sessions between autonomous systems, but the routers also exchange routing information as if they were IBGP peers.

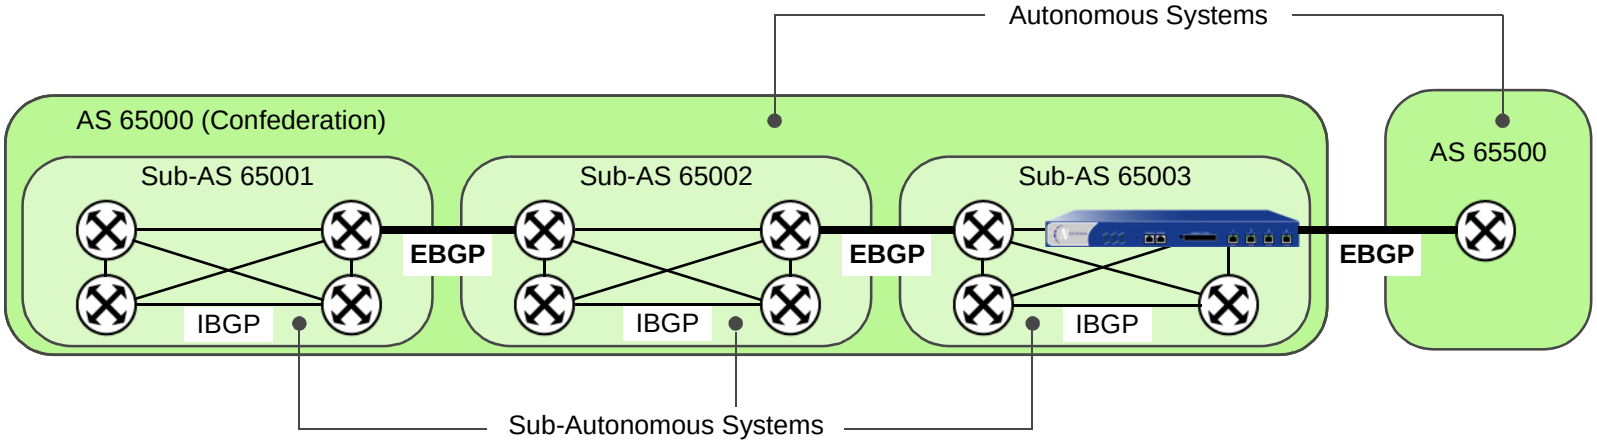

For each router in a confederation, you need to specify the following:

- The sub-AS number (this is the AS number that you specify when you create the BGP routing instance)
- The confederation to which the sub-AS belongs (this is the AS number that is visible to BGP routers outside the confederation)
- The peer sub-AS numbers in the confederation
- Whether the confederation supports RFC 1965 (the default) or RFC 3065[4]

4. The AS-Path attribute (see “Path Attributes” on page 89) is normally composed of a sequence. ASs traversed by the routing update. RFC 3065 allows for the AS-Path attribute to include the member ASs in the local confederation traversed by the routing update.

Example: Configuring a Confederation

In this example, the NetScreen device is a BGP router in sub-AS 65003 that belongs to the confederation 65000. The peer sub-ASs in confederation 65000 are 65002 and 65003.

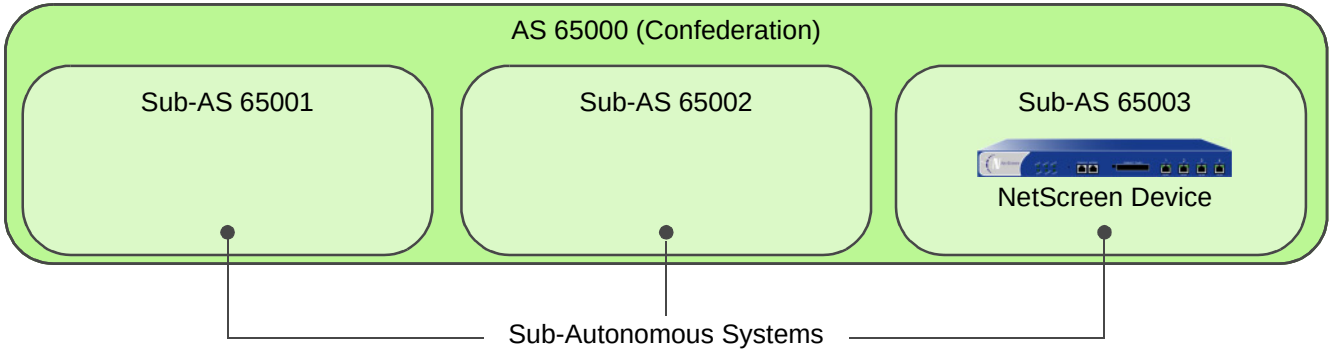

WebUI

Network > Routing > Virtual Routers > Edit (for trust-vr) > Create BGP Instance: Enter the following, and then click **Apply**:

AS Number (required): 65003
BGP Enabled: (select)

> Confederation: Enter the following, and then click **Apply**:

Enable: (select)
ID: 65000
Supported RFC: RFC 1965 (select)

Enter the following, and then click **Add**:
Peer member area ID: 65001

Enter the following, and then click **Add**:
Peer member area ID: 65002

CLI

```
set vrouter trust-vr protocol bgp 65003
set vrouter trust-vr protocol bgp confederation id 65000
set vrouter trust-vr protocol bgp confederation peer 65001
set vrouter trust-vr protocol bgp confederation peer 65002
save
```

BGP Communities

The communities path attribute provides a way of grouping destinations (called communities), which a BGP router can then use to control the routes it accepts, prefers, or redistributes to peers. A BGP router can either append communities to a route (if the route does not have a communities path attribute) or modify the communities in a route (if the route contains a communities path attribute). The communities path attribute provides an alternative to distributing route information based on IP address prefixes or AS path attribute. You can use the communities path attribute in many ways, but its primary purpose is to simplify configuration of routing policies in complex networking environments.

RFC 1997 describes the operation of BGP communities. An AS administrator can assign the same community to a set of routes that require the same routing decisions; this is sometimes called *route coloring*. For example, you can assign one community value to routes that receive access to the Internet and a different community value to routes that do not.

There are two forms of communities:

- A *specific community* consists of the AS identifier and a community identifier. The community identifier is defined by the AS administrator.
- A *well-known community* signifies special handling for routes that contain these community values. The following are well-known community values that you can specify for BGP routes on the NetScreen device:
  - **no-export**: Routes with this communities path attribute are not advertised outside a BGP confederation.
  - **no-advertise**: Routes with this communities path attribute are not advertised to other BGP peers.
  - **no-export-subconfed**: Routes with this communities path attribute are not advertised to EBGP peers.

You can use a route map to filter routes that match a specified community list, remove or set the communities path attributes in routes, or add or delete communities from the route.

For example, if an ISP provides Internet connectivity to its customers, then all routes from those customers can be assigned a specific community number. Those customer routes are then advertised to peer ISPs. Routes from other ISPs are assigned different community numbers and are not advertised to peer ISPs.

Index

A

access list for routes 24

B

BGP
- AS-path access list 106
- authenticating neighbors 102
- communities 113
- confederations 110
- configuration steps 91
- configuring peer group 95
- configuring peers 95
- creating instance in VR 92
- enabling in VR 92
- enabling on interface 94
- external BGP 90
- internal BGP 90
- message types 89
- parameters 104
- path attributes 89
- protocol overview 88
- redistributing routes 105
- regular expressions 106
- rejecting default routes 103
- route reflection 107
- security configuration 102
- verifying configuration 100

C

character types, ScreenOS supported x
CLI
- conventions vi

conventions
- CLI vi
- illustration ix
- names x
- WebUI vii

E

exporting routes 28

I

illustration
- conventions ix

importing routes 28

N

names
- conventions x

O

OSPF
- areas 34
- assigning interface to area 42
- authenticating neighbors 60
- backup designated router 36
- broadcast network 36
- configuration steps 38
- creating instance in VR 39
- defining area 41
- designated router 36
- enabling on interface 44
- filtering neighbors 62
- global parameters 51
- hello protocol 35
- interface parameters 57
- link-state advertisements 34, 37
- not so stubby area 35
- point-to-point network 36
- protecting against flooding 64
- redistributing routes 49
- rejecting default routes 63
- router adjacency 35
- router types 35
- security configuration 60
- stub area 35
- summarizing redistributed routes 50
- virtual links 53

R

RIP
- authenticating neighbors 80
- configuration steps 69
- creating instance in VR 70
- enabling on interface 72
- filtering neighbors 82
- global parameters 76
- interface parameters 78
- protecting against flooding 84
- protocol overview 68
- redistributing routes 73
- rejecting default routes 83
- security configuration 80

route filtering 24
route map 22
route metric 17
route redistribution 21
routing
- route preference 15
- route selection 15

routing table
- route selection 15

S

source-based routing 17

V

VRs 3–28
- access lists 24
- BGP 91–101
- custom 7
- exporting routes 28
- forwarding traffic between 4
- importing routes 28
- maximum routing table entries 14

modifying 12
on vsys 9
OSPF 38–65
predefined 3
RIP 69–86
route filtering 24
route map 22
route metric 17
route preference 15
route redistribution 21
route selection 15
router ID 12
source-based routing 17
using two VRs 3, 4

W

WebUI
    conventions vii

NetScreen Concepts & Examples

ScreenOS Reference Guide

Volume 7: Virtual Systems

ScreenOS 5.0.0

P/N 093-0930-000

Rev. B

Copyright Notice

NetScreen Technologies, Inc.
Building #3
805 11th Avenue
Sunnyvale, CA 94089
www.netscreen.com

FCC Statement

The following information is for FCC compliance of Class A devices: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. The equipment generates, uses, and can radiate radio-frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case users will be required to correct the interference at their own expense.

The following information is for FCC compliance of Class B devices: The equipment described in this manual generates and may radiate radio-frequency energy. If it is not installed in accordance with NetScreen's installation instructions, it may cause interference with radio and television reception. This equipment has been tested and found to comply with the limits for a Class B digital device in accordance with the specifications in part 15 of the FCC rules. These specifications are designed to provide reasonable protection against such interference in a residential installation. However, there is no guarantee that interference will not occur in a particular installation.

If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Consult the dealer or an experienced radio/TV technician for help.
- Connect the equipment to an outlet on a circuit different from that to which the receiver is connected.

Caution: Changes or modifications to this product could void the user's warranty and authority to operate this device.

Disclaimer

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR NETSCREEN REPRESENTATIVE FOR A COPY.

Contents

Preface iii
Conventions iv
CLI Conventions iv
WebUI Conventions v
Illustration Conventions vii
Naming Conventions and Character Types viii
NetScreen Documentation ix
Chapter 1 Virtual Systems 1
Creating a Vsys Object 3
Example: Vsys Objects and Admins 3
Virtual Routers 6
Zones 7
Interfaces 8
Traffic Sorting 10
Traffic Destined for the NetScreen Device 10
Through Traffic 11
Dedicated and Shared Interfaces 15
Dedicated Interfaces 15
Shared Interfaces 15
Importing and Exporting Physical Interfaces 18
Example: Importing a Physical Interface to a Virtual System 18
Example: Exporting a Physical Interface from a Virtual System 19
VLAN-Based Traffic Classification 21
VLANs 22
Defining Subinterfaces and VLAN Tags 23
Example: Defining Three Subinterfaces and VLAN Tags 25
Communicating between Virtual Systems 28
Example: InterVsys Communication 28
IP-Based Traffic Classification 33
Example: Configuring IP-Based Traffic Classification 35
Logging On as a Vsys Admin 38
Example: Logging On and Changing Your Password 38
Index ..... IX-I

Preface

You can logically partition a single NetScreen security system into multiple virtual systems to provide multi-tenant services. Each virtual system (vsys) is a unique security domain and can be managed by its own administrators (called “virtual system administrators” or “vsys admins”) who can individualize their security domain by setting their own address books, user lists, custom services, VPNs, and policies.

Volume 7, “Virtual Systems” describes virtual systems, dedicated and shared interfaces, and VLAN-based and IP-based traffic classification. This volume also describes how to create a vsys (you must have root-level administrator privilege) and define vsys admins.

CONVENTIONS

This document contains several types of conventions, which are introduced in the following sections:

- "CLI Conventions"
- "WebUI Conventions" on page v
- "Illustration Conventions" on page vii
- "Naming Conventions and Character Types" on page viii

CLI Conventions

The following conventions are used when presenting the syntax of a command line interface (CLI) command:

- Anything inside square brackets [ ] is optional.
- Anything inside braces { } is required.
- If there is more than one choice, each choice is separated by a pipe (|). For example,

  ```
  set interface { ethernet1 | ethernet2 | ethernet3 } manage
  ```

  means "set the management options for the ethernet1, ethernet2, or ethernet3 interface".
- Variables appear in *italic.* For example:

  ```
  set admin user name password
  ```

When a CLI command appears within the context of a sentence, it is in **bold** (except for variables, which are always in *italic*). For example: "Use the **get system** command to display the serial number of a NetScreen device."

Note: *When typing a keyword, you only have to type enough letters to identify the word uniquely. For example, typing* ***set adm u joe j12fmt54*** *is enough to enter the command* ***set admin user joe j12fmt54****. Although you can use this shortcut when entering commands, all the commands documented here are presented in their entirety.*

WebUI Conventions

Throughout this book, a chevron (>) is used to indicate navigation through the WebUI by clicking menu options and links. For example, the path to the address configuration dialog box is presented as **Objects > Addresses > List > New**. This navigational sequence is shown below.

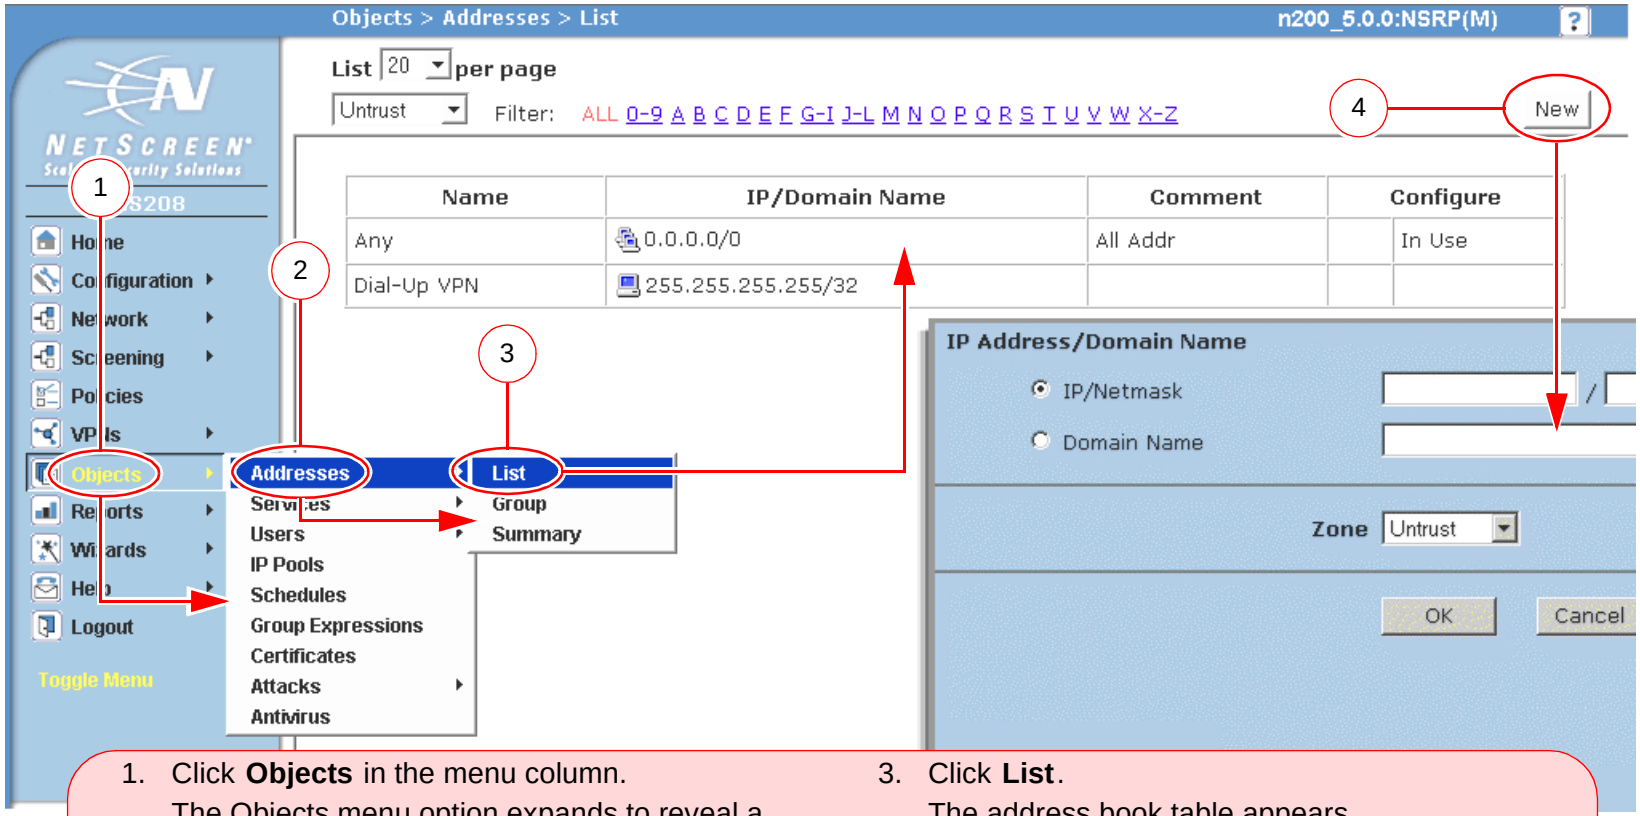

1. Click **Objects** in the menu column.
   The Objects menu option expands to reveal a subset of options for Objects.
2. (Applet menu) Hover the mouse over **Addresses**.
   (DHTML menu) Click **Addresses**.
   The Addresses option expands to reveal a subset of options for Addresses.
3. Click **List**.
   The address book table appears.
4. Click the **New** link.
   The new address configuration dialog box appears.

To perform a task with the WebUI, you must first navigate to the appropriate dialog box where you can then define objects and set parameters. The set of instructions for each task is divided into two parts: a navigational path and configuration details. For example, the following set of instructions includes the path to the address configuration dialog box and the settings for you to configure:

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: addr_1

IP Address/Domain Name:

IP/Netmask: (select), 10.2.2.5/32

Zone: Untrust

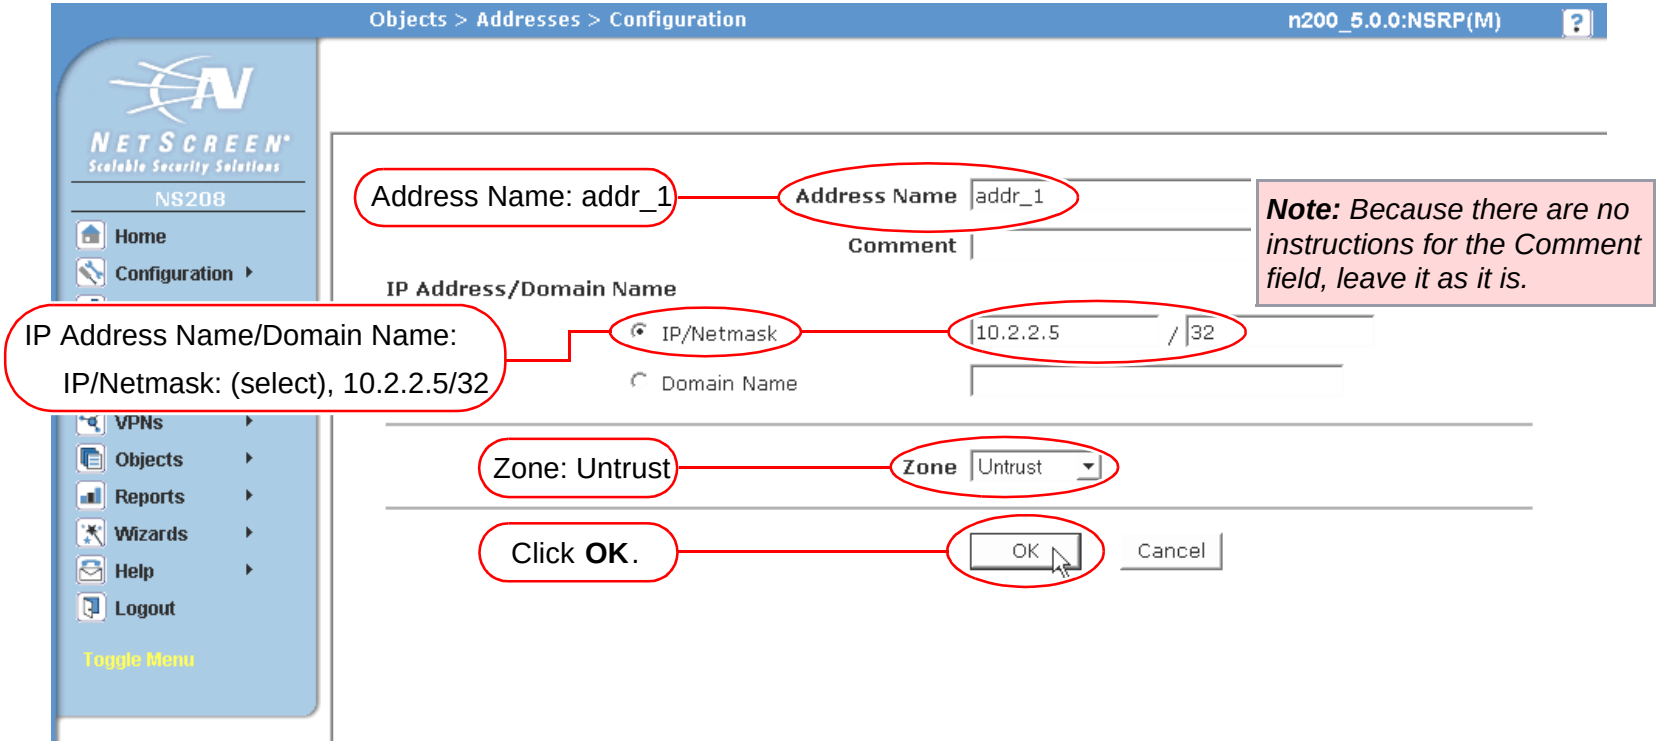

Illustration Conventions

The following graphics make up the basic set of images used in illustrations throughout this book:

Generic NetScreen Device

Virtual Routing Domain

Security Zone

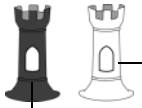

Security Zone Interfaces

White = Protected Zone Interface (example: Trust Zone)

Black = Outside Zone Interface (example: Untrust Zone)

Tunnel Interface

VPN Tunnel

Router Icon

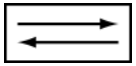

Switch Icon

Local Area Network (LAN) with a Single Subnet

(example: 10.1.1.0/24)

Internet

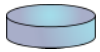

Dynamic IP (DIP) Pool

Desktop Computer

Laptop Computer

Generic Network Device

(examples: NAT server, Access Concentrator)

Server

Naming Conventions and Character Types

ScreenOS employs the following conventions regarding the names of objects—such as addresses, admin users, auth servers, IKE gateways, virtual systems, VPN tunnels, and zones—defined in ScreenOS configurations.

- If a name string includes one or more spaces, the entire string must be enclosed within double quotes ("); for example, **set address trust "local LAN" 10.1.1.0/24**.
- NetScreen trims any spaces leading or trailing text within a set of double quotes; for example, **" local LAN "** becomes **"local LAN"**.
- NetScreen treats multiple consecutive spaces as a single space.
- Name strings are case sensitive, although many CLI key words are case insensitive. For example, **"local LAN"** is different from **"local lan"**.

ScreenOS supports the following character types:

- Single-byte character sets (SBCS) and multiple-byte character sets (MBCS). Examples of SBCS are ASCII, European, and Hebrew. Examples of MBCS—also referred to as double-byte character sets (DBCS)—are Chinese, Korean, and Japanese.

  Note: *A console connection only supports SBCS. The WebUI supports both SBCS and MBCS, depending on the character sets that your Web browser supports.*

- ASCII characters from 32 (0x20 in hexidecimals) to 255 (0xff), except double quotes ("), which have special significance as an indicator of the beginning or end of a name string that includes spaces.

NETSCREEN DOCUMENTATION

To obtain technical documentation for any NetScreen product, visit www.netscreen.com/resources/manuals/.

To obtain the latest software version, visit www.netscreen.com/services/download_soft. Select a category of software product from the dropdown list, then follow the displayed instructions. (You must be a registered user to download NetScreen software.)

If you find any errors or omissions in the following content, please contact us at the e-mail address below:

techpubs@netscreen.com

1

Virtual Systems

You can logically partition a single NetScreen security system[1] into multiple virtual systems to provide multi-tenant services. Each virtual system (vsys) is a unique security domain and can have its own administrators (called “virtual system administrators” or “vsys admins”) who can individualize their security domain by setting their own address books, user lists, custom services, VPNs, and policies (although only a root-level administrator can set firewall security options, create virtual system administrators, and define interfaces and subinterfaces).

Note: *For more information on the various levels of administration that NetScreen supports, see “Levels of Administration” on page* ***3****-37.*

NetScreen virtual systems support two kinds of traffic classifications: VLAN-based and IP-based, both of which can function exclusively or concurrently. This chapter discusses the following concepts and implementation of virtual systems:

- “Creating a Vsys Object” on page 3
  - “Virtual Routers” on page 6
  - “Zones” on page 7
  - “Interfaces” on page 8
- “Traffic Sorting” on page 10
  - “Traffic Destined for the NetScreen Device” on page 10
  - “Through Traffic” on page 11
  - “Dedicated and Shared Interfaces” on page 15
  - “Importing and Exporting Physical Interfaces” on page 18

1. NetScreen devices are divided into two general categories: security systems and appliances. Only NetScreen security systems can support virtual systems. Refer to the NetScreen marketing literature to see which platforms support this feature.

- "VLAN-Based Traffic Classification" on page 21
  - "VLANs" on page 22
  - "Defining Subinterfaces and VLAN Tags" on page 23
  - "Communicating between Virtual Systems" on page 28
- "IP-Based Traffic Classification" on page 33
- "Logging On as a Vsys Admin" on page 38

CREATING A VSYS OBJECT

The root administrator or root-level read/write admin must complete the following tasks to create a vsys object:

- Define a virtual system
- (Optional) Define one or more vsys admins[2]
- Select the virtual router that you want the vsys to use for its Trust-*vsysname* zone, Untrust-Tun-*vsysname* zone, and Global-*vsysname* zone

After creating a vsys object, you—as the root-level admin—need to perform other configurations to make it a functional vsys. You must configure subinterfaces or interfaces for the vsys, and possibly shared virtual routers and shared security zones. The subsequent configurations depend on whether the vsys is intended to support VLAN-based or IP-based traffic classifications, or a combination of both. After completing these configurations, you can then exit the virtual system and allow a vsys admin, if defined, to log on and begin configuring addresses, users, services, VPNs, routes, and policies.

Example: Vsys Objects and Admins

In this example, as a root-level admin, you create three vsys objects: vsys1, vsys2, vsys3. For vsys1, you create vsys admin Alice with password wIEaS1v1[3]. For vsys2, you create vsys admin Bob with password pjF56Ms2. For vsys3, you do not define a vsys admin. Instead, you accept the admin definition that the NetScreen device automatically generates. In the case of vsys3, the NetScreen device creates the admin “vsys_vsys3” with password “vsys_vsys3”.

Note: *Vsys names, admin names, and passwords are case-sensitive. “Vsys abc” is different from “vsys ABC.”*

For vsys1 and vsys2, you use the default virtual router. For vsys3, you choose the sharable root-level untrust-vr.

2. A root-level administrator can define one vsys admin with read-write privileges and one vsys admin with read-only privileges per vsys.
3. Only a root-level administrator can create a vsys admin’s profile (user name and password). Because the NetScreen device uses the user name to determine the vsys to which a user belongs, a vsys admin cannot change his or her user name. However, a vsys admin can (and should) change his or her password.

After you create a vsys through the WebUI, you remain at the root level. Entering the newly created vsys requires a separate step:

Vsys: Click **Enter** (for the virtual system you want to enter).

The WebUI pages of the vsys you have entered appear, with the name of the vsys above the central display area—Vsys:*Name*.

When you create a vsys through the CLI, you immediately enter the system that you have just created. (To enter an existing vsys from the root level, use the **enter vsys** *name_str* command.) When you enter a vsys, note that the CLI command prompt changes to include the name of the system in which you are now issuing commands.

WebUI

1. Vsys1

   Vsys > New: Enter the following, and then click **OK**:

   Vsys Name: vsys1

   Vsys Admin Name: Alice

   Vsys Admin New Password: wIEaS1v1

   Confirm New Password: wIEaS1v1

   Virtual Router:

   Create a default virtual router: (select)

2. Vsys2

   Vsys > New: Enter the following, and then click **OK**:

   Vsys Name: vsys2

   Vsys Admin Name: Bob

   Vsys Admin New Password: pjF56Ms2

   Confirm New Password: pjF56Ms2

   Virtual Router:

   Create a default virtual router: (select)

3. Vsys3

Vsys > New: Enter the following, and then click **OK**:

Vsys Name: vsys3

Virtual Router:

Select an existing virtual router: (select), untrust-vr

CLI

1. Vsys1

```
ns-> set vsys vsys1
ns(vsys1)-> set admin name Alice
ns(vsys1)-> set admin password wIEaS1v1
ns(vsys1)-> save[4]
ns(vsys1)-> exit
```

2. Vsys2

```
ns-> set vsys vsys2
ns(vsys2)-> set admin name Bob
ns(vsys2)-> set admin password pjF56Ms2
ns(vsys2)-> save
ns(vsys2)-> exit
```

3. Vsys3

```
ns-> set vsys vsys3 vrouter share untrust-vr
ns(vsys3)-> save
```

4. After issuing any commands, you must issue a **save** command before issuing an **exit** command or the NetScreen device loses your changes

Virtual Routers

When a root-level admin creates a vsys object, the vsys automatically has the following virtual routers available for its use:

- All shared root-level virtual routers, such as the untrust-vr

  In the same way that a vsys and the root system share the Untrust zone, they also share the untrust-vr, and any other virtual routers defined at the root level as sharable.

- Its own virtual router

  By default, a vsys-level virtual router is named *vsysname-vr*. You can also customize its name to make it more meaningful. This is a vsys-specific virtual router that, by default, maintains the routing table for the Trust-*vsysname* zone. All vsys-level virtual routers are non-sharable.

You can select any shared virtual router or the vsys-level virtual router as the default virtual router for a vsys. To change the default virtual router, enter a vsys and use the following CLI command: **set vrouter** *name* **default-vrouter**.

If you, as a root-level administrator, want all of the vsys zones to be in the untrust-vr routing domain—for example, if all the interfaces bound to the Trust -*vsysname* zone are in Route mode—you can dispense with the vsysname-vr by changing the vsys-level security zone bindings from the *vsysname*-vr to the untrust-vr. For more information on virtual routers, see “Virtual Routers” on page **6**-1.

Note: *This release of ScreenOS supports user-defined virtual routers within a virtual system.*

Zones

Each virtual system (vsys) is a unique security domain and can share security zones with the root system and have its own security zones. When a root-level admin creates a vsys object, the following zones are automatically inherited or created:

- All shared zones (inherited from the root system)
- Shared Null zone (inherited from the root system)
- Trust-*vsys_name* zone
- Untrust-Tun-*vsys_name* zone
- Global-*vsys_name* zone

> ***Note:*** *For information on each of these zone types, see "Zones" on page* ***2****-45.*

Each vsys can also support extra user-defined security zones. You can bind these zones to any shared virtual routers defined at the root level or to the virtual router dedicated to that vsys. To create a security zone for a vsys named vsys1, do either of the following:

WebUI

Vsys > Enter (for vsys1)

Network > Zones > New: Enter the following, and then click **OK**:

Zone Name: (type a name for the zone)

Virtual Router Name: (select a virtual router from the drop-down list)

Zone Type: Layer 3

CLI

```
ns-> enter vsys vsys1
ns (vsys1)-> set zone name name_str
ns(vsys1)-> set zone vrouter vrouter
ns(vsys1)-> save
```

The maximum number of security zones that a vsys or the root system can contain is limited only by the number of security zones available at the device level[5]. It is possible for a single vsys to consume all available security zones if the root admin or a root-level read/write admin assigns all of them to that particular vsys. Conversely, if all virtual systems share root-level security zones and do not make use of any user-defined vsys-level zones, then all security zones are available for root-level use.

Interfaces

A vsys can support the following three kinds of interfaces for their Untrust and Trust zones:

Untrust Zone Interface Types	Trust Zone Interface Types
• Dedicated Physical Interface	• Dedicated Physical Interface
• Subinterface (with VLAN tagging as a means for trunking* inbound and outbound traffic)	• Subinterface (with VLAN tagging)
• Shared Interface (physical, subinterface, redundant interface, aggregate interface) with Root System	• Shared Physical Interface with Root System (and IP-based traffic classification†)

* For information about VLAN tagging and trunking concepts, see "VLANs" on page 22.

† For more information about IP-based traffic classification, see "IP-Based Traffic Classification" on page 33.

You can bind one, two, or all three of the above interface types to a security zone concurrently. You can also bind multiple interfaces of each type to a zone.

5. The total number of user-definable (or "custom") security zones available at the device level is the sum of the number of root-level custom zones—as defined by one or more zone license keys—and the number of custom zones permitted by the vsys license key.

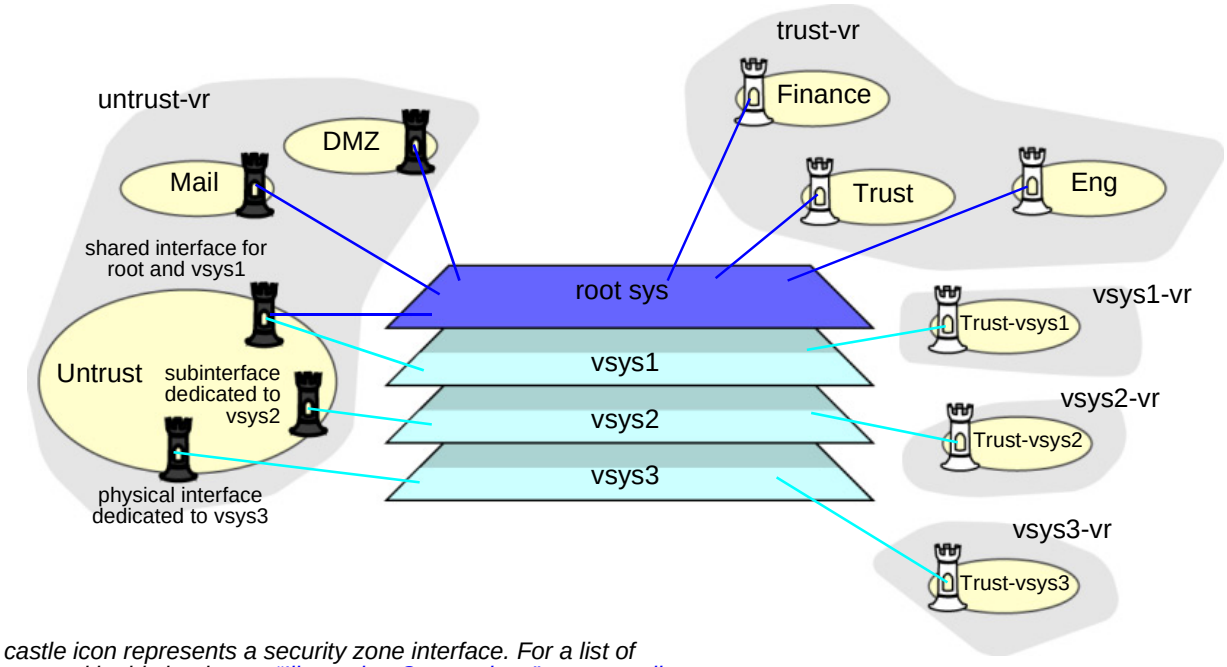

Note: *The castle icon represents a security zone interface. For a list of graphic icons used in this book, see "Illustration Conventions" on page vii.*

TRAFFIC SORTING

The NetScreen device must sort every packet it receives for delivery to the proper system. A NetScreen device receives two kinds of user traffic, which it sorts in two different ways:

- Traffic destined for an IP address on the device itself, such as encrypted VPN traffic and traffic destined for a MIP or VIP
- Traffic destined for an IP address beyond the device

Traffic Destined for the NetScreen Device

For traffic destined for an object (VPN, MIP, or VIP) on the NetScreen device, the device determines the system to which the traffic belongs through the association of the object with the system in which it was configured.

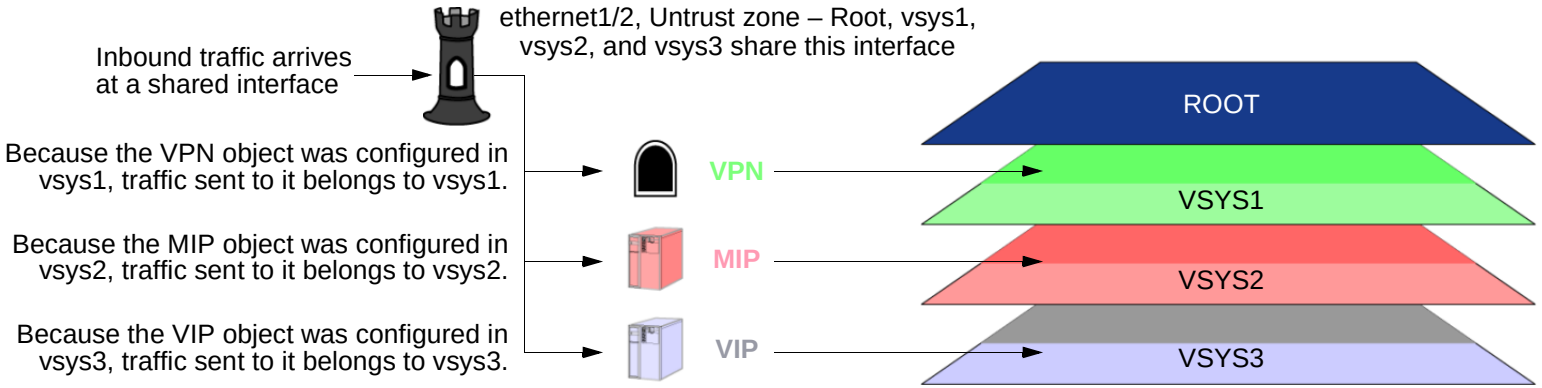

Inbound traffic can also reach a vsys via VPN tunnels; however, if the outgoing interface is a shared interface, you cannot create an AutoKey IKE VPN tunnel for a vsys and the root system to the same remote site.

Through Traffic

For traffic destined for an IP address beyond the NetScreen device (also known as “through traffic”), the device uses techniques made possible by VLAN-based and IP-based traffic classifications. VLAN-based traffic classification uses VLAN tags[6] in frame headers to identify the system to which inbound traffic belongs. IP-based traffic classification uses the source and destination IP address in IP packet headers to identify the system to which traffic belongs. The procedure that the NetScreen device uses to determine the system to which a packet belongs progresses through the following three steps:

1. Ingress Interface/Source IP Traffic Classification

The NetScreen device checks if the ingress interface is a dedicated interface or a shared interface[7].

1. If the ingress interface is dedicated to a vsys (“v-i”, for example), the NetScreen device associates the traffic with the system to which the interface is dedicated.
2. If the ingress interface is a shared interface, the NetScreen device uses IP classification to check if the source IP address is associated with a particular vsys.
   - If the source IP address is not associated with a particular vsys, ingress IP classification fails.
   - If the source IP address is associated with a particular vsys, ingress IP classification succeeds.

2. Egress Interface/Destination IP Traffic Classification

The NetScreen device checks if the egress interface is shared or dedicated.

1. If the egress interface is dedicated to a vsys (“v-e”, for example), the NetScreen device associates the traffic with the system to which the interface is dedicated.
2. If the egress interface is a shared interface, the NetScreen device uses IP classification to check if the destination IP address is associated with a particular vsys.
   - If the destination IP address is not associated with a particular vsys, egress IP classification fails.
   - If the destination IP address is associated with a particular vsys, egress IP classification succeeds.

6. VLAN tagging requires the use of subinterfaces. A subinterface must be dedicated to a system, in contrast to a shared interface, which is shared by all systems.
7. For more information about shared and dedicated interfaces, see “Dedicated and Shared Interfaces” on page 15.

When a packet arrives at a NetScreen device that has virtual systems, it performs the following steps to associate the packet with a vsys.

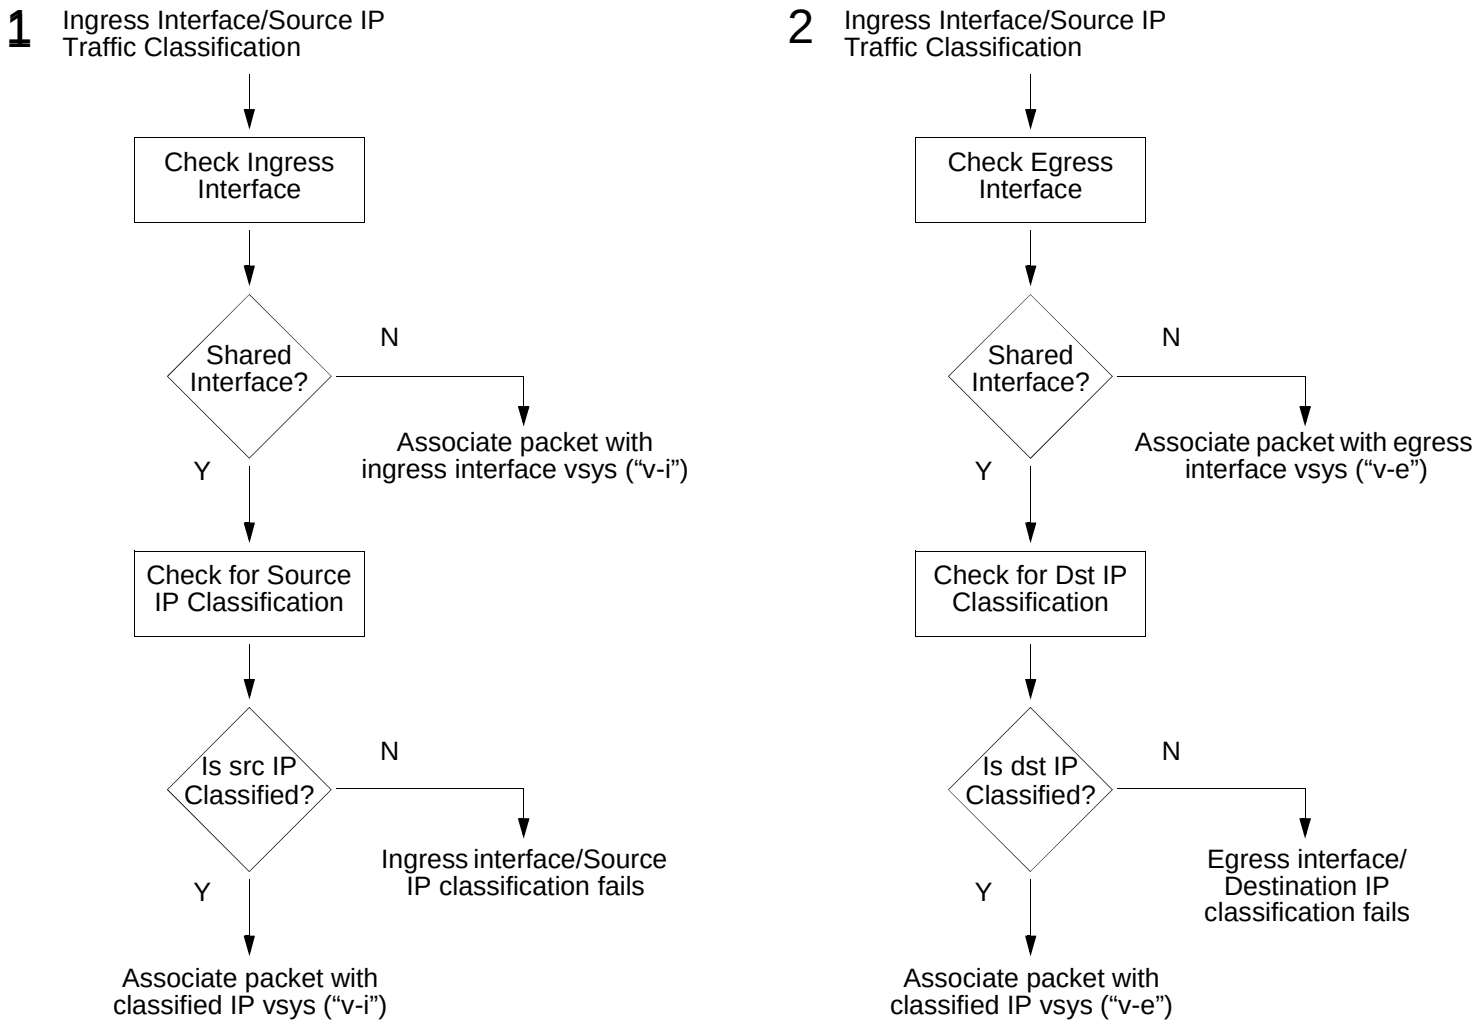

3. Vsys Traffic Assignment

Based on the outcome of the ingress interface/source IP (I/S) and egress interface/destination IP (E/D) traffic classifications, the NetScreen device determines the vsys to which traffic belongs.

- If I/S traffic classification succeeds, but E/D traffic classification fails, the NetScreen device uses the policy set and route table for the vsys associated with the ingress interface or source IP address (a vsys named “v-i”, for example).

  I/S traffic classification is particularly useful when permitting outbound traffic from a vsys to a public network such as the Internet.

- If E/D traffic classification succeeds, but I/S traffic classification fails, the NetScreen device uses the policy set and route table for the vsys associated with the egress interface or destination IP address (a vsys named “v-e”, for example).

  E/D traffic classification is particularly useful when permitting inbound traffic to one or more servers in a vsys from a public network such as the Internet.

- If both classification attempts succeed and the associated virtual systems are the same, the NetScreen device uses the policy set and route table for that vsys.

  You can use both I/S and E/D IP traffic classification to permit traffic from specific addresses in one zone to specific addresses in another zone of the same vsys.

- If both classification attempts succeed, the associated virtual systems are different, and the interfaces are bound to the same shared security zone, the NetScreen first uses the policy set and route table for the I/S vsys, and then uses the policy set and route table for the E/D vsys.

  NetScreen supports intrazone intervsys traffic when the traffic occurs in the same shared zone. The NetScreen device first applies the “v-i” policy set and route table, loops the traffic back on the Untrust interface, and then applies the “v-e” policy set and route table. Such intrazone traffic might be common if a single company uses one shared internal zone with different virtual systems for different internal departments and wants to allow traffic between the different departments.

- If both classification attempts succeed, the associated virtual systems are different, and the interfaces are bound to different shared security zones, the NetScreen device drops the packet.

  NetScreen does not support interzone intervsys traffic between shared security zones.

- If both classification attempts succeed, the associated virtual systems are different, and the ingress and egress interfaces are bound to zones dedicated to different virtual systems, the NetScreen device first applies the “v-i” policy set and route table. It then loops the traffic back on the Untrust interface and

applies the “v-e” policy set and route table. (See “Example: InterVsys Communication” on page 28.) NetScreen supports interzone intervsys traffic between dedicated security zones.

– If both classification attempts fail, the NetScreen device drops the packet.

3 After performing the ingress interface/source IP (I/S) and egress interface/destination IP (E/D) classifications, the NetScreen device uses its findings to determine the traffic classification.

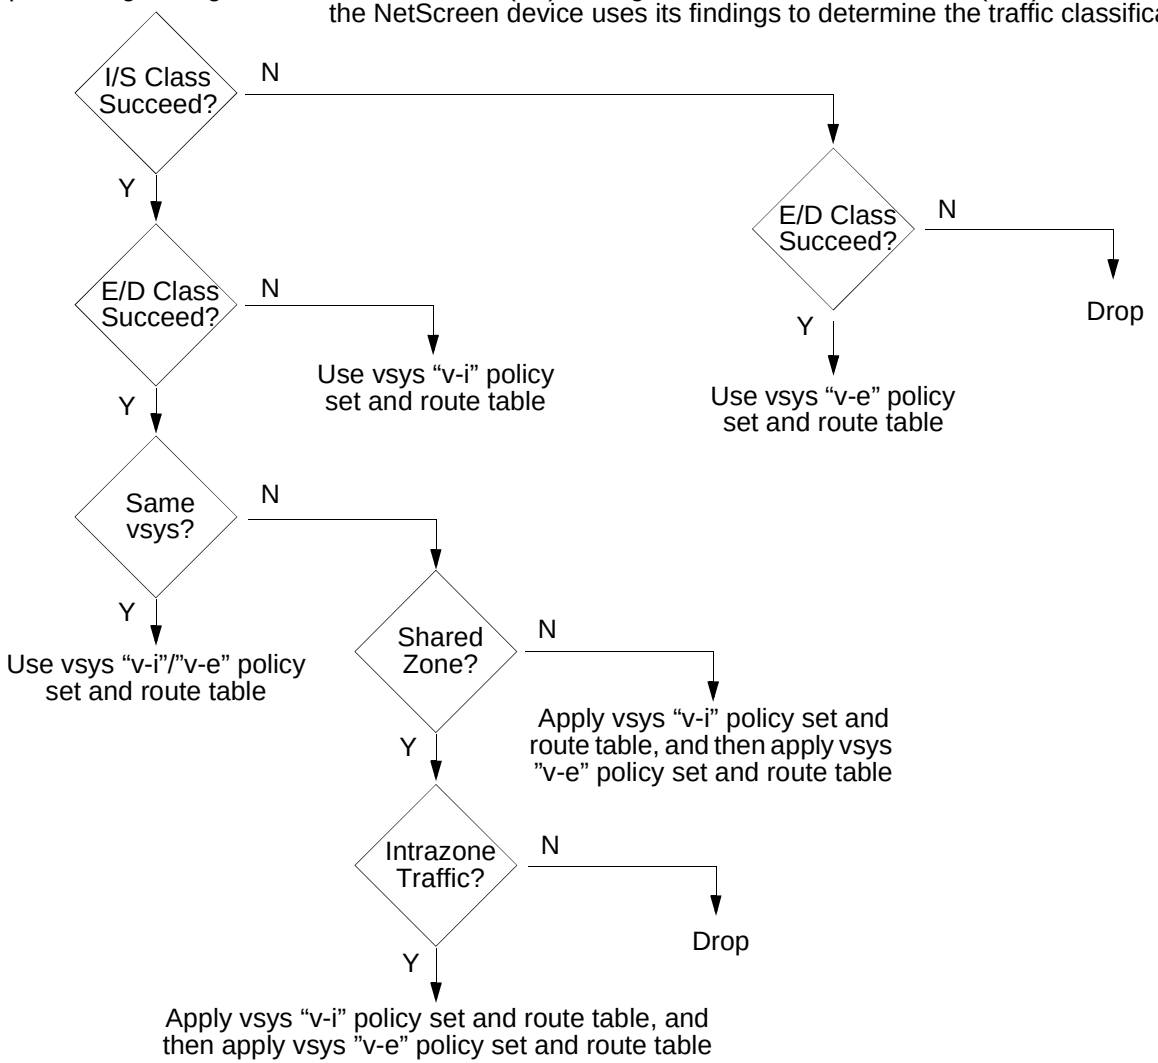

Dedicated and Shared Interfaces

There are two kinds of interfaces that affect how a NetScreen device can correctly sort inbound traffic to the right system: dedicated and shared.

Dedicated Interfaces

A system—virtual and root—can have multiple interfaces or subinterfaces dedicated exclusively to its own use. Such interfaces are not sharable by other systems. You can dedicate an interface to a system as follows:

- When you configure a physical interface, subinterface, redundant interface, or aggregate interface in the root system and bind it to a non-sharable zone, that interface remains dedicated to the root system.
- When you import a physical or aggregate interface into a vsys and bind it to either the shared Untrust zone or the Trust-*vsys_name* zone, that interface becomes a dedicated interface for that vsys.
- When you configure a subinterface in a vsys, it belongs to that vsys.

Note: *When a system has a dedicated subinterface, the NetScreen device must employ VLAN-based traffic classification to properly sort inbound traffic.*

Shared Interfaces

A system—virtual and root—can share an interface with another system. For an interface to be sharable, you must configure it at the root level and bind it to a shared zone in a shared virtual router. By default, the predefined untrust-vr is a shared virtual router, and the predefined Untrust zone is a shared zone. Consequently, a vsys can share any root-level physical interface, subinterface, redundant interface, or aggregate interface that you bind to the Untrust zone.

To create a shared interface in a zone other than the Untrust zone, you must define the zone as a shared zone at the root level[8]. To do that, the zone must be in a shared virtual router, such as the untrust-vr or any other root-level virtual router that you define as sharable. Then, when you bind a root-level interface to the shared zone, it automatically becomes a shared interface.

Note: *To create a virtual router, you need to obtain a vsys license key, which provides you with the ability to define virtual systems, virtual routers, and security zones for use either in a vsys or in the root system.*

8. For the shared zone option to be available, the NetScreen device must be operating at Layer 3, which means that you must previously assign an IP address to at least one root-level interface.

A shared virtual router can support both shared and non-sharable root-level security zones. You can define a root-level zone bound to a shared virtual router as sharable or not. Any root-level zone that you bind to a shared virtual router and define as sharable becomes a shared zone, available for use by virtual systems too. Any root-level zone that you bind to a shared virtual router and define as non-sharable remains a dedicated zone for use by the root system alone. If you bind a vsys-level zone to either the virtual router dedicated to that vsys or to a shared virtual router created in the root system, the zone remains a dedicated zone, available for use only by the vsys for which you created it.

A shared zone can support both shared and dedicated interfaces. Any root-level interface that you bind to a shared zone becomes a shared interface, available for use by virtual systems also. Any vsys-level interface that you bind to a shared zone remains a dedicated interface, available for use only by the vsys for which you created it.

A non-sharable zone can only be used by the system in which you created it and can only support dedicated interfaces for that system. All vsys-level zones are non-sharable.

To create a shared interface, you must create a shared virtual router (or use the predefined untrust-vr), create a shared security zone (or use the predefined Untrust zone), and then bind the interface to the shared zone. You must do all three steps in the root system.

The options in the WebUI and CLI are as follows:

1. To create a shared virtual router:

   ### WebUI

   Network > Routing > Virtual Routers > New: Select the **Shared and accessible by other vsys** option, and then click **Apply**.

   ### CLI

   set vrouter name *name_str*

   set vrouter *name_str* **shared**

   (You cannot modify an existing shared virtual router to make it unshared unless you first delete all virtual systems. However, you can modify a virtual router from unshared to shared at any time.)

2. To create a shared zone, do the following at the root level:

WebUI

Note: *At the time of this release, you can only define a shared zone through the CLI.*

CLI

set zone name *name_str*

set zone *zone* **vrouter** *sharable_vr_name_str*

set zone *zone* **shared**

3. To create a shared interface, do the following at the root level:

WebUI

Network > Interfaces > New (or Edit for an existing interface): Configure the interface and bind it to a shared zone, and then click **OK**.

CLI

set interface *interface* **zone** *shared_zone_name_str*

When two or more systems share an interface, the NetScreen device must employ IP-based traffic classification to properly sort inbound traffic. (For more information about IP-based traffic classification, including an example showing how to configure it for several vsys, see "IP-Based Traffic Classification" on page 33.)

Importing and Exporting Physical Interfaces

You can dedicate one or more physical interfaces to a vsys. In effect, you import a physical interface from the root system to a virtual system. After importing a physical interface to a vsys, the vsys has exclusive use of it.

Note: *Before you can import an interface to a virtual system, it must be in the Null zone at the root level.*

Example: Importing a Physical Interface to a Virtual System

In this example, you—as the root admin—import the physical interface ethernet4/1 to vsys1. You bind it to the Untrust zone and assign it the IP address 1.1.1.1/24.

WebUI

1. Entering Vsys1

   Vsys: Click **Enter** (for vsys1).

2. Importing and Defining the Interface

   Network > Interfaces: Click **Import** (for ethernet4/1).

   Network > Interfaces > Edit (for ethernet4/1): Enter the following, and then click **OK**:

   Zone Name: Untrust

   IP Address/Netmask: 1.1.1.1/24

3. Exiting Vsys1

   Click the **Exit Vsys** button (at the bottom of the menu column) to return to the root level.

CLI

1. Entering Vsys1

```
ns-> enter vsys vsys1
```

2. Importing and Defining the Interface

```
ns(vsys1)-> set interface ethernet4/1 import
ns(vsys1)-> set interface ethernet4/1 zone untrust
ns(vsys1)-> set interface ethernet4/1 ip 1.1.1.1/24
ns(vsys1)-> save
```

3. Exiting Vsys1

```
ns(vsys1)-> exit
```

Example: Exporting a Physical Interface from a Virtual System

In this example, you bind the physical interface ethernet4/1 to the Null zone in vsys1 and assign it the IP address 0.0.0.0/0. Then you export interface ethernet4/1 to the root system.

WebUI

1. Entering Vsys1

   Vsys: Click **Enter** (for vsys1).

2. Exporting the Interface

   Network > Interfaces > Edit (for ethernet4/1): Enter the following, and then click **OK**:

   Zone Name: Null

   IP Address/Netmask: 0.0.0.0/0

   Network > Interfaces: Click **Export** (for ethernet4/1).

   (Interface ethernet4/1 is now available for use in the root system or in another vsys.)

3. Exiting Vsys1

Click the **Exit Vsys** button (at the bottom of the menu column) to return to the root level.

CLI

1. Entering Vsys1

```
ns-> enter vsys vsys1
```

2. Exporting the Interface

```
ns(vsys1)-> unset interface ethernet4/1 ip
ns(vsys1)-> unset interface ethernet4/1 zone
ns(vsys1)-> unset interface ethernet4/1 import
This command will remove all objects associated with interface, continue? y/[n] y
ns(vsys1)-> save
```

(Interface ethernet4/1 is now available for use in the root system or in another vsys.)

3. Exiting Vsys1

```
ns(vsys1)-> exit
```

VLAN-BASED TRAFFIC CLASSIFICATION

With the VLAN-based traffic classification, a NetScreen device uses VLAN tagging[9] to direct traffic to various subinterfaces bound to different systems[10]. By default, a vsys has two security zones—a shared Untrust zone and its own Trust zone. Each vsys can share the Untrust zone interface with the root system and with other virtual systems. A vsys can also have its own subinterface or a dedicated physical interface (imported from the root system) bound to the Untrust zone.

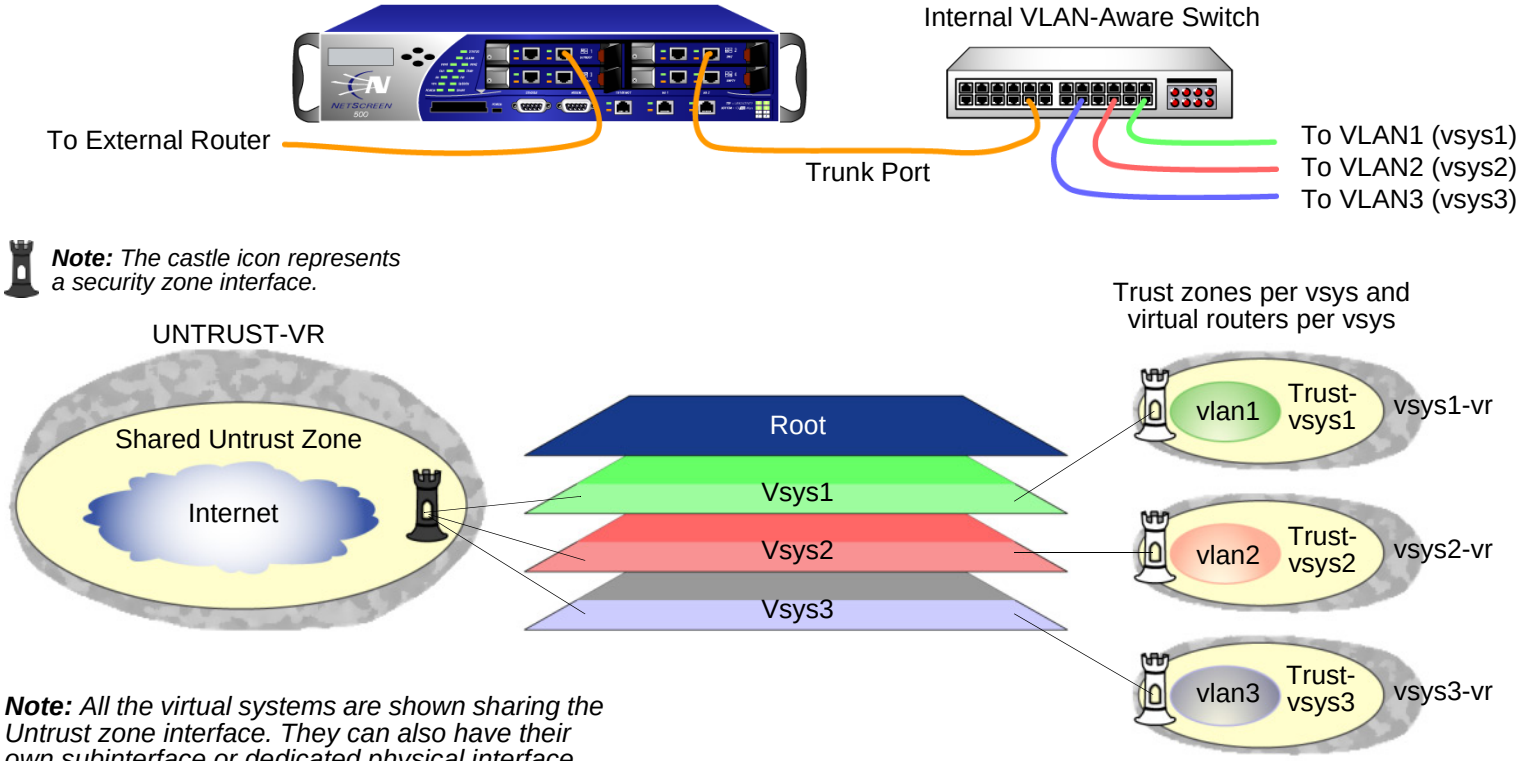

9. NetScreen supports VLANs compliant with the IEEE 802.1Q VLAN standard.
10. You can dedicate a physical interface to a virtual system by importing it from the root system to the virtual system. (See "Importing and Exporting Physical Interfaces" on page 18.) When using physical interfaces, VLAN tagging is unnecessary for traffic on that interface.

VLANs

Each VLAN is bound to a system through a subinterface. If a vsys shares the Untrust zone interface with the root system and has a subinterface bound to its Trust-*vsys_name* zone, the vsys must be associated with a VLAN in the Trust-*vsys_name* zone. If the vsys also has its own subinterface bound to the Untrust zone, the vsys must also be associated with another VLAN in the Untrust zone.

A subinterface stems from a physical interface, which then acts as a trunk port. A trunk port allows a Layer 2 network device to bundle traffic from several VLANs through a single physical port, sorting the various packets by the VLAN identifier (VID) in their frame headers. VLAN trunking allows one physical interface to support multiple logical subinterfaces, each of which must be identified by a unique VLAN tag. The VLAN identifier (tag) on an incoming ethernet frame indicates its intended subinterface—and hence the system—to which it is destined. When you associate a VLAN with an interface or subinterface, the NetScreen device automatically defines the physical port as a trunk port. When using VLANs at the root level in Transparent mode, you must manually define all physical ports as trunk ports with the following CLI command: **set interface vlan1 vlan trunk**.

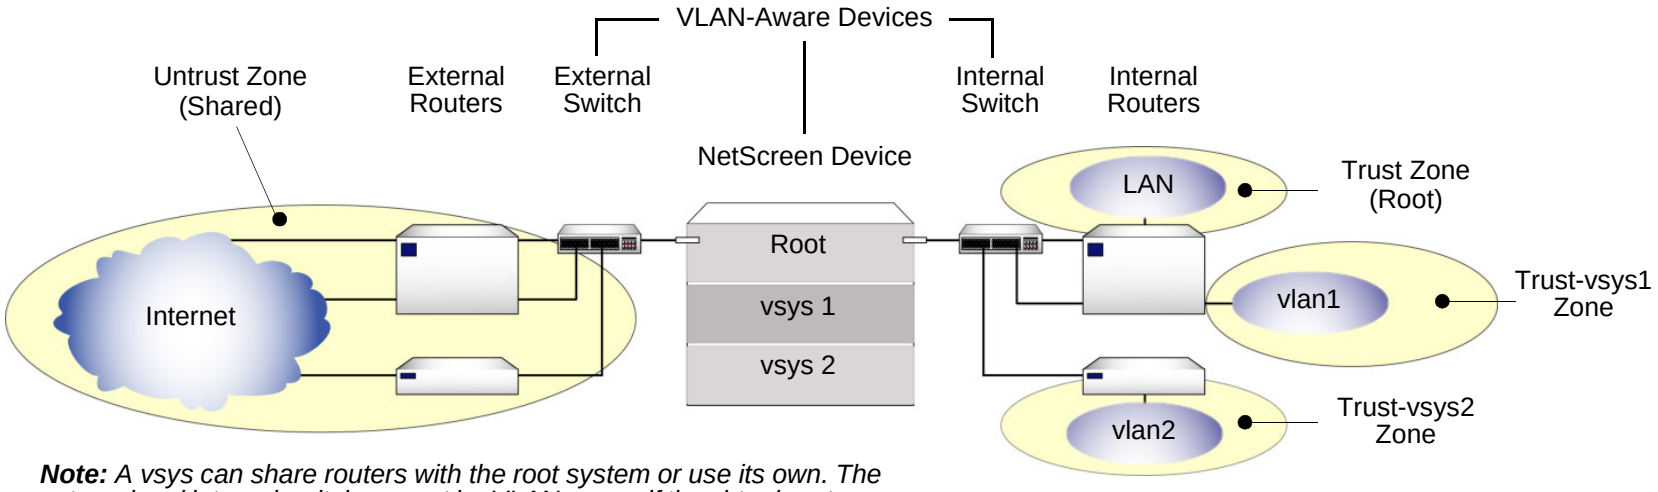

Note: *A vsys can share routers with the root system or use its own. The external and internal switches must be VLAN-aware if the virtual systems have subinterfaces bound to the Untrust and Trust-*vsys_name *zones.*

When a vsys uses a subinterface (not a dedicated physical interface) bound to the Trust-*vsys_name* zone, the internal switch and internal router in the Trust-*vsys_name* zone must have VLAN-support capabilities. If you create more than one subinterface on a physical interface, then you must define the connecting switch port as a trunk port and make it a member of all VLANs that use it.

When a vsys uses a subinterface (not a shared interface or a dedicated physical interface) bound to the shared Untrust zone, the external switch and external router that receives its inbound and outbound traffic must have VLAN-support capabilities. The router tags the incoming frames so that when they reach the NetScreen device, it can direct them to the correct subinterface.

Although a vsys cannot be in Transparent mode, because it requires unique interface or subinterface IP addresses, the root system can be in Transparent mode[11]. For the root system to support VLANs while operating in Transparent mode, use the following CLI command to enable the physical interfaces bound to Layer 2 security zones to act as trunk ports: **set interface vlan1 vlan trunk**.

Defining Subinterfaces and VLAN Tags

The Trust-*vsys_name* zone subinterface links a vsys to its internal VLAN. The Untrust zone subinterface links a vsys to the public WAN, usually the Internet. A subinterface has the following attributes:

- A unique VLAN ID (from 1 to 4095)
- A public or private IP address[12] (the IP address is private by default)
- A netmask for a class A, B, or C subnet
- An associated VLAN

A vsys can have a single Untrust zone subinterface and multiple Trust-*vsys_name* zone subinterfaces. If a virtual system does not have its own Untrust zone subinterface, it shares the root level Untrust zone interface. NetScreen devices also support subinterfaces and VLANs at the root level.

11. When the root system is in Transparent mode, it cannot support virtual systems. It can, however, support root-level VLANs while in Transparent mode.
12. For information about public and private IP addresses, see “Public IP Addresses” on page **2**-77 and “Private IP Addresses” on page **2**-78.

vsys1 shares the Untrust zone interface with the root system. **vsys2** and **vsys100** have their own dedicated subinterfaces bound to the Untrust zone.

The **root system** has a physical interface and a subinterface bound to its Trust zone.

vsys1 has three subinterfaces bound to its Trust-vsys1 zone, each leading to a different VLAN.

vsys2 has two subinterfaces bound to its Trust-vsys2 zone, each leading to a different VLAN.

vsys100 has one subinterface bound to its Trust-vsys100 zone.

Note: *All VLAN IDs must be unique per physical interface.*

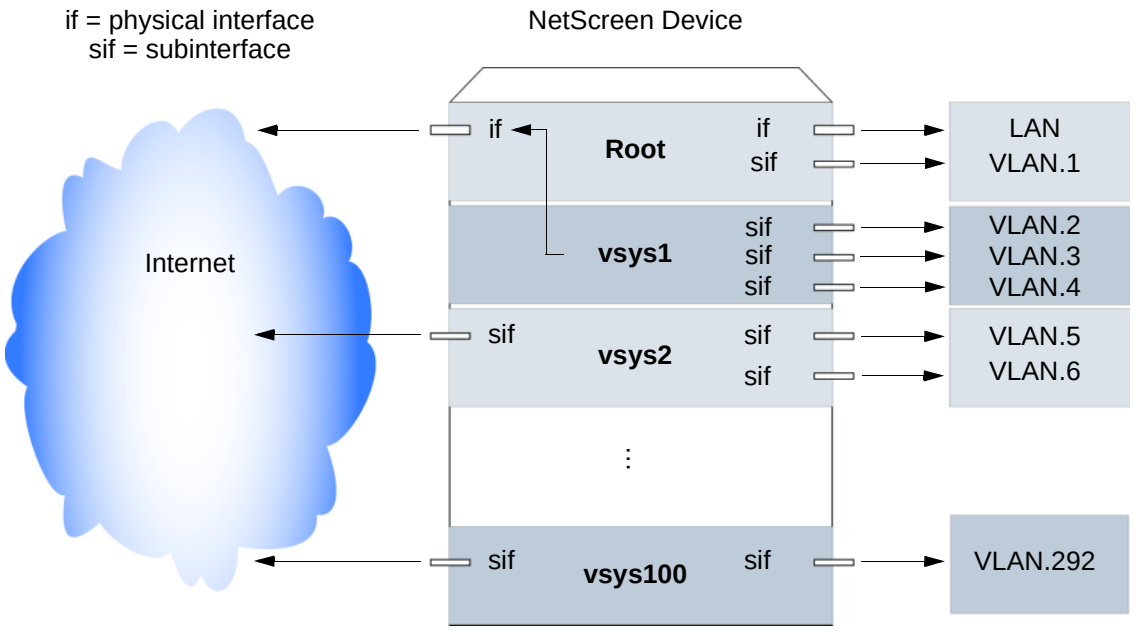

The NetScreen device supports IEEE 802.1Q-compliant VLAN tags. A tag is an added bit in the Ethernet frame header that indicates membership in a particular VLAN. By binding a VLAN to a vsys, the tag also determines to which vsys a frame belongs, and consequently, which policy is applied to that frame. If a VLAN is not bound to a vsys, policies set in the root system of the NetScreen device are applied to the frame.

A root-level administrator can create a VLAN, assign members to it, and bind it to a vsys. (The assigning of members to a VLAN can be done by several methods—protocol type, MAC address, port number—and is beyond the scope of this document.) The vsys admin, if there is one, then manages the vsys through the creation of addresses, users, services, VPNs, and policies. If there is no vsys admin, then a root-level administrator performs these tasks.

Note: *If the root-level admin does not associate a VLAN to a vsys, the VLAN operates within the NetScreen device root system.*

There are three tasks that a root-level administrator must perform to create a VLAN for a vsys: Enter a virtual system, define a subinterface, and associate it with a VLAN.

Note: *All subnets in a vsys must be disjointed; that is, there must be no overlapping IP addresses among the subnets in the same vsys. For example: Subinterface1 – 10.2.2.1 255.255.255.0 and Subinterface2 – 10.2.3.1 255.255.255.0 are disjointed, and therefore, link to acceptable subnets.*

However, subnets with the following subinterfaces overlap, and are unacceptable within the same vsys: subinterface1 – 10.2.2.1 255.255.0.0 and subinterface2 – 10.2.3.1 255.255.0.0.

The address ranges of subnets in different virtual systems can overlap.

Example: Defining Three Subinterfaces and VLAN Tags

In this example, you define subinterfaces and VLAN tags for the three virtual systems that you created in "Example: Vsys Objects and Admins" on page 3—vsys1, vsys2, and vsys3. The first two subinterfaces are for two private virtual systems operating in NAT mode, and the third subinterface is for a public virtual system operating in Route mode. The subinterfaces are 10.1.1.1/24, 10.2.2.1/24, and 1.3.3.1/24. You create all three subinterfaces on ethernet3/2.

All three virtual systems share the Untrust zone and its interface (ethernet1/1; 1.1.1.1/24) with the root system. The Untrust zone is in the untrust-vr routing domain.

WebUI

1. Vsys1 Subinterface and VLAN Tag

Vsys: Click **Enter** (for vsys1).

Network > Interfaces > New Sub-IF (for ethernet3/2): Enter the following, and then click **OK**:

Interface Name: ethernet3/2.1

Zone Name: Trust-vsys1

IP Address / Netmask: 10.1.1.1/24

VLAN Tag: 1[13]

2. Vsys2 Subinterface and VLAN Tag

Vsys: Click **Enter** (for vsys2).

Network > Interfaces > New Sub-IF (for ethernet3/2): Enter the following, and then click **OK**:

Interface Name: ethernet3/2.2

Zone Name: Trust-vsys2

IP Address / Netmask: 10.2.2.1/24

VLAN Tag: 2

3. Vsys3 Subinterface and VLAN Tag

Vsys: Click **Enter** (for vsys3).

Network > Interfaces > New Sub-IF (for ethernet3/2): Enter the following, and then click **Apply**:

Interface Name: ethernet3/2.3

Zone Name: Trust-vsys3

IP Address / Netmask: 1.3.3.1/24

VLAN Tag: 3

Select **Interface Mode: Route**, and then click **OK**.

Click **Exit Vsys** to return to the root level.

13. You can define virtual systems to operate in Route mode or NAT mode. The default is NAT mode, and thus unnecessary to specify when creating the first two subinterfaces in this example.

CLI

1. Vsys1 Subinterface and VLAN Tag

```
ns-> enter vsys vsys1
ns(vsys1)-> set interface ethernet3/2.1 zone trust-vsys1
ns(vsys1)-> set interface ethernet3/2.1 ip 10.1.1.1/24 tag 1[14]
ns(vsys1)-> save
ns(vsys1)-> exit
```

2. Vsys2 Subinterface and VLAN Tag

```
ns-> enter vsys vsys2
ns(vsys2)-> set interface ethernet3/2.2 zone trust-vsys2
ns(vsys2)-> set interface ethernet3/2.2 ip 10.2.2.1/24 tag 2
ns(vsys2)-> save
ns(vsys2)-> exit
```

3. Vsys3 Subinterface and VLAN Tag

```
ns-> enter vsys vsys3
ns(vsys3)-> set interface ethernet3/2.3 zone trust-vsys3
ns(vsys3)-> set interface ethernet3/2.3 ip 1.3.3.1/24 tag 3
ns(vsys3)-> set interface ethernet3/2.3 route
ns(vsys3)-> save
ns(vsys3)-> exit
```

14. You can define virtual systems to operate in Route mode or NAT mode. The default is NAT mode, and thus unnecessary to specify when creating the first two subinterfaces in this example.

Communicating between Virtual Systems

The members of a VLAN in a vsys have unrestricted communication access with each other. The VLAN members of different virtual systems cannot communicate with one another unless the participating vsys administrators specifically configure policies allowing the members of their respective systems to do so.

Traffic between root-level VLANs operates within the parameters set by root-level policies. Traffic between virtual system VLANs operates within the parameters set by the participating virtual system policies[15]. The NetScreen device passes only traffic allowed to leave the originating virtual system and allowed to enter the destination virtual system. In other words, the vsys admins of both virtual systems must set policies allowing the traffic to flow in the appropriate direction—outgoing and incoming.

Example: InterVsys Communication

In this example, the admins for vsys1 and vsys2—see "Example: Defining Three Subinterfaces and VLAN Tags" on page 25—set up policies to enable traffic between a workstation (work_js with the IP address 10.1.1.10/32) in VLAN 1 and a server (ftp_server with the IP address 10.2.2.20/32) in VLAN 2. The connection is possible if the following two conditions are met:

- The vsys admin for vsys1 has set a policy permitting traffic from the workstation in Trust-vsys1 to the server in its Untrust zone.
- The vsys admin for vsys2 has set a policy permitting traffic from the workstation in its Untrust zone to the server in Trust-vsys2.

Notice that the network device in front of the internal interface on the NetScreen device is a Layer 2 switch. This forces traffic from VLAN 1 going to VLAN 2 to go through the switch to the NetScreen device for Layer 3 routing. If the network device were a Layer 3 router, traffic between VLAN1 and VLAN2 could pass through the router, bypassing all policies set on the NetScreen device.

The vsys1 and vsys2 admins also set up the appropriate routes. The shared Untrust zone is in the untrust-vr and the Trust zones in vsys1 and vsys2.

15. Policies set in the root system do not affect policies set in virtual systems, and vice versa.

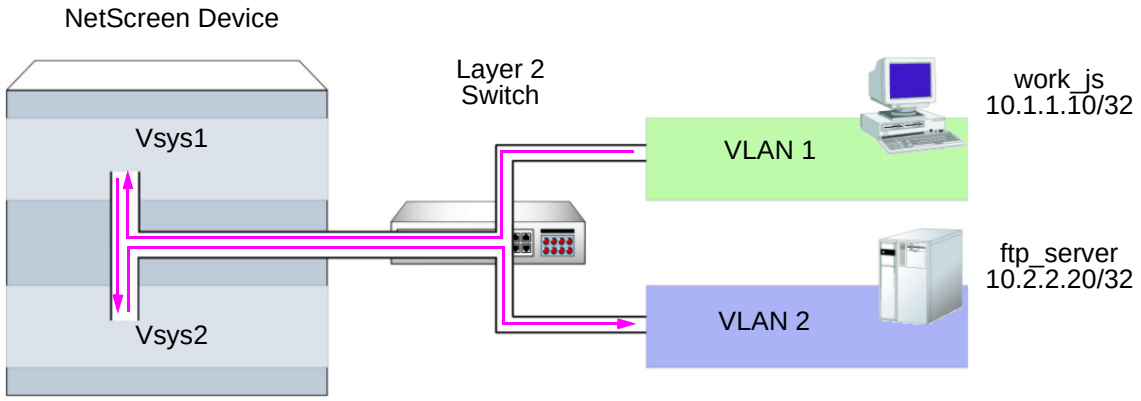

WebUI

1. Vsys1

Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: work_js

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.10/32

Zone: Trust-vsys1

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: ftp_server

IP Address/Domain Name:

IP/Netmask: (select), 10.2.2.20/32

Zone: Untrust

Routes

Network > Routing > Routing Entries > untrust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 10.1.1.0/24

Next Hop Virtual Router Name: (select); vsys1-vr

Network > Routing > Routing Entries > vsys1-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Next Hop Virtual Router Name: (select); untrust-vr

Policy

Policies > (From: Trust-vsys1, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), work_js

Destination Address:

Address Book Entry: (select), ftp_server

Service: FTP-Get

Action: Permit

2. Vsys2

Addresses

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: ftp_server

IP Address/Domain Name:

IP/Netmask: (select), 10.2.2.2/032

Zone: Trust-vsys2

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: work_js

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.10/32

Zone: Untrust

Routes

Network > Routing > Routing Entries > untrust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 10.2.2.0/24

Next Hop Virtual Router Name: (select); vsys2-vr

Network > Routing > Routing Entries > vsys2-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Next Hop Virtual Router Name: (select); untrust-vr

Policy

Policies > (From: Untrust, To: Trust-vsys2) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), work_js

Destination Address:

Address Book Entry: (select), ftp_server

Service: FTP-Get

Action: Permit

CLI

1. Vsys1

   Addresses

   ```
   set address trust-vsys1 work_js 10.1.1.10/32
   set address untrust ftp_server 10.2.2.20/32
   ```

   Routes

   ```
   set vrouter untrust-vr route 10.1.1.0/24 vrouter vsys1-vr
   set vrouter vsys1-vr route 0.0.0.0/0 vrouter untrust-vr
   ```

   Policy

   ```
   set policy from trust-vsys1 to untrust work_js ftp_server ftp-get permit
   save
   ```

2. Vsys2

   Addresses

   ```
   set address trust-vsys2 ftp_server 10.2.2.20/32
   set address untrust work_js 10.1.1.10/32
   ```

3. Routes

   ```
   set vrouter untrust-vr route 10.2.2.0/24 vrouter vsys2-vr
   set vrouter vsys2-vr route 0.0.0.0/0 vrouter untrust-vr
   ```

4. Vsys2 Policy

   ```
   set policy from untrust to trust-vsys2 work_js ftp_server ftp-get permit
   save
   ```

IP-BASED TRAFFIC CLASSIFICATION

IP-based traffic classification allows you to use virtual systems without VLANs. Instead of VLAN tags, the NetScreen device uses IP addresses to sort traffic, associating a subnet or range of IP addresses with a particular system—root or vsys. Using IP-based traffic classification exclusively to sort traffic, all systems share the following:

- The untrust-vr and a user-defined internal-vr
- The Untrust zone and a user-defined internal zone
- An Untrust zone interface and a user-defined internal zone interface[16]

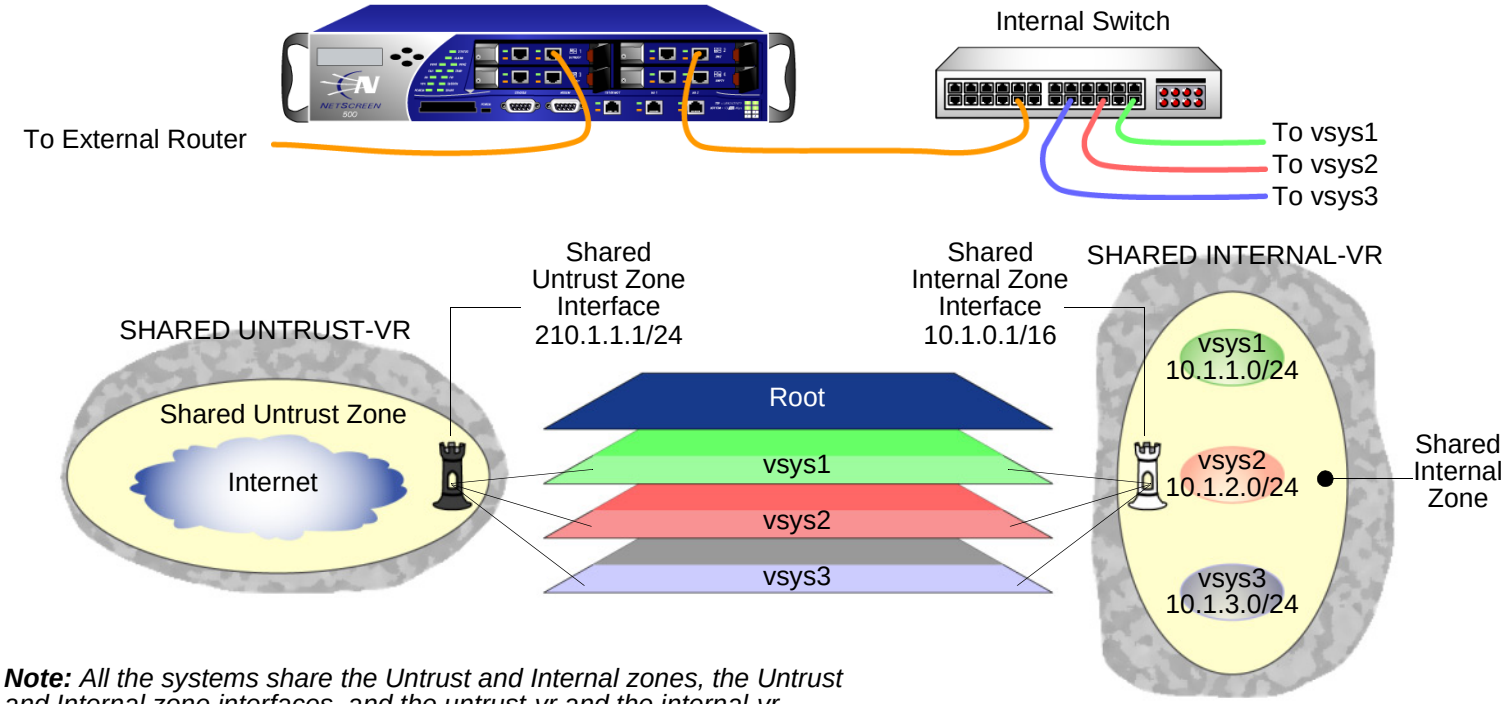

Note: *All the systems share the Untrust and Internal zones, the Untrust and Internal zone interfaces, and the untrust-vr and the internal-vr.*

16. Even when using VLAN-based traffic classification for internal traffic, for external traffic all systems use the shared Untrust zone—and, unless a system has a dedicated interface, a shared Untrust zone interface. Using a shared interface on one side and a dedicated interface (with VLAN tagging) on the other constitutes a hybrid approach. VLAN-based and IP-based traffic classification can coexist within the same system or among different systems simultaneously.

To designate a subnet or range of IP addresses to the root system or to a previously created virtual system, you must do either of the following at the root level:

WebUI

Network > Zones > Edit (for *zone*) > IP Classification: Enter the following, and then click **OK**:

System: (select **root** or ***vsys_name_str***)

Address Type: (select **Subnet** and enter ***ip_addr/mask***, or select **Range** and enter ***ip_addr1*** – ***ip_addr2***)

CLI

set zone *zone* **ip-classification net** *ip_addr/mask* { **root** | **vsys** *name_str* }

set zone *zone* **ip-classification range** *ip_addr1-ip_addr2* { **root** | **vsys** *name_str* }

Because IP-based traffic classification requires the use of a shared security zone, virtual systems cannot use overlapping internal IP addresses, as is possible with VLAN-based traffic classification. Also, because all the systems share the same internal interface, the operational mode for that interface must be either NAT or Route mode; you cannot mix NAT and Route modes for different systems. In this regard, the addressing scheme of an IP-based approach is not as flexible as that allowed by the more commonly used VLAN-based approach.

Furthermore, sharing virtual routers, security zones, and interfaces is inherently less secure than dedicating an internal virtual router, internal security zone, and internal and external interfaces to each vsys. When all virtual systems share the same interfaces, it is possible for a vsys admin in one vsys to use the **snoop** command to gather information about the traffic activities of another vsys. Also, because IP spoofing is possible on the internal side, NetScreen recommends that you disable the IP spoofing SCREEN option on the shared internal interface. When deciding which traffic classification scheme to use, you must weigh the ease of management offered by the IP-based approach against the increased security and greater addressing flexibility offered by the VLAN-based approach.

Example: Configuring IP-Based Traffic Classification

In this example, you set up IP-based traffic classification for the three virtual systems created in "Example: Vsys Objects and Admins" on page 3. You define the trust-vr as sharable. You create a new zone, name it *Internal*, and bind it to the trust-vr. You then make the Internal zone sharable. You bind ethernet3/2 to the shared Internal zone, assign it IP address 10.1.0.1/16, and select NAT mode.

You bind ethernet1/2 to the shared Untrust zone and assign it IP address 210.1.1.1/24. The IP address of the default gateway in the Untrust zone is 210.1.1.250. Both the Internal and Untrust zones are in the shared trust-vr routing domain.

The subnets and their respective vsys associations are as follows:

- 10.1.1.0/24 – vsys1
- 10.1.2.0/24 – vsys2
- 10.1.3.0/24 – vsys3

WebUI

1. Virtual Routers, Security Zones, and Interfaces

Network > Routing > Virtual Routers > Edit (for trust-vr): Select the **Shared and accessible by other vsys** check box, and then click **OK**.

Network > Zones > New: Enter the following, and then click **OK**:

Zone Name: Internal

Virtual Router Name: trust-vr

Zone Type: Layer 3

Network > Zones > Edit (for Internal): Select the **Share Zone** check box, and then click **OK**.

Network > Interfaces > Edit (for ethernet3/2): Enter the following, and then click **OK**:

Zone Name: Internal

IP Address/Netmask: 10.1.0.1/16

Network > Interfaces > Edit (for ethernet1/2): Enter the following, and then click **OK**:

Zone Name: Untrust

IP Address/Netmask: 210.1.1.1/24

2. Route

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet1/2

Gateway IP Address: 210.1.1.250

3. IP Classification of the Trust Zone

Network > Zones > Edit (for Internal) > IP Classification: Enter the following, and then click **OK**:

System: vsys1

Address Type:

Subnet: (select); 10.1.1.0/24

Network > Zones > Edit (for Internal) > IP Classification: Enter the following, and then click **OK**:

System: vsys2

Address Type:

Subnet: (select); 10.1.2.0/24

Network > Zones > Edit (for Internal) > IP Classification: Enter the following, and then click **OK**:

System: vsys3

Address Type:

Subnet: (select); 10.1.3.0/24

Network > Zones > Edit (for Internal): Select the **IP Classification** check box, and then click **OK**.

CLI

1. Virtual Routers, Security Zones, and Interfaces

```
set vrouter trust-vr shared
set zone name Internal
set zone Internal shared
set interface ethernet3/2 zone Internal
set interface ethernet3/2 ip 10.1.0.1/16
set interface ethernet3/2 nat
set interface ethernet1/2 zone untrust
set interface ethernet1/2 ip 210.1.1.1/24
```

2. Route

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet1/2 gateway 210.1.1.250
```

3. IP Classification of the Trust Zone

```
set zone Internal ip-classification net 10.1.1.0/24 vsys1
set zone Internal ip-classification net 10.1.2.0/24 vsys2
set zone Internal ip-classification net 10.1.3.0/24 vsys3
set zone Internal ip-classification
save
```

LOGGING ON AS A VSYS ADMIN

Whereas a root-level administrator enters a vsys from the root level, a vsys admin enters his or her vsys directly. When a root-level administrator exits a vsys, he or she exits to the root system. When a vsys admin exits a vsys, the connection is immediately severed.

The following example shows how to log on to a vsys as a vsys admin, change your password, and log out.

Example: Logging On and Changing Your Password

In this example, you, as a vsys admin, log on to vsys1 by entering your assigned login name jsmith and password Pd50iH10. You change your password to I6Dls13guh, and then log out.

Note: *A vsys admin cannot change his or her login name (user name) because the NetScreen device uses that name, which must be unique among all vsys admins, to route the login connection to the appropriate vsys.*

WebUI

1. Logging On

In the URL field in your Web browser, enter the Untrust zone interface IP address for vsys1.

When the Network Password dialog box appears, enter the following, and then click **OK**:

User Name: jsmith

Password: Pd50iH10

2. Changing your Password

Configuration > Admin > Administrators: Enter the following, and then click **OK**:

Vsys Admin Old Password: Pd50iH10

Vsys Admin New Password: I6Dls13guh

Confirm New Password: I6Dls13guh

3. Logging Out

Click **Logout**, located at the bottom of the menu column.

CLI

1. Logging On

From a Secure Command Shell (SCS), Telnet, or HyperTerminal session command-line prompt, enter the Untrust zone interface IP address for vsys1.

Log on with the following user name and password:

- User Name: jsmith
- Password: Pd50iH10

2. Changing your Password

```
set admin password I6Dls13guh
save
```

3. Logging Out

```
exit
```

Index

A

administration
- vsys admin 38

C

character types, ScreenOS supported viii
CLI
- conventions iv

conventions
- CLI iv
- illustration vii
- names viii
- WebUI v

D

defining
- subinterfaces 25

I

IEEE 802.1Q VLAN standard 21
illustration
- conventions vii

interfaces
- dedicated 15, 33
- exporting from vsys 19
- importing to vsys 18
- shared 15, 33

IP-based traffic classification 33

L

logging in
- vsys 33, 38

M

MIP
- virtual systems 10

N

names
- conventions viii

P

password
- vsys admin 38

ports
- trunk 23

S

ScreenOS
- virtual systems, VRs 6
- virtual systems, zones 7

security zones
- *See* zones

software
- key, vsys 15

subinterfaces 23
- configuring (vsys) 23
- creating (vsys) 23
- defining 25
- multiple subinterfaces per vsys 23

T

traffic
- classification, IP-based 33
- classification, VLAN-based 21
- through traffic, vsys sorting 11–14

trunk ports 23
- defined 22
- manually setting 22

V

VIP
- virtual systems 10

virtual system 1–39
- admin types 3
- admins iii, 1
- basic functional requirements 3
- changing admin's password 3, 38
- creating a vsys object 3
- exporting a physical interface 19
- importing a physical interface 18
- interfaces 8
- IP-based traffic classification 33–37
- manageability and security 34
- MIP 10
- overlapping address ranges 25, 34
- overlapping subnets 25
- shared VR 15
- shared zone 15
- software key 15
- traffic sorting 10–17
- Transparent mode 22
- VIP 10
- VLAN-based traffic classification 21–32
- VRs 6
- zones 7

VLANs
- communicating with another VLAN 28–32
- creating 25–27
- subinterfaces 23
- tag 23, 24
- Transparent mode 22, 23
- trunking 22
- VLAN-based traffic classification 21

VRs
- creating a shared VR 16
- shared 15

W

WebUI
- conventions v

Z

zones
- shared 15
- vsys 7

NetScreen Concepts & Examples

ScreenOS Reference Guide

Volume 8: High Availability

ScreenOS 5.0.0

P/N 093-0931-000

Rev. B

Copyright Notice

NetScreen Technologies, Inc.
Building #3
805 11th Avenue
Sunnyvale, CA 94089
www.netscreen.com

FCC Statement

The following information is for FCC compliance of Class A devices: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. The equipment generates, uses, and can radiate radio-frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case users will be required to correct the interference at their own expense.

The following information is for FCC compliance of Class B devices: The equipment described in this manual generates and may radiate radio-frequency energy. If it is not installed in accordance with NetScreen's installation instructions, it may cause interference with radio and television reception. This equipment has been tested and found to comply with the limits for a Class B digital device in accordance with the specifications in part 15 of the FCC rules. These specifications are designed to provide reasonable protection against such interference in a residential installation. However, there is no guarantee that interference will not occur in a particular installation.

If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Consult the dealer or an experienced radio/TV technician for help.
- Connect the equipment to an outlet on a circuit different from that to which the receiver is connected.

Caution: Changes or modifications to this product could void the user's warranty and authority to operate this device.

Disclaimer

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR NETSCREEN REPRESENTATIVE FOR A COPY.

Contents

Preface v
Conventions vi
CLI Conventions vi
WebUI Conventions vii
Illustration Conventions ix
Naming Conventions and Character Types x
NetScreen Documentation xi
Chapter 1 NSRP 1
NSRP Overview 3
NSRP and NetScreen Operational Modes 8
Basic Active/Passive NSRP Configuration 8
Default Settings 9
Example: NSRP for an Active/Passive Configuration 10
NSRP Clusters 15
Cluster Name 17
Example: Creating an NSRP Cluster 18
Run-Time Objects 21
RTO Mirror States 22
VSD Groups 23
Preempt Option 23
VSD Group Member States 24
Heartbeat Messages 25
Example: Creating Two VSD Groups 26
VSIs and Static Routes 28
Example: Trust and Untrust Zone VSIs 29
Synchronization 33
Synchronizing Configurations 33
Synchronizing Files 34
Synchronizing RTOs 34
Example: Manually Resynchronizing RTOs 35
Example: Adding a Device to an Active NSRP Cluster 36
Synchronizing System Clocks 37
Dual HA Interfaces 38
Control Messages 39
Data Messages (Packet Forwarding) 40
Dynamic Routing Advisory 41
Dual HA Link Probes 42
Example: Sending Link Probes Manually 43
Example: Sending Link Probes Automatically 44
Setup Procedure 45
Cabling for a Full-Mesh Configuration 45
Active/Active NSRP Configuration 49
Example: NSRP for an Active/Active Configuration 49
Chapter 2 Interface Redundancy 57
Redundant Interfaces 58
Example: Creating Redundant Interfaces for VSIs 60
Aggregate Interfaces 65
Example: Configuring an Aggregate Interface 66

Dual Untrust Interfaces 67
Interface Failover 68
Example: Manually Forcing Traffic from the Primary to the Backup Interface 68
Example: Manually Forcing Traffic from the Backup to the Primary Interface 68
Example: Automatically Switching Traffic between the Primary and Backup Interface 69
Determining Interface Failover 69
Interface Failover with IP Tracking 70
Example: Configuring Automatic Failover with IP Tracking 71
Interface Failover with VPN Tunnel Monitoring 75
Example: Configuring Automatic Failover with VPN Tunnel Monitoring 76
Serial Interface 82
Modem Settings 83
Example: Configuring Modem Settings 84
ISP Configuration 85
Example: Configuring ISP Information 86
Serial Interface Failover 87
Example: Configuring Dial Backup in the Trust-Untrust Mode 88
Example: Deleting a Default Route for the Serial Interface 91
Example: Adding a Default Route for the Serial Interface 91
Example: Specifying a Policy as Inactive for Serial Interface Failover 92
Chapter 3 Failover 93
Device Failover (NSRP) 94
VSD Group Failover (NSRP) 95
Configuring Object Monitoring for Device or VSD Group Failover 96
Configuring Monitored Objects 98
Physical Interface Objects 98
Example: Monitoring an Interface 98
Zone Objects 99
Example: Monitoring an Interface 99
Tracked IP Objects 100
Example: Track IP for Device Failover 103
Virtual System Failover 108
Example: VSIs for Inter-Virtual System Load Sharing 108
Chapter 4 NSRP-Lite 115
Introduction to NSRP-Lite 117
Clusters and VSD Groups 118
Default Settings 119
Cluster 120
Cluster Name 121
Authentication and Encryption 122
VSD Group 123
VSD Group Member States 123
Heartbeat Messages 124
Preempt Option 125
Cabling and Configuring NSRP-Lite 126
Example: Configuring NSRP-Lite 127
Configuration and File Synchronization 134
Synchronizing Configurations 134
Synchronizing Files 135
Example: Adding a Device to an Active NSRP Cluster 135

Disabling Configuration and File Synchronization....136
Path Monitoring....137
Setting Thresholds....138
Weighting Tracked IP Addresses....138
IP Tracking for VPN Tunnel Failover....139
Example: IP Tracking through a VPN Tunnel....140
Index....IX-I

Preface

Volume 8, “High Availability” presents an overview of the NetScreen Redundancy Protocol (NSRP) operations and describes how to cable, configure, and manage NetScreen devices in a redundant group to provide high availability using NSRP. This volume also describes the various ways in which interface redundancy is provided on NetScreen devices and how to configure the devices for failover when there are redundant components.

CONVENTIONS

This document contains several types of conventions, which are introduced in the following sections:

- "CLI Conventions"
- "WebUI Conventions" on page vii
- "Illustration Conventions" on page ix
- "Naming Conventions and Character Types" on page x

CLI Conventions

The following conventions are used when presenting the syntax of a command line interface (CLI) command:

- Anything inside square brackets [ ] is optional.
- Anything inside braces { } is required.
- If there is more than one choice, each choice is separated by a pipe (|). For example,

  ```
  set interface { ethernet1 | ethernet2 | ethernet3 } manage
  ```

  means "set the management options for the ethernet1, ethernet2, or ethernet3 interface".
- Variables appear in *italic.* For example:

  ```
  set admin user name password
  ```

When a CLI command appears within the context of a sentence, it is in **bold** (except for variables, which are always in *italic*). For example: "Use the **get system** command to display the serial number of a NetScreen device."

Note: *When typing a keyword, you only have to type enough letters to identify the word uniquely. For example, typing* ***set adm u joe j12fmt54*** *is enough to enter the command* ***set admin user joe j12fmt54****. Although you can use this shortcut when entering commands, all the commands documented here are presented in their entirety.*

WebUI Conventions

Throughout this book, a chevron (>) is used to indicate navigation through the WebUI by clicking menu options and links. For example, the path to the address configuration dialog box is presented as **Objects > Addresses > List > New**. This navigational sequence is shown below.

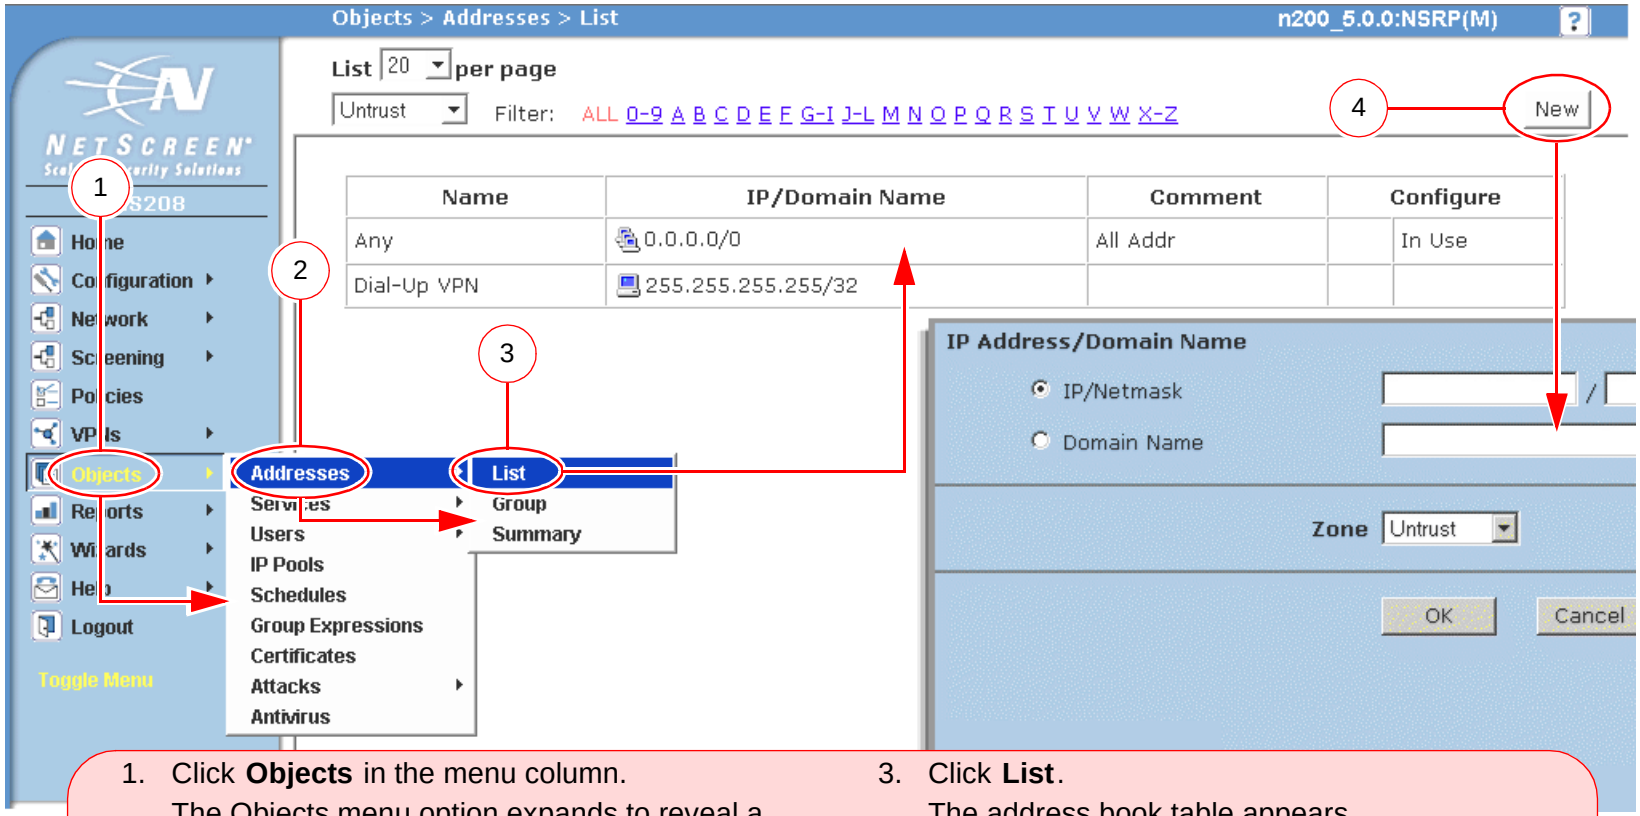

1. Click **Objects** in the menu column.
   The Objects menu option expands to reveal a subset of options for Objects.
2. (Applet menu) Hover the mouse over **Addresses**.
   (DHTML menu) Click **Addresses**.
   The Addresses option expands to reveal a subset of options for Addresses.
3. Click **List**.
   The address book table appears.
4. Click the **New** link.
   The new address configuration dialog box appears.

To perform a task with the WebUI, you must first navigate to the appropriate dialog box where you can then define objects and set parameters. The set of instructions for each task is divided into two parts: a navigational path and configuration details. For example, the following set of instructions includes the path to the address configuration dialog box and the settings for you to configure:

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: addr_1

IP Address/Domain Name:

IP/Netmask: (select), 10.2.2.5/32

Zone: Untrust

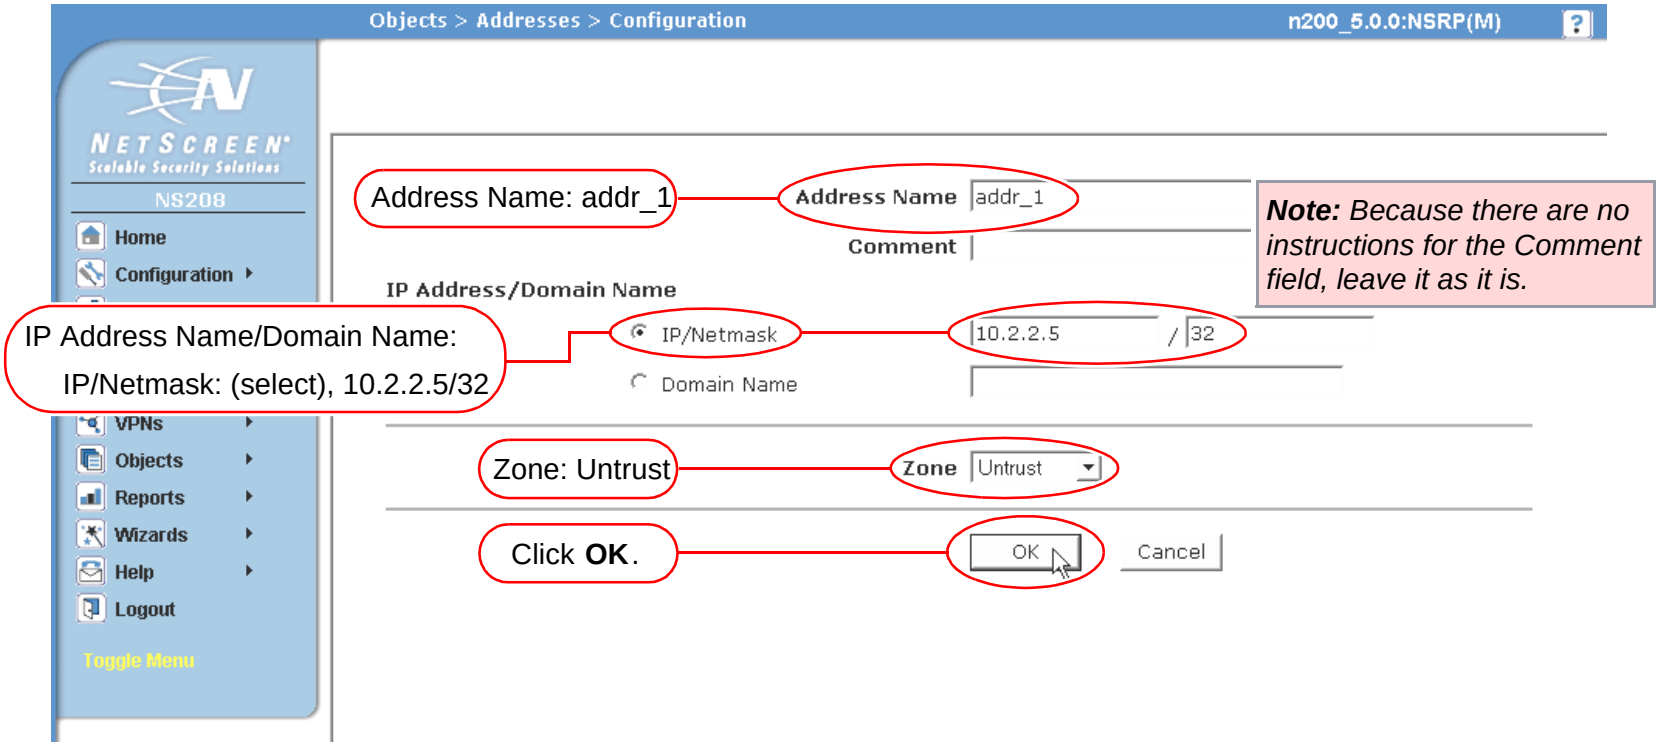

Illustration Conventions

The following graphics make up the basic set of images used in illustrations throughout this book:

Generic NetScreen Device

Virtual Routing Domain

Security Zone

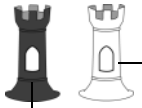

Security Zone Interfaces

White = Protected Zone Interface (example: Trust Zone)

Black = Outside Zone Interface (example: Untrust Zone)

Tunnel Interface

VPN Tunnel

Router Icon

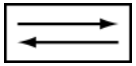

Switch Icon

Local Area Network (LAN) with a Single Subnet

(example: 10.1.1.0/24)

Internet

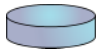

Dynamic IP (DIP) Pool

Desktop Computer

Laptop Computer

Generic Network Device

(examples: NAT server, Access Concentrator)

Server

Naming Conventions and Character Types

ScreenOS employs the following conventions regarding the names of objects—such as addresses, admin users, auth servers, IKE gateways, virtual systems, VPN tunnels, and zones—defined in ScreenOS configurations.

- If a name string includes one or more spaces, the entire string must be enclosed within double quotes ("); for example, **set address trust "local LAN" 10.1.1.0/24**.
- NetScreen trims any spaces leading or trailing text within a set of double quotes; for example, **" local LAN "** becomes **"local LAN"**.
- NetScreen treats multiple consecutive spaces as a single space.
- Name strings are case sensitive, although many CLI key words are case insensitive. For example, **"local LAN"** is different from **"local lan"**.

ScreenOS supports the following character types:

- Single-byte character sets (SBCS) and multiple-byte character sets (MBCS). Examples of SBCS are ASCII, European, and Hebrew. Examples of MBCS—also referred to as double-byte character sets (DBCS)—are Chinese, Korean, and Japanese.

  Note: *A console connection only supports SBCS. The WebUI supports both SBCS and MBCS, depending on the character sets that your Web browser supports.*

- ASCII characters from 32 (0x20 in hexidecimals) to 255 (0xff), except double quotes ("), which have special significance as an indicator of the beginning or end of a name string that includes spaces.

NETSCREEN DOCUMENTATION

To obtain technical documentation for any NetScreen product, visit www.netscreen.com/resources/manuals/.

To obtain the latest software version, visit www.netscreen.com/services/download_soft. Select a category of software product from the dropdown list, then follow the displayed instructions. (You must be a registered user to download NetScreen software.)

If you find any errors or omissions in the following content, please contact us at the e-mail address below:

techpubs@netscreen.com

1

NSRP

NetScreen Redundancy Protocol (NSRP) is a proprietary protocol that is supported on select NetScreen devices to provide high availability (HA) services. This chapter explains the components of NSRP and describes how to configure a NetScreen device for HA using NSRP. The specific topics covered are as follows:

- "NSRP Overview" on page 3
- "NSRP and NetScreen Operational Modes" on page 8
  - "Basic Active/Passive NSRP Configuration" on page 8
- "NSRP Clusters" on page 15
  - "Cluster Name" on page 17
  - "Run-Time Objects" on page 21
- "VSD Groups" on page 23
  - "Preempt Option" on page 23
  - "VSD Group Member States" on page 24
  - "Heartbeat Messages" on page 25
  - "VSIs and Static Routes" on page 28
- "Synchronization" on page 33
  - "Synchronizing Configurations" on page 33
  - "Synchronizing Files" on page 34
  - "Synchronizing RTOs" on page 34
  - "Synchronizing System Clocks" on page 37

- "Dual HA Interfaces" on page 38
  - "Control Messages" on page 39
  - "Data Messages (Packet Forwarding)" on page 40
  - "Dual HA Link Probes" on page 42
- "Setup Procedure" on page 45
  - "Cabling for a Full-Mesh Configuration" on page 45
  - "Active/Active NSRP Configuration" on page 49

NSRP OVERVIEW

To function properly as a network firewall, a NetScreen device must be placed at the single point through which all inter-zone traffic must pass.

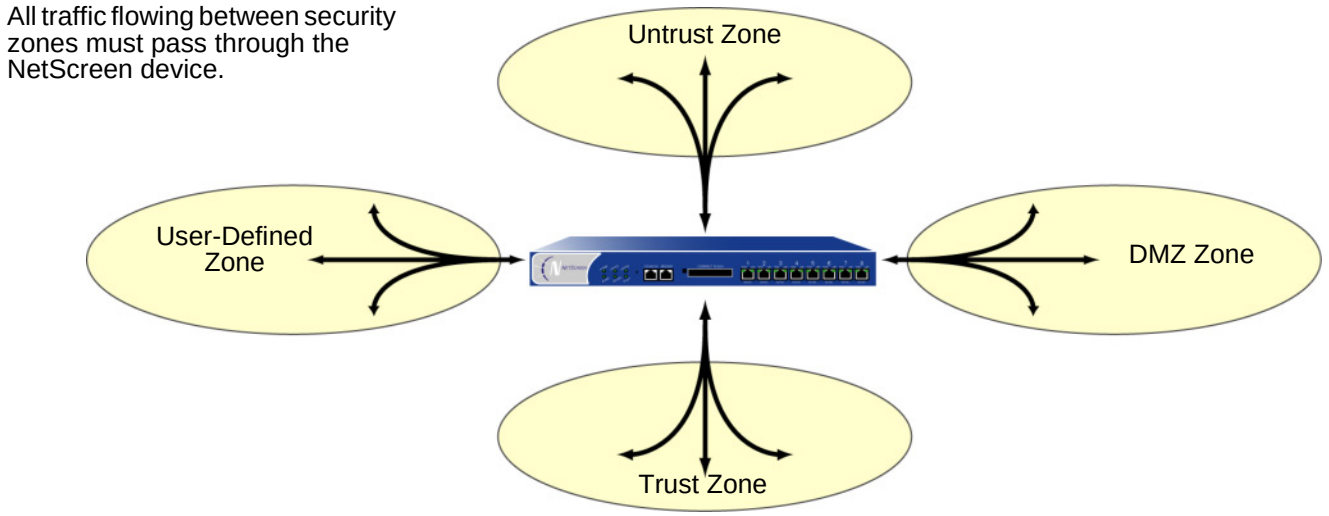

Because the NetScreen device is the single point through which all inter-zone traffic must pass, it is vital that the traffic flow remain uninterrupted, even in the event of a device or network failure.

To assure a continuous traffic flow, you can cable and configure two NetScreen devices in a redundant cluster, with one device acting as a master and the other as its backup. The master propagates all its network and configuration settings and the current session information to the backup. Should the master fail, the backup is promoted to master and takes over the traffic processing.

Note: To simplify the failover concept, only Trust and Untrust zones are shown.

Active/Passive Failover

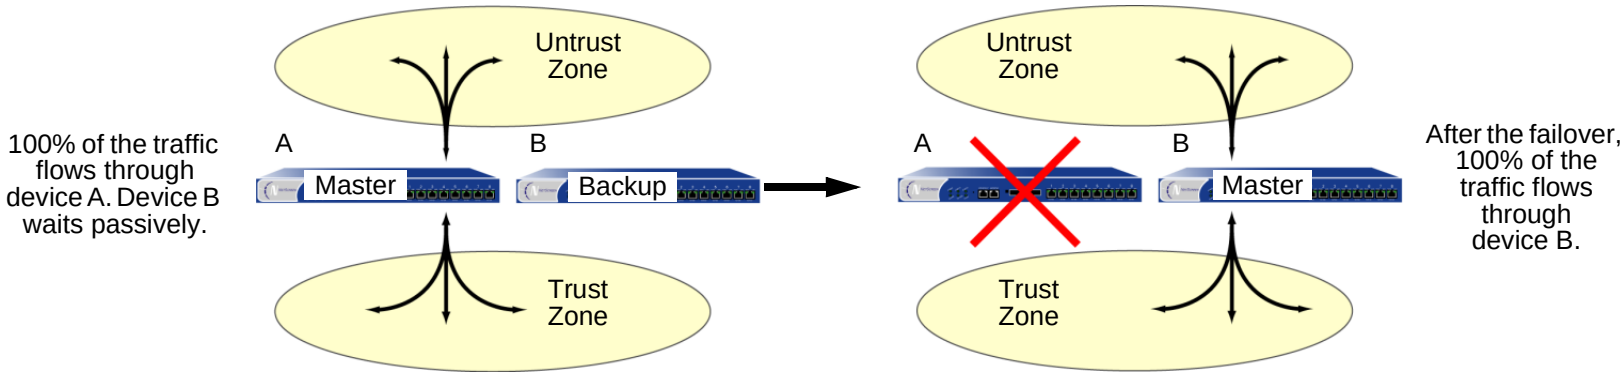

In this case, the two devices are in an active/passive configuration; that is, the master is active, handling all firewall and VPN activities, and the backup is passive[1], waiting to take over when the master steps down.

1. Although the backup is passive in the sense that it is not processing traffic, it is quite active maintaining its synchronization with the configuration settings and session information it continuously receives from the master.

With the NetScreen device in Route or NAT mode, you can configure both devices in a redundant cluster to be active, sharing the traffic distributed between them by routers with load-balancing capabilities running a protocol such as the Virtual Router Redundancy Protocol (VRRP). This is accomplished using the NetScreen Redundancy Protocol (NSRP) to create two virtual security devices (VSD) groups, each with its own virtual security interfaces (VSIs). Device A acts as the master of VSD group 1 and as the backup of VSD group 2. Device B acts as the master of VSD group 2 and as the backup of VSD group 1. This configuration is known as active/active (see illustration below). Because of device redundancy, there is no single point of failure.

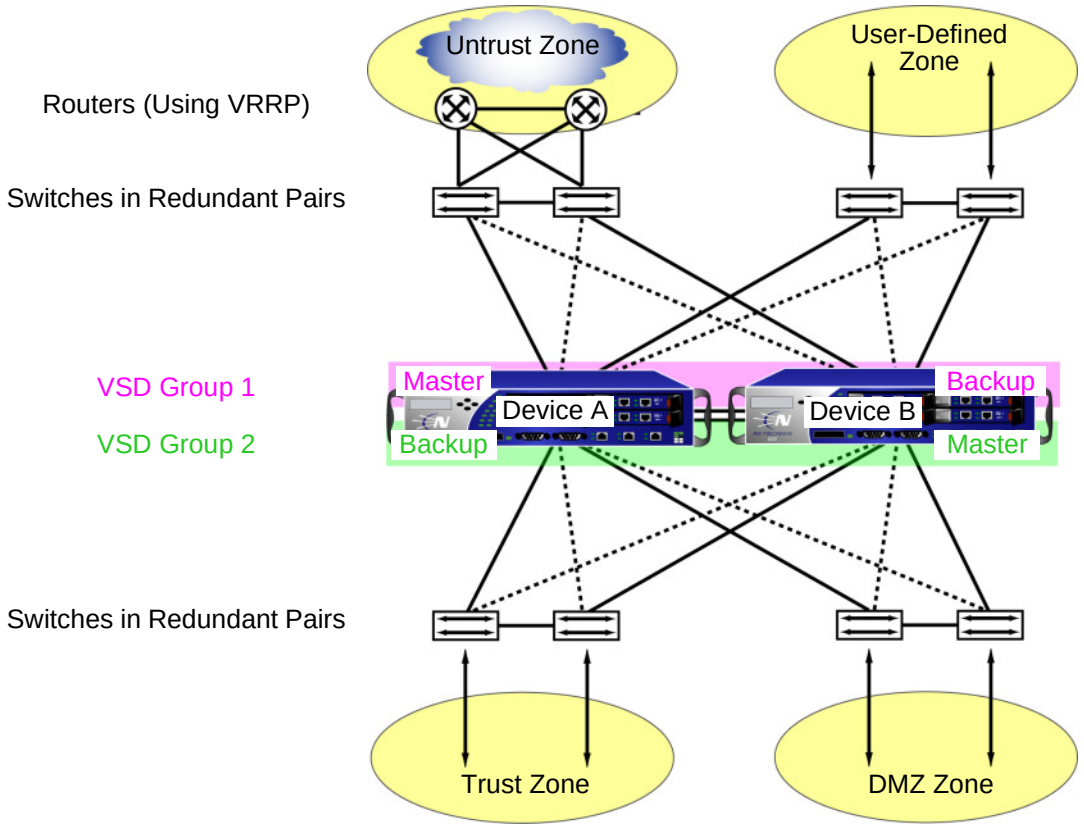

Devices A and B each receive 50% of the network and VPN traffic. Should device A fail, device B becomes the master of VSD group 1, as well as continuing to be the master of VSD group 2, and handles 100% of the traffic. Traffic redirection resulting from a failover in an active/active configuration is shown in the next illustration.

Note: To simplify the failover concept, only Trust and Untrust zones are shown.

Active/Active Failover

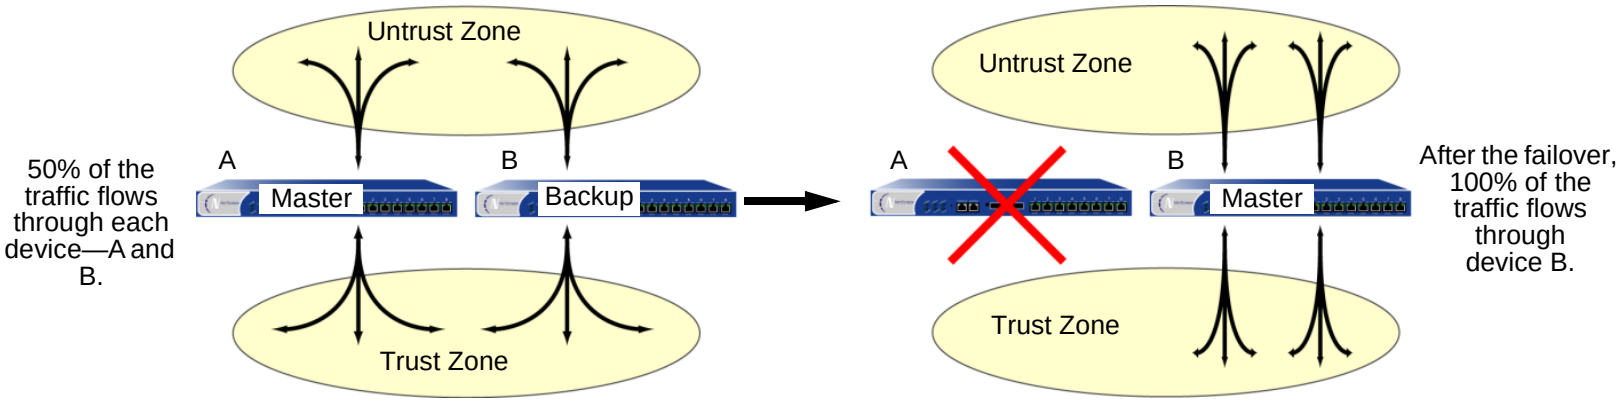

Although the total number of sessions divided between the two devices in an active/active configuration cannot exceed the capacity of a single NetScreen device (otherwise, in the case of a failover, the excess sessions might be lost[2]), the addition of a second device doubles the available bandwidth potential. A second active device also guarantees that both devices have functioning network connections.

2. Each device in an active/active configuration can tolerate traffic bursts exceeding 50% of the capacity of a single device for short periods of time; however, should a failover occur during that period, the excess traffic might be lost.

In addition to NSRP clusters, which are primarily responsible for propagating configurations among group members and advertising each member's current VSD group states, you can configure devices A and B as members in an RTO mirror group, which is responsible for maintaining the synchronicity of run-time objects (RTOs)[3] between a pair of devices. When the master steps down, the backup can immediately assume mastership with minimal service downtime by maintaining all current sessions.

Because of the sensitive nature of NSRP communications, you can secure all NSRP traffic through encryption and authentication. For encryption and authentication, NSRP supports the DES and MD5 algorithms respectively. (For more information about these algorithms, see "Protocols" on page **5**-7.)

Note: *If the HA cables run directly from one NetScreen device to another (that is, not through a switch forwarding other kinds of network traffic), it is unnecessary to use encryption and authentication.*

If you want to use Simple Network Management Protocol (SNMP) to monitor the NetScreen device, private NSRP MIBs are available for download at www.netscreen.com/services/download_soft. (For more information about SNMP, see "SNMP" on page **3**-91.)

NSRP consists of two basic elements, which are fully explained in the following sections:

- "NSRP Clusters" on page 15
- "VSD Groups" on page 23

For an example of a basic active/passive NSRP configuration, see "Example: NSRP for an Active/Passive Configuration" on page 10. For an example of an active/active NSRP configuration, see "Example: NSRP for an Active/Active Configuration" on page 49.

3. RTOs are objects created dynamically in the NetScreen device memory during the normal operation of the device. RTOs allow the device to understand the network around it and enforce its policies. Examples of RTOs are TCP/UDP sessions, IPSec Phase 2 security associations (SAs), DHCP allocations, RSA and DSS key pairs, ARP tables, and DNS caches.

NSRP and NetScreen Operational Modes

NetScreen device interfaces can run in one of three modes: NAT, Route, and Transparent mode. When interfaces are in NAT or Route mode, the NetScreen device operates at Layer 3 in the OSI model. The security zone interfaces have IP addresses, and the NetScreen device forwards traffic like a Layer 3 router. When interfaces are in Transparent mode, the NetScreen device operates at Layer 2. The security zone interfaces do not have IP addresses, and the NetScreen device forwards traffic like a Layer 2 switch.

When a NetScreen device is operating at Layer 3 (NAT or Route mode), it can be in an active/active or active/passive NSRP configuration. To manage a backup device, you must use the manage IP address that you set per security zone interface[4].

When a NetScreen device is operating at Layer 2 (Transparent mode), it can only be in an active/passive NSRP configuration. To manage a backup device, you use the manage IP address that you set on the VLAN1 interface.

Basic Active/Passive NSRP Configuration

Performing the most basic active/passive NSRP configuration is quite easy. You can put a device in an NSRP cluster and VSD group with a single CLI command—**set nsrp cluster id** *number*—or in the WebUI by typing a single number for the NSRP cluster ID.

You can enable automatic RTO synchronization with the CLI command **set nsrp rto sync all**, or in the WebUI by selecting the **NSRP RTO Synchronization** option on the Network > NSRP > Synchronization page and then clicking **Apply**.

Next, you must also select the ports that you want the devices to monitor, so that if they detect a loss of network connectivity on one of the monitored ports, the device fails over.

Note: *Before NSRP can function, you must first cable two NetScreen devices together as explained in* *"Cabling for a Full-Mesh Configuration" on page 45**. Also, if you want to maintain network connectivity for administrative traffic to one or more physical interfaces on a NetScreen device in an NSRP cluster, first set the manage IP address for those interfaces as explained in "Manage IP" on page* ***3****-34 before you enable NSRP.*

4. You cannot set a manage IP address on a VSI for any VSD group except VSD group 0.

Default Settings

The basic NSRP configuration uses the following default settings:

- VSD Group Information
  - VSD group ID: 0
  - Device priority in the VSD group: 100
  - Preempt option: disabled
  - Preempt hold-down time: 0 seconds
  - Initial state hold-down time: 5 seconds
  - Heartbeat interval: 1000 milliseconds
  - Lost heartbeat threshold: 3
- RTO Mirror Information
  - RTO synchronization: disabled
  - Heartbeat interval: 4 seconds
  - Lost heartbeat threshold: 16
- NSRP Link Information
  - Number of gratuitous ARPs: 4
  - NSRP encryption: disabled
  - NSRP authentication: disabled
  - Interfaces monitored: none
  - Secondary path: none

When you set a NetScreen device in an NSRP cluster, the NetScreen device automatically creates VSD group 0 and transforms physical interfaces into Virtual Security Interfaces (VSIs) for VSD group 0[5].

5. The convention for indicating a VSI is **<interface_name>:<VSD_group_ID>**. For example, the following indicates that the redundant interface *red1* is a VSI for VSD group 1: **red1:1**. However, if the VSD group ID is 0, no VSD group ID is specified. For example, if the redundant interface *red2* is a VSI for VSD group 0, it appears simply as **red2**.

Example: NSRP for an Active/Passive Configuration

In the following example, you cable ethernet7 on NetScreen-A to ethernet7 on NetScreen-B. You cable the ethernet8 interfaces likewise. Then you bind ethernet7 and ethernet8 to the HA zone[6]. You set manage IP addresses for the Trust zone interfaces on both devices—10.1.1.20 for NetScreen-A and 10.1.1.21 for NetScreen-B. You then assign each device to NSRP cluster ID 1. When the devices become members of the NSRP cluster, the IP addresses of their physical interfaces automatically become the IP addresses of the Virtual Security Interfaces (VSIs) for VSD group ID 0. Each VSD member has a default priority of 100, the device with the higher unit ID becomes the VSD group master.

You configure the devices to monitor ports ethernet1 and ethernet3, so that loss of network connectivity on either of those ports triggers a device failover. You also enable the automatic synchronization of RTOs.

Note: *This is an overly simplistic example and is included to illustrate the basic elements of NSRP configuration. For a more fully developed configuration, see* *"Example: NSRP for an Active/Active Configuration" on page 49**.*

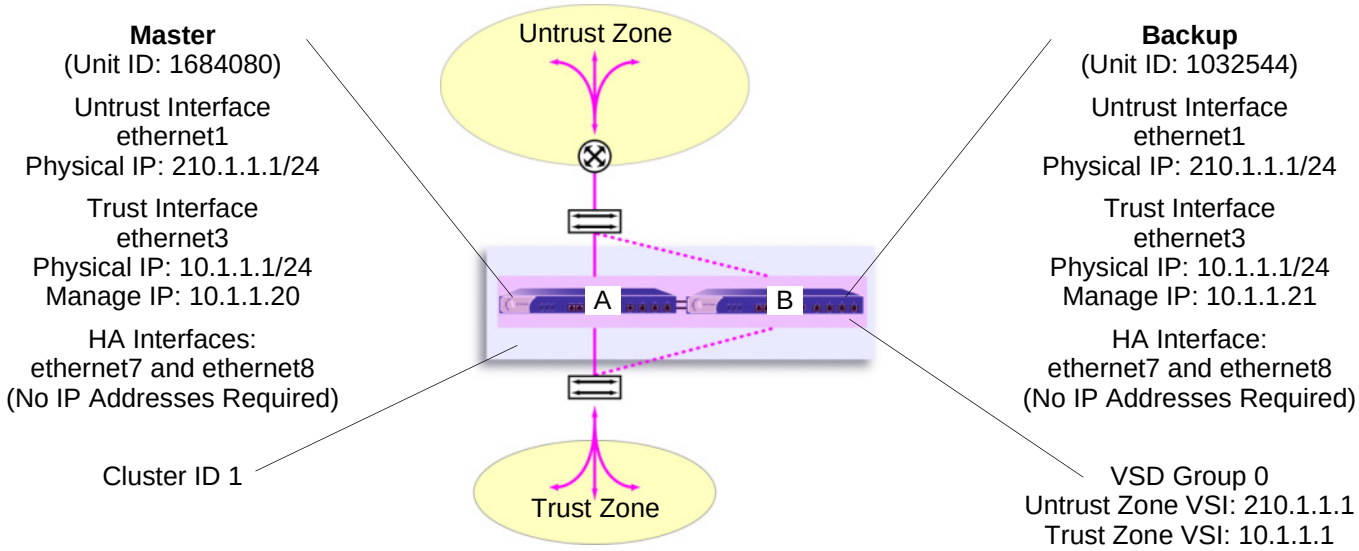

6. By default, ethernet8 is bound to the HA zone. Binding it to the HA zone is only necessary if you have previously bound it to a different zone.

WebUI (NetScreen-A)

1. Interfaces

Network > Interfaces > Edit (for ethernet7): Enter the following, and then click **OK**:

Zone Name: HA

Network > Interfaces > Edit (for ethernet8): Enter the following, and then click **OK**:

Zone Name: HA

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 210.1.1.1/24

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Manage IP: 10.1.1.20

Enter the following, and then click **OK**:

Interface Mode: NAT

2. NSRP

Network > NSRP > Monitor > Interface > VSD ID: Device Edit Interface: Select **ethernet1** and **ethernet3**, and then click **Apply**.

Network > NSRP > Synchronization: Select **NSRP RTO Synchronization**, and then click **Apply**[7].

Network > NSRP > Cluster: In the Cluster ID field enter **1**, and then click **Apply**.

7. If you do not enable the automatic RTO synchronization option, you can manually synchronize RTOs with the CLI command **exec nsrp sync rto all**.

WebUI (NetScreen-B)

3. Interfaces

Network > Interfaces > Edit (for ethernet7): Enter the following, and then click **OK**:

Zone Name: HA

Network > Interfaces > Edit (for ethernet8): Enter the following, and then click **OK**:

Zone Name: HA

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 210.1.1.1/24

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **Apply**:

Zone Name: Trust

Static IP: (select this option when present)

IP Address/Netmask: 10.1.1.1/24

Manage IP: 10.1.1.21

Enter the following, and then click **OK**:

Interface Mode: NAT

4. NSRP

Network > NSRP > Monitor > Interface > VSD ID: Device Edit Interface: Select **ethernet1** and **ethernet3**, and then click **Apply**.

Network > NSRP > Synchronization: Select **NSRP RTO Synchronization**, and then click **Apply**.

Network > NSRP > Cluster: In the Cluster ID field enter **1**, and then click **Apply**.

CLI (NetScreen-A)

1. Interfaces

```
set interface ethernet7 zone ha
set interface ethernet8 zone ha
set interface ethernet1 zone untrust
set interface ethernet1 ip 210.1.1.1/24
set interface ethernet3 zone trust
set interface ethernet3 ip 10.1.1.1/24
set interface ethernet3 manage-ip 10.1.1.20
set interface ethernet3 nat
```

2. NSRP

```
set nsrp rto-mirror sync[8]
set nsrp monitor interface ethernet1
set nsrp monitor interface ethernet3
set nsrp cluster id 1
save
```

CLI (NetScreen-B)

3. Interfaces

```
set interface ethernet7 zone ha
set interface ethernet8 zone ha
set interface ethernet1 zone untrust
set interface ethernet1 ip 210.1.1.1/24
set interface ethernet3 zone trust
set interface ethernet3 ip 10.1.1.1/24
set interface ethernet3 manage-ip 10.1.1.21
set interface ethernet3 nat
```

8. If you do not enable the automatic RTO synchronization option, you can manually synchronize RTOs with the CLI command **exec nsrp sync rto all**.

4. NSRP

```
set nsrp rto-mirror sync
set nsrp monitor interface ethernet1
set nsrp monitor interface ethernet3
set nsrp cluster id 1
save
```

Note: *After performing this configuration, type the* ***get nsrp*** *command to check the default NSRP settings that the device automatically creates, and which are noted on page 8.*

NSRP CLUSTERS

An NSRP cluster consists of a group of NetScreen devices that enforce the same overall security policy and share the same configuration settings. When you assign a NetScreen device to an NSRP cluster, any changes made to the configuration on one member of the cluster propagate to the other. Members of the same NSRP cluster maintain identical settings for the following:

- Policies and policy objects (such as addresses, services, VPNs, users, and schedules)
- System parameters (such as settings for authentication servers, DNS, SNMP, syslog, URL blocking, firewall detection options, and so forth)

Members of a cluster do not propagate the following configuration settings:

Non-Propagating Commands	
NSRP	• set/unset nsrp cluster id *number* • set/unset nsrp auth password *pswd_str* • set/unset nsrp encrypt password *pswd_str* • set/unset nsrp monitor interface *interface* • set/unset nsrp vsd-group id *id_num* { mode *string* | preempt | priority *number* } • set/unset nsrp rto-mirror …
Interface	• set/unset interface *interface* manage-ip *ip_addr* • set/unset interface *interface* phy … • set/unset interface *interface* bandwidth *number* • set/unset interface redundant *number* phy primary *interface* • All commands pertaining to local interfaces
IP Tracking	• All IP tracking commands (set/unset nsrp track-ip …)
Console Settings	• All console commands (set/unset console …)
Hostname	• set/unset hostname *name_str*

Non-Propagating Commands	
SNMP	• set/unset snmp name *name_str*
Virtual Router	• set/unset vrouter *name_str* router-id *ip_addr*
Clear*	• All clear commands (clear admin, clear dhcp, …)
Debug†	• All debug commands (debug alarm, debug arp, …)

* By default, NSRP cluster members do not propagate the **clear** commands. To propagate a clear command to all devices in an NSRP cluster, insert the keyword **cluster** into the command. For example, **clear cluster admin …**, **clear cluster dhcp …**

† By default, NSRP cluster members do not propagate the **debug** commands. To propagate a debug command to all devices in an NSRP cluster, insert the keyword **cluster** into the **debug** command. For example, **debug cluster alarm …**, **debug cluster arp …**

Before two NetScreen devices can provide redundant network connectivity, you must group them in the same NSRP cluster by assigning a cluster ID[9] between 1 and 7. When a NetScreen device becomes a member of a cluster, it automatically becomes a member of VSD group 0, and all interfaces become VSIs for VSD group 0. If you want to retain some interfaces as local interfaces and create VSIs from select interfaces, you must do the following:

1. Remove VSD group 0.

   All the interfaces on all cluster members become local interfaces.
2. Create another VSD group, such as VSD group 1.
3. Create VSIs for that VSD group.

For more information about VSD groups, see "VSD Groups" on page 23.

Cluster members can also synchronize run-time objects (RTOs), which allows a newly elected VSD group master to maintain uninterrupted network and VPN services after a failover. (For more information about RTOs, see "Run-Time Objects" on page 21.)

9. Assigning an ID of 0 removes a device from a cluster.

Cluster Name

Because NSRP cluster members can have different host names, a failover can disrupt SNMP communication and the validity of digital certificates because SNMP communication and certificates rely on the host name of a device to function properly.

To define a single name for all cluster members, type the following CLI command:

```
set nsrp cluster name name_str
```

Use the cluster name when configuring the SNMP host name for the NetScreen device (**set snmp name** *name_str*) and when defining the common name in a PKCS10 certificate request file.

The use of a single name for all cluster members allows SNMP communication and digital certificate use to continue without interruption after a device failover.

Example: Creating an NSRP Cluster

In this example, you group devices A and B within NSRP cluster ID 1 with cluster name “cluster1”. You also specify the following settings on each device:

NSRP communication security: Assign passwords—725dCaIgDL and WiJoaw4177—for creating authentication and encryption keys to secure NSRP communications.

After you have grouped both devices in the same cluster and given them the same authentication and encryption passwords, you can enter the following settings on either device A or B. (Most settings entered on one device in a cluster propagate to the other device. For a list on non-propagating commands, see “Non-Propagating Commands” on page 15.)

- **Interface monitoring:** Select the ethernet1 (bound to the Untrust zone) and ethernet2 (bound to the Trust zone) for monitoring layer 2 network connectivity.
- **Secondary link:** Specify that the ethernet2 interface carry VSD heartbeats should both HA1 and HA2 links go down. The purpose of this feature is to prevent multiple VSD group masters when both HA links fail.
- **Gratuitous ARP broadcasting:** Specify the number of ARP broadcasts as 5 (the default is 4). ARP broadcasts notify surrounding network devices of the MAC address of a new master after a failover has occurred.

(All the interfaces on these devices become VSIs for VSD group 0. In the “VSD Groups” section, you create a second VSD group for these devices. See “Example: Creating Two VSD Groups” on page 26.)

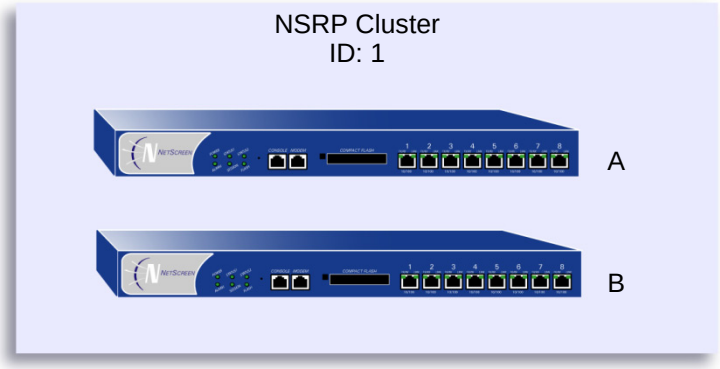

WebUI (NetScreen-A)

1. NSRP Cluster and Communication Security

Network > NSRP > Cluster: Enter the following[10], and then click **Apply**:

Cluster ID: 1

NSRP Authentication Password: (select) 725dCaIgDL

NSRP Encryption Password: (select) WiJoaw4177

WebUI (NetScreen-B)

2. NSRP Cluster and Communication Security

Network > NSRP > Cluster: Enter the following, and then click **Apply**:

Cluster ID: 1

Number of Gratuitous ARPs to Resend: 5

NSRP Authentication Password: (select) 725dCaIgDL

NSRP Encryption Password: (select) WiJoaw4177

3. NSRP Settings

Network > NSRP > Monitor > Interface > VSD ID: Device Edit Interface: Select **ethernet1** and **ethernet2**, and then click **Apply**.

Network > NSRP > Link: Select **ethernet2** from the Secondary Link drop-down list, and then click **Apply**.

10. You can only set a cluster name through the CLI.

CLI (NetScreen-A)

1. NSRP Cluster and Communication Security

```
set nsrp cluster id 1
set nsrp auth password 725dCaIgDL
set nsrp encrypt password WiJoaw4177
save
```

CLI (NetScreen-B)

2. NSRP Cluster and Communication Security

```
set nsrp cluster id 1
set nsrp auth password 725dCaIgDL
set nsrp encrypt password WiJoaw4177
save
```

3. NSRP Settings

```
set nsrp cluster name cluster1
set nsrp monitor interface ethernet1
set nsrp monitor interface ethernet2
set nsrp secondary-path ethernet2
set nsrp arp 5
save
```

Run-Time Objects

Run-time objects (RTOs) are code objects created dynamically in memory during normal operation. Some examples of RTOs are session table entries, ARP cache entries, DHCP leases, and IPSec security associations (SAs). In the event of a failover, it is critical that the current RTOs be maintained by the new master to avoid service interruption[11]. To accomplish this, RTOs are backed up by the members of an NSRP cluster. Working together, each member backs up the RTOs from the other, which allows RTOs to be maintained should the master of either VSD group in an active/active HA scheme step down.

In the current ScreenOS release, you do not have to configure one or more RTO mirror groups to synchronize RTOs among members of an NSRP cluster. Defining a NetScreen device as a member of a cluster and specifying RTO synchronization automatically enables the local device to send and receive RTOs.

By default, NSRP cluster members do not synchronize RTOs. Before enabling RTO synchronization, you must first synchronize the configurations between the cluster members. Unless the configurations on both members in the cluster are identical, RTO synchronization might fail. (For examples of the synchronization procedure, see "Example: Adding a Device to an Active NSRP Cluster" on page 36 and "Example: NSRP for an Active/Active Configuration" on page 49.)

To enable RTO synchronization, do either of the following:

WebUI

Network > NSRP > Synchronization: Select the **NSRP RTO Synchronization** check box, and then click **Apply**.

CLI

```
set nsrp rto-mirror sync
save
```

11. Using policies, you can specify which sessions to backup and which not to backup. For traffic whose sessions you do not want backed up, apply a policy with the HA session backup option disabled. In the WebUI, clear the **HA Session Backup** check box. In the CLI, use the **no-session-backup** argument in the **set policy** command. By default, the backing up of sessions is enabled.

RTO Mirror States

The procedure for two NSRP cluster members to initiate their RTO mirror relationship develops through two operational states—set and active. The devices progress through these states as follows:

1. After you add the first device to a group, its state is set. In the set state, the device waits for its peer to join the group. As the receiver of RTOs, it periodically transmits an r-ready message (receiver-ready), announcing its own availability. As the sender of RTOs, it waits until it gets an r-ready message from a device with the same cluster ID.
2. After you add the peer and the two devices are correctly cabled for HA (see "Cabling for a Full-Mesh Configuration" on page 45), then the following occurs:
   a. The receiver sends an r-ready message.
   b. The sender gets the r-ready message, and immediately sends a group-active message to inform its peer that its state is now active.
   c. The receiver then changes its state to active as well.

In addition to passing RTOs from sender to receiver, both active mirrors send RTO heartbeats at user-defined intervals to communicate their operational status. To define the interval, use the following CLI command: **set nsrp rto-mirror hb-interval** *number*.

If a device does not receive a specified number of consecutive heartbeats from its peer, it changes its state from active to set. To define the lost heartbeat threshold required to impel a state changeover, use the following CLI command: **set nsrp rto-mirror hb-threshold** *number*.

Note: *To maintain identical RTO heartbeat settings, the* ***set nsrp rto-mirror hb-interval*** *number and* ***set nsrp rto-mirror hb-threshold*** *number are propagated.*

You can use the following command to disable RTO session synchronization on the device acting as sender in an NSRP cluster: **set nsrp rto-mirror session off**. Issuing this command on a device only disables session synchronization from that device to others in the cluster.

VSD GROUPS

A Virtual Security Device (VSD) group is a pair of physical NetScreen devices that collectively comprise a single VSD. One physical device acts as the master of the VSD group. The virtual security interface (VSI) of the VSD is bound to the physical interface of the master. The other physical device acts as the backup[12]. If the master device fails, the VSD fails over to the backup and the VSI binding is transferred to the physical interface on the backup, which is instantly promoted to master.

By grouping two NetScreen devices into two VSD groups, with each physical device being the master in one group and the backup in the other, both devices can actively process traffic as masters while backing up each other in the event of a failover.

Upon initial NSRP configuration, the VSD group member with the priority number closest to 0 becomes the master. (The default is 100.) If two devices have the same priority value, the device with the lowest MAC address becomes master.

Preempt Option

You can determine whether a better priority number (closer to zero) can initiate a failover by setting the device that you want to be master in preempt mode. If you enable the preempt option on that device, it becomes the master of the VSD group if the current master has a lesser priority number (farther from zero). If you disable this option, a master with a lesser priority than a backup can keep its position (unless some other factor, such as an internal problem or faulty network connectivity, causes a failover).

Using the hold-down time to delay a failover can prevent a flurry of rapid failovers in the event of port-flickering on an adjacent switch and also ensure that surrounding network devices have sufficient time to negotiate new links before the new master becomes available. To enable or disable the preempt option, use the following CLI command:

```
set/unset nsrp vsd-group id number preempt
```

You can use the following CLI command to set the hold-down time—used for delaying the preempted failover—to any length from 0 to 600 seconds:

```
set nsrp vsd-group id number preempt hold-down number
```

12. In the current release, a VSD group can have two members. In later releases, there might be more than two members, in which case, one device acts as a master, another as a primary backup, and the remaining VSD group members as backups.

VSD Group Member States

The members of a VSD group can be in one of six states:

- Master – The state of a VSD group member that processes traffic sent to the VSI.
- Primary Backup – The state of a VSD group member that becomes the master should the current master step down. The election process uses device priorities to determine which member to promote. Note that when electing a new master, an RTO peer has precedence over any other VSD group member, even if that member has a better priority rating.
- Backup – The state of a VSD group member that monitors the status of the primary backup and elects one of the backup devices to primary backup if the current one steps down.
- Initial – The transient state of a VSD group member while it joins a VSD group, either when the device boots up or when it is added via the **set nsrp vsd-group id** *id_num* command.

  You can specify how long a VSD group member stays in the initial state with the **set nsrp vsd-group init-hold** *number* command. The default (and minimum) setting is 5. To determine the initial state hold-down time, multiply init-hold value by the VSD heartbeat-interval (init-hold x hb-interval = initial state hold-down time). For example, if the init-hold is 5 and the hb-interval is 1000 milliseconds, then the initial state hold-down time is 5,000 milliseconds, or 5 seconds (5 x 1000 = 5000).

  Note: *If you reduce the VSD heartbeat interval, you should increase the init-hold value. For information on configuring the heartbeat interval, see* *"Heartbeat Messages" on page 25**.*

- Ineligible – The state that an administrator purposefully assigns to a VSD group member so that it cannot participate in the election process. To do this, use the **set nsrp vsd-group id** *id_num* **mode ineligible** command.
- Inoperable – The state of a VSD group member after a system check determines that the device has an internal problem (such as no processing boards) or a network connection problem (such as when an interface link fails).

  Note: *When the device returns from either the ineligible state (when you use the* ***exec nsrp vsd-group id*** *id_num* ***mode { backup | init | master | pb }*** *command) or inoperable state (when the system or network problem has been corrected), it must first pass through the initial state.*

You can determine the state of a device by observing the HA LED. The meanings of the various colors—dark, green, yellow, red—are as follows:

- Dark: The device is not enabled for NSRP.
- Green: The device is enabled for NSRP; it is the master in one or more VSD groups; and it is not in inoperable mode.
- Yellow: The device is enabled for NSRP; it is not the master in any VSD group; and it is not in inoperable mode.
- Red: The device is enabled for NSRP, but it is currently in inoperable mode.

Heartbeat Messages

Every VSD group member—even if it is in the initial, ineligible, or inoperable state—communicates with its group members by sending a heartbeat message every second[13]. These messages allow every member to know the current state of every other member. The heartbeat message includes the following information:

- Unit ID of the device
- VSD group ID
- VSD group member status (master, primary backup, or backup)
- Device priority
- RTO peer information

The interval for sending VSD heartbeats is configurable (200, 600, 800, or 1000 milliseconds; 1000 ms is the default). The CLI command—which applies globally to all VSD groups—is **set nsrp vsd-group hb-interval** *number*. You can also configure the lost heartbeat threshold that is used to determine when a VSD group member is considered as missing. The CLI command, which also applies globally to all VSD groups, is **set nsrp vsd hb-threshold** *number*. The minimum value for the lost heartbeat threshold is 3.

The heartbeat messages are sent over the HA1 link. For more information about the HA1 and HA2 interfaces and the kinds of messages communicated over each, see “Dual HA Interfaces” on page 38.

13. If a device is in the inoperable state with all HA links down, it can neither send nor receive VSD heartbeat messages unless you have configured a secondary path for these messages. For more information about configuring a secondary path, see “Example: Creating an NSRP Cluster” on page 18.

Example: Creating Two VSD Groups

This example continues with the configuration of devices A and B, which are already members of the same NSRP cluster and VSD group 0 (see “Example: Creating an NSRP Cluster” on page 18).

In this example, you create a second VSD group—Group 1. You assign device A priority 1 in Group 0 and the default priority (100) in Group 1. You assign device B priority 1 in Group 1 and the default priority (100) in Group 0. In both VSD groups, you enable the preempt option on the master and set the preempt hold-down time to 10 seconds. If both devices are active, device A is always the master of Group 1 and B the master of Group 2.

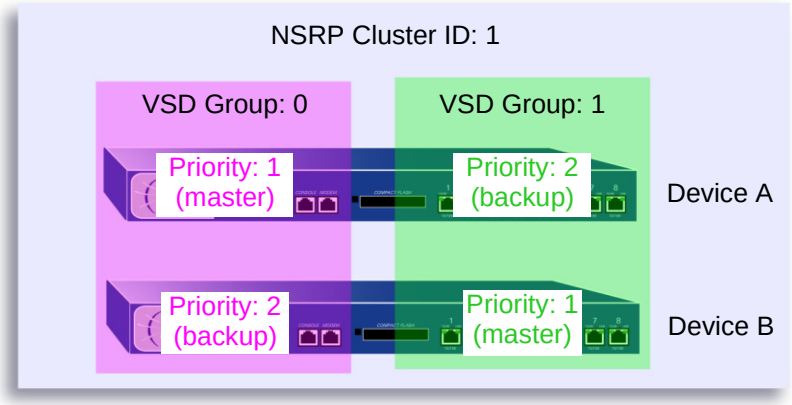

WebUI

1. Device A

   Network > NSRP > VSD Group > Edit (for VSD group 0): Enter the following, and then click **OK**:

   Priority: 1

   Enable Preempt: (select)

   Preempt Hold-Down Time (sec): 10

   Network > NSRP > VSD Group > New: In the Group ID field, type **1**, and then click **OK**.

2. Device B

Network > NSRP > VSD Group > Edit (for VSD group 1): Enter the following, and then click **OK**:

Priority: 1

Enable Preempt: (select)

Preempt Hold-Down Time (sec): 10

CLI

3. Device A

```
set nsrp vsd-group id 0 priority 1
set nsrp vsd-group id 0 preempt hold-down 10
set nsrp vsd-group id 0 preempt
set nsrp vsd-group id 1
save
```

4. Device B

```
set nsrp vsd-group id 1 priority 1
set nsrp vsd-group id 1 preempt hold-down 10
set nsrp vsd-group id 1 preempt
save
```

VSIs and Static Routes

After you create a VSD group, you must bind Virtual Security Interfaces (VSIs) to the VSD. When you put a NetScreen device in an NSRP cluster, all the security zone interfaces become VSIs of VSD group 0. You must manually assign VSIs to VSDs with other IDs for each security zone configured on the NetScreen device.

By default, the NetScreen device adds an entry to its routing table for the immediate subnet of a VSI. For static routes to addresses beyond the immediate subnet, you must manually make route table entries for each VSI through which you want the NetScreen device to forward traffic to those addresses. For example, if you have two VSDs and you want to configure a default route to a router in the Untrust zone, you must make a routing table entry for the Untrust zone VSI of both VSDs. If you set the default route on only one VSD (for example, VSD 0), the NetScreen device acting as the master of the other VSD (for example, VSD 1) must pass all outbound traffic sent to it across the HA data link to the device acting as the master of VSD 0.

If the default route is set only on VSD 0, NetScreen-B, as the master of VSD 1, must forward outbound traffic received on its Trust zone VSI across the HA data link to NetScreen-A.

NetScreen-A sends it out its Untrust zone VSI to the external router.

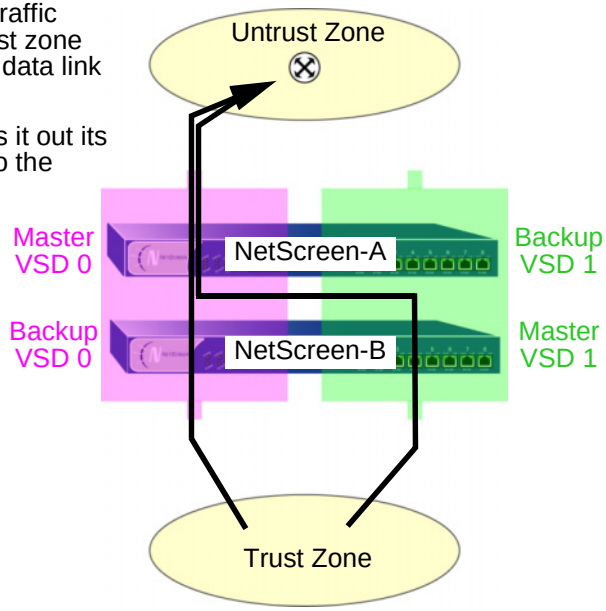

If the default route is set on both VSD 0 and 1, both NetScreen devices forward outbound traffic received on their Trust zone VSIs out their own Untrust zone VSIs to the external router.

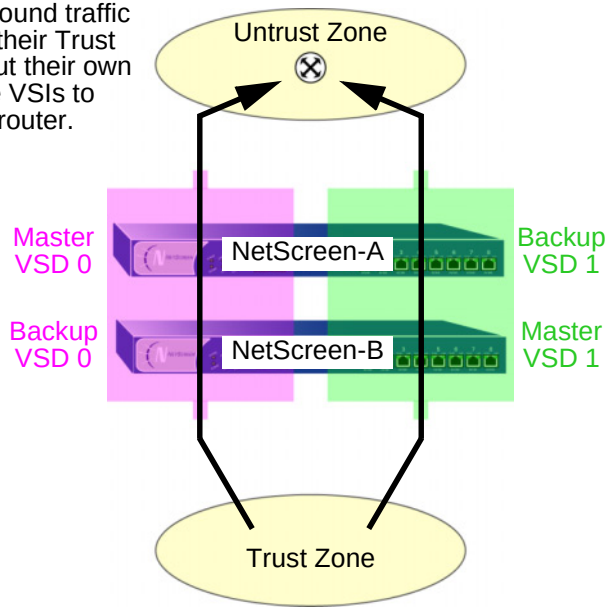

Example: Trust and Untrust Zone VSIs

This example builds on the previous example, “Example: Creating Two VSD Groups” on page 26 and assumes that you have already done the following on devices A and B:

- Put both devices in NSRP cluster 1
- Created VSD group 1 (the NetScreen device created VSD group 0 automatically when you put the device in NSRP cluster 1)

You bind ethernet1 to the Untrust zone and assign it IP address 210.1.1.1/24. You bind ethernet3 to the Trust zone, put it in NAT mode, and assign it IP address 10.1.1.1/24. You define 10.1.1.21 as the manage IP on ethernet3 for device A, and 10.1.1.22 as the manage IP on ethernet3 for device B. Then you create the following VSIs for VSD group 1:

- Untrust zone VSI ethernet1:1 (210.1.1.2/24)
- Trust zone VSI ethernet3:1 (10.1.1.2/24)

The NetScreen device creates VSIs for VSD group 0 automatically, using the IP addresses already assigned to the local interfaces at the time you put the device in an NSRP cluster. In this example, the VSD group 0 Untrust zone VSI is ethernet1[14] with IP address 210.1.1.1/24. The VSD group 0 Trust zone VSI is ethernet3 with IP address 10.1.1.1/24.

Finally, you set two default routes to the external router in the Untrust zone at 210.1.1.250—one route for the Untrust zone VSI on VSD 0 and another for the Untrust zone VSI on VSD 1. All security zones are in the trust-vr routing domain.

14. The VSD group ID “0” does not appear in the names of VSIs in VSD 0. Instead of *ethernet1:0,* the VSI is identified simply as *ethernet1.*

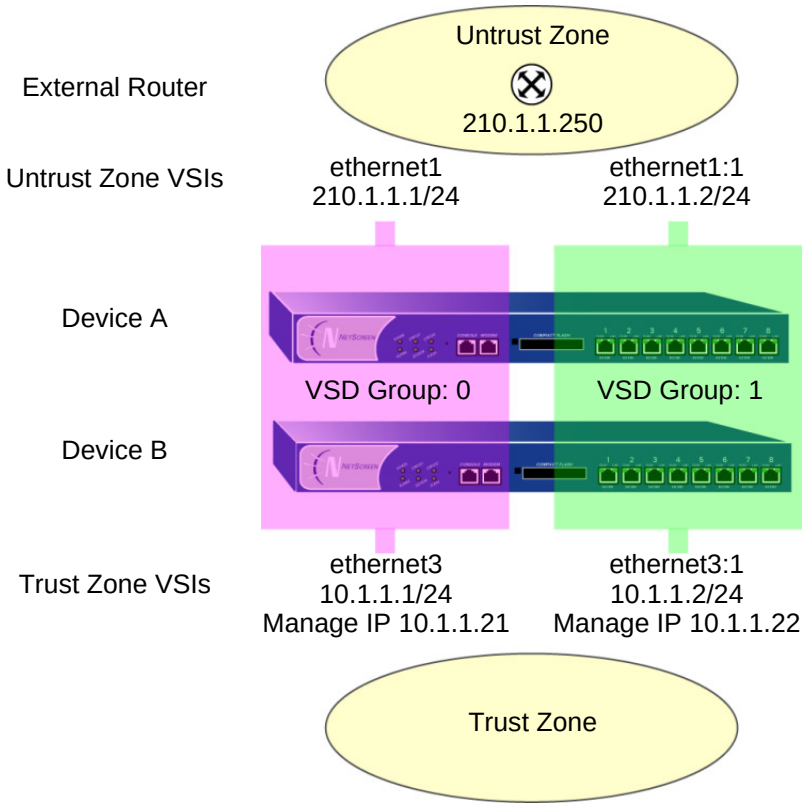

WebUI (Device A)

1. Interfaces (VSIs for VSD Group 0)

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **Apply**:

Zone Name: Trust

IP Address/Netmask: 10.1.1.1/24

Manage IP: 10.1.1.21

Enter the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select this option when present)

IP Address/Netmask: 210.1.1.1/24

WebUI (Device B)

2. Manage IP Address

Network > Interfaces > Edit (for ethernet3): Enter **10.1.1.22** in the Manage IP field, and then click **Apply**.

3. VSIs for VSD Group 1

Network > Interfaces > New VSI IF: Enter the following, and then click **OK**:

Interface Name: VSI Base: ethernet1

VSD Group: 1

IP Address / Netmask: 210.1.1.2/24

Network > Interfaces > New VSI IF: Enter the following, and then click **OK**:

Interface Name: VSI Base: ethernet3

VSD Group: 1

IP Address / Netmask: 10.1.1.2/24

4. Routes

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address: 0.0.0.0

Netmask: 0.0.0.0

Gateway: (select)

Interface: ethernet1:1

Gateway IP Address: 210.1.1.250

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address: 0.0.0.0

Netmask: 0.0.0.0

Gateway: (select)

Interface: ethernet1:2

Gateway IP Address: 210.1.1.250

CLI (Device A)

1. Interfaces

```
set interface ethernet3 zone trust
set interface ethernet3 ip 10.1.1.1/24
set interface ethernet3 manage-ip 10.1.1.21
set interface ethernet3 nat

set interface ethernet1 zone untrust
set interface ethernet1 ip 210.1.1.1/24
```

CLI (Device B)

2. Manage IP Address

```
set interface ethernet3 manage-ip 10.1.1.22
```

3. Virtual Security Interfaces

```
set interface ethernet1:1 ip 210.1.1.2/24
set interface ethernet3:1 ip 10.1.1.1.2/24
```

4. Routes

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet1 gateway 210.1.1.250
set vrouter trust-vr route 0.0.0.0/0 interface ethernet1:1 gateway 210.1.1.250
save
```

SYNCHRONIZATION

When you add a new device to an active NSRP cluster, you must synchronize the configuration and files (such as PKI public/private key files) from the master of the VSD group or groups to the new device. After the configurations and files are synchronized, you must then synchronize the run-time objects (RTOs). You must also synchronize configurations, files, and RTOs after a member of a cluster becomes unsynchronized for any reason.

Synchronizing Configurations

If you make any configuration changes on one device while another in the cluster reboots (or if all HA links fail), it is possible that the configuration settings can become unsynchronized. To discover if the configuration of one device is out of sync with that of another, use the **exec nsrp sync global-config check-sum** command. The output states whether the configurations of the two devices are in or out of sync and provides the checksums of the local and remote devices.

If the configurations are out of sync, use the following command to resynchronize them: **exec nsrp sync global-config save** (and then reboot the device) or **exec nsrp sync global-config run** (which does not require rebooting the device). If you do not use the **unset all** command on the local device before synchronizing the configurations, the local device appends the settings from the remote device to its existing settings. However, after synchronizing the configurations, every duplicate setting produces an error message. To avoid generating error messages while synchronizing configurations, you can do the following:

1. Download both the local and remote configurations to a workstation.
2. Use an application such as WinDiff to discern the differences between the files.
3. Manually enter the settings on the local device that had been added, modified, or deleted on the remote device.

Note: *Because NetScreen devices use the NetScreen Reliable Transport Protocol (NRTP), which is very similar to TCP—only more lightweight—configurations on active devices in a cluster rarely become unsynchronized.*

Synchronizing Files

If you need to synchronize a specific file, enter the following command on the device to which you are synchronizing the file: **exec nsrp sync file name** *name_str* **from peer**. If you need to synchronize all files, enter **exec nsrp sync file from peer**.

You can synchronize PKI objects—such as local and CA certificates, key pairs, and CRLs—either with an RTO sync or a configuration sync operation:

- If RTO synchronization is enabled enter **exec nsrp sync global-config run** (which does not require rebooting the device), and then **exec nsrp sync rto pki from peer**
- If RTO synchronization is disabled: **exec nsrp sync global-config save** and then reboot the device.

Synchronizing RTOs

If you have enabled RTO mirror synchronization on a device in a cluster (see "Run-Time Objects" on page 21), when the device reboots, the RTOs automatically resynchronize. However, if you disable RTO mirror synchronization—perhaps to perform some debugging or maintenance on the device—when you again enable RTO synchronization, you must manually resync all the RTOs. To do that, you can use the **exec nsrp sync rto all** command. To resync only selected RTOs such as ARP, DNS, sessions, or VPNs—you can use the following CLI command: **exec nsrp sync rto** { **arp** | **auth-table** | **dhcp** | **dns** | **l2tp** | **phase1-sa** | **pki** | **rm** | **session** | **vpn** }.

To enable RTO synchronization to begin automatically when a member in an NSRP cluster detects another member in the cluster, use the **set nsrp rto-mirror sync** command. When you need to synchronize RTOs manually, use the **exec nsrp sync rto** { **all** | **arp** | **auth-table** | **dhcp** | **dns** | **l2tp** | **phase1-sa** | **pki** | **rm** | **session** | **vpn** } command.

Example: Manually Resynchronizing RTOs

In this example, devices A and B are in NSRP cluster 1 and VSD groups 1 and 2. Device A is the master of VSD group 1 and the backup in VSD group 2. Device B is the master of VSD group 2 and the backup in VSD group 1.

You want to do some troubleshooting on device B, and you do not want to disconnect it from the network. You force device B to become the backup in VSD group 2, and then you disable RTO synchronization. Device A becomes the master of both VSD groups. After you finish troubleshooting device B, you again enable RTO mirror synchronization and then manually resync the RTOs from device A to device B. Finally, you reassign device B as the master of VSD group 2.

WebUI

> ***Note:*** *The manual synchronization of RTOs is only available through the CLI.*

CLI

Device B

```
exec nsrp vsd-group id 2 mode backup
unset nsrp rto-mirror sync
```

Device B is no longer processing traffic nor synchronizing RTOs with device A. At this point, you can troubleshoot device B without affecting the traffic-processing performance of device A.

```
set nsrp rto-mirror sync
exec nsrp sync rto all from peer
exec nsrp vsd-group id 2 mode master
```

Example: Adding a Device to an Active NSRP Cluster

In this example, you add device A, which had previously been functioning as a single security appliance, to VSD groups 0 and 1 in NSRP cluster with cluster ID 1 and name “cluster1”. You must unset the previous configurations on device A, reboot it, and then synchronize the configuration, files, and RTOs from the master of both VSD groups. You then assign device A as the master of VSD group 0.

WebUI

> ***Note:*** *The cold start synchronization feature is only available through the CLI.*

CLI

Device A

`unset all`[15]

The following prompt appears: “Erase all system config, are you sure y / [n]?”

Press the **Y** key.

The system configuration is returned to the factory default settings.

`reset`

The following prompt appears: “Configuration modified, save? [y] / n”

Press the **N** key.

The following prompt appears: “System reset, are you sure? y / [n] n”

Press the **Y** key.

The system reboots.

`set nsrp cluster id 1`

15. If you do not first use the **unset all** command, the **exec nsrp sync global-config** command appends new configuration settings to existing settings. (Note: The NetScreen device generates an error message for each duplicate setting that is synchronized.)

```
set nsrp cluster name cluster1
exec nsrp sync file
exec nsrp sync global-config
set nsrp rto-mirror sync
exec nsrp vsd-group id 0 mode master
save all[16]
```

Synchronizing System Clocks

NSRP contains a mechanism for synchronizing the system clocks of NSRP cluster members. When you set the system clock manually, the NSRP time synchronization mechanism keeps the members' clocks properly synchronized. However, when you use the Network Time Protocol (NTP) to set the system clocks on all the cluster members, and then use NSRP to synchronize the time among them, the time can become unsynchronized. Although the resolution for NSRP synchronization is in seconds, NTP has sub-second resolution. Because the time on each cluster member might differ by a few seconds due to processing delays, NetScreen recommends that you disable NSRP time synchronization when NTP is enabled on all cluster members and each member can update its system clock from an NTP server. To disable the NSRP time synchronization function, enter the following command:

```
set ntp no-ha-sync
```

16. Using the **save all** command saves the configurations in all virtual systems as well as at the root level. Using the **save** command saves the configuration at the root level only.

DUAL HA INTERFACES

The basic principle of NSRP is that there be no single point of failure. In addition to redundant devices, NetScreen devices either have dedicated physical redundant HA interfaces (HA1 and HA2) or you can bind two generic interfaces to the HA zone to provide HA interface redundancy.

In addition, you can create redundant security zone interfaces.

All NSRP information passes between cluster members via the two HA interfaces. To better distribute the out-of-band bandwidth, HA1 handles the NSRP control messages while HA2 handles the network data messages. If either port fails on a NetScreen device with gigabit HA1 and HA2 interfaces, the remaining active port assumes both kinds of traffic. For NetScreen devices that must use a 100-megabit interface for the data link, a failure of the data link results in one active HA link for control messages only. If the control link fails on such devices, then the data link becomes the control link and sends and receives control messages only.

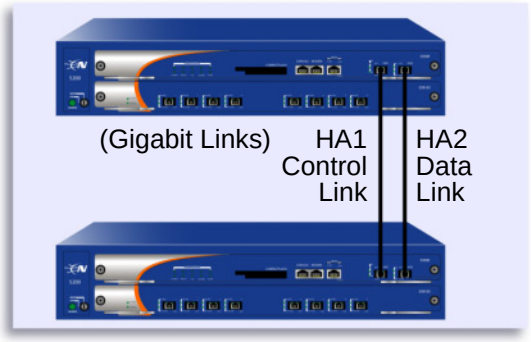

If either HA1 or HA2 fails, then the control and data messages are both sent over the remaining HA link.

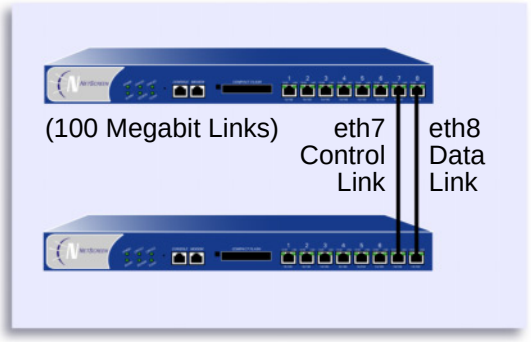

If either ethernet7 or ethernet8 fails, then only control messages are sent over the remaining HA link.

Note: *If you use a switch between HA ports, use port-based VLANs, which do not conflict with the VLAN tags on the forwarded packets.*

On NetScreen devices that do not have dedicated HA interfaces, you must bind one or two physical ethernet interfaces to the HA zone. If you bind a single gigabit interface to the HA zone, the HA link supports both control and data messages. If you bind one 100-megabit interface to the HA zone, the HA link supports control messages only.

If you bind two interfaces (gigabit or 100-megabit) to the HA zone, the interface with the lower number becomes the control link, and the interface with the higher number becomes the data link. For example, if you bind only ethernet 8 to the HA zone, it becomes the control link. If you then bind ethernet7 to the HA zone, it becomes the control link (because it has a lower number than ethernet8), and ethernet8 changes to the data link. (For information on binding an interface to a zone, see “Binding an Interface to a Security Zone” on page **2**-76.)

The order in which you cable the HA interfaces together also affects which becomes the control and data links. If ethernet7 and ethernet8 are both bound to the HA zone, but you only cable the ethernet8 interfaces together, then ethernet8 becomes the control link. If you then cable the ethernet7 interfaces together, ethernet7 becomes the control link (because it is active and has a lower number than ethernet8) and ethernet8 becomes the data link. The same principle also applies to the HA1 and HA2 interfaces.

On NetScreen devices that do not have dedicated HA interfaces, you can also designate an interface bound to a security zone to handle HA control messages. Use the CLI command **set nsrp interface** *interface*.

Control Messages

There are two kinds of control messages: heartbeats and HA messages.

Heartbeats: Heartbeats are sent periodically to establish and sustain communications among the NSRP cluster members, VSD group members, and RTO mirrors. The heartbeats continually advertise the sender’s member status, and the health of its system and network connectivity. The three kinds of heartbeat messages are as follows:

- HA physical link heartbeats
- VSD heartbeats
- RTO heartbeats

HA physical link heartbeats are broadcast messages from the HA1 and HA2 interfaces of each member of an NSRP cluster to the other member. The purpose of these messages is to monitor the health of the HA interfaces. If, for example, one member does not receive three consecutive heartbeats from HA1, both devices transfer transmission of the control messages to HA2.

VSD heartbeats are broadcast from the HA1 interface of each member of a VSD group. The VSD group uses these messages to monitor the membership status of all its members. If, for example, the master advertises that it has become inoperable, the primary backup immediately becomes the VSD group master.

Each member of a mirror group broadcasts RTO heartbeats from the HA1 interface. The purpose of these messages is to locate an active peer and then maintain the mirror relationship by sending group active messages. If, for example, a device does not receive 16 consecutive RTO heartbeats from its peer, it changes its state from active to set.

Note: *If you remove a device from a mirror group, it enters the undefined state and transmits a "group detach" message to its peer. The peer immediately changes its state from active to set without waiting for the missing heartbeats to exceed the threshold.*

HA Messages: The two kinds of HA messages are as follows:

- Configuration messages – The network and configuration settings that the master sends to the other VSD group member
- RTO messages – The RTOs that the master sends to the other RTO mirror

The HA messages contain the information that enables the backup to become the master without causing a service interruption.

Data Messages (Packet Forwarding)

Data messages are IP packets traversing the firewall that the backup in a VSD group must forward to the device acting as master. When a packet arrives at the interface of a NetScreen device in an active/active configuration, the device first identifies which VSD group must process the packet. If the device that receives the packet is the master of the identified VSD group, it processes the packet itself. If the device is not the master, it forwards the packet over the HA data link to the master.

For example, a load-balancing router might send the first packet in a session to device A (master of VSD group 1), which creates an entry in its session table. If the router performs load balancing by sending packets round-robin (that is, the router sends each packet to a NetScreen device in turn), the router might send the next packet to

device B (backup of VSD group 1). Because a session entry exists in device A, device B forwards the packet across the data link[17] to device A, which processes it.

Inbound packet forwarding across the data link occurs only when the NetScreen devices are in an active/active configuration in Route mode. When in NAT mode, the router always sends the incoming packets to the MIP, VIP, or VPN tunnel gateway, although the NetScreen device that receives the returning outbound packet might forward it across the data link to the device that has the session entry to which the packet belongs. This kind of packet forwarding produces an h-shaped path. Like the down stroke in the letter *h*, the inbound packet goes straight through one device, but the outbound packet might be sent halfway through the other device and then forwarded across the data link to the first device.

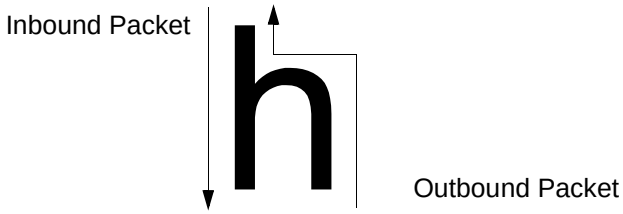

Dynamic Routing Advisory

If an NSRP cluster is in a dynamic routing environment and you disable packet forwarding (**unset nsrp data-forwarding**), traffic arriving at an inactive interface can be lost[18]. Because the NetScreen device cannot forward traffic across the data link to the NetScreen device on which the interface is active, it drops it. To avoid this when you disable packet forwarding, the NetScreen device indicates the status of interfaces belonging to a non-master VSD on a device as “down” instead of just “inactive”. This status signals routers not to send traffic to these interfaces.

17. If there is no data link, the NetScreen device that receives the packet drops it immediately.

18. An inactive interface is an interface belonging to a VSD that is not the master on that device.

Dual HA Link Probes

You can connect the redundant HA interfaces by directly cabling HA ports on one device to the HA ports on another device. Or, you can connect the HA ports on two devices through one or more switched networks. In the following configuration, the HA1 port on the device NetScreen-A is connected to the HA1 port on NetScreen-B via two switches, Switch-1a and Switch-1b. To provide a redundant HA interface, the HA2 port on NetScreen-A is connected to the HA2 port on NetScreen-B via Switch-2a and Switch-2b. In this configuration, the link between the HA1 ports on NetScreen-A and NetScreen-B handles NSRP control messages, while the HA2 link handles network data messages. If the link between the HA1 port on NetScreen-A and Switch-1a goes down, NetScreen-A transfers transmission of the control messages to its HA2 port. However, NetScreen-B does not recognize the failure of the HA1 link as its HA1 port is still active, and rejects the NSRP control messages sent by NetScreen-A on the HA2 link.

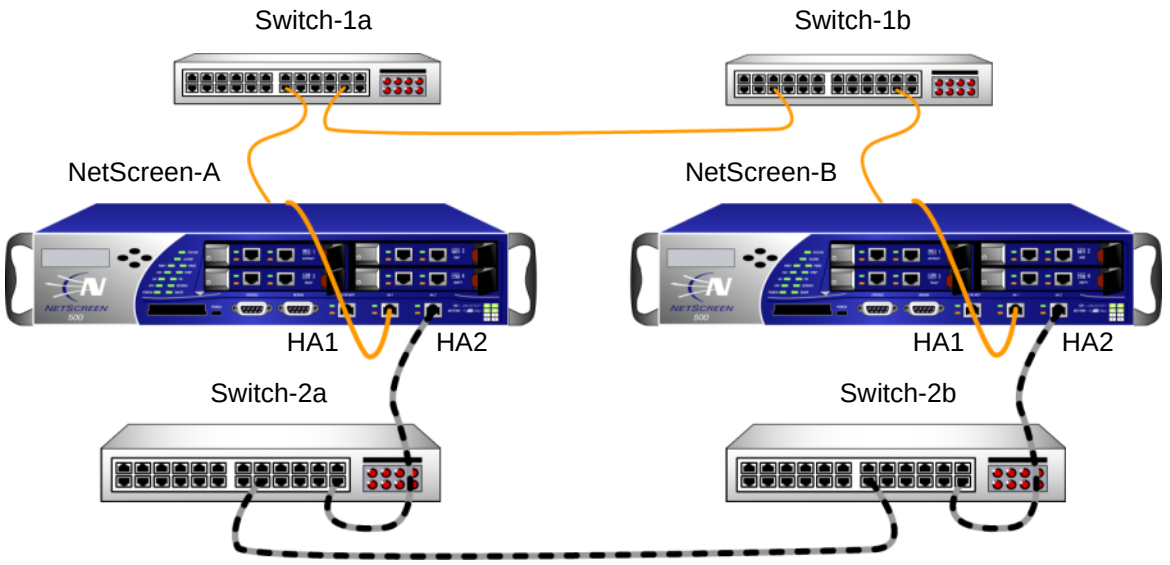

Control Link = Solid Orange Cable

Data Link = Dashed Black/Gray Cable

To prevent this situation, you can configure a NetScreen device to monitor the status of a HA link by sending NSRP probe requests on the HA link to its peer. If a reply is received from the peer on the HA link, the request is considered successful and the HA link is assumed to be up. If no reply is received from the peer within the constraints specified, the HA link is considered to be down. This enables NetScreen devices to switch transmission of control messages to an available HA link when necessary, even if there is no physical failure on the HA ports on either device.

There are two ways that probe requests can be sent on an HA link:

- **Manually by the administrator** Probes are sent on a specific HA links once every second for a specified number of times. If no reply is received from the peer after the specified number of probes are sent, the HA link is considered to be down. Probes are sent out immediately after you execute the command.
- **Automatically by ScreenOS** Probes are sent by on all HA links once every second. (You can optionally specify the HA zone interface and the interval at which probes are sent.) By default, if five consecutive probes are sent without receiving a reply from the peer, the link is considered to be down; you can specify a different threshold value for determining when the link is down. Note that even when a primary HA link is down, the NetScreen device continues to send probes on that link. If the primary HA link connection is restored and peer responses are once again received on the link, the NetScreen devices can switch transmission of control messages back to the primary HA link.

Example: Sending Link Probes Manually

In this example, the ethernet7 and ethernet8 interfaces on the NetScreen device are bound to the HA zone. You configure 5 link probes to be sent on the ethernet8 interface to the peer's MAC address 00e02000080. (Note that if you do not specify a MAC address, the default NSRP MAC address is used.)

WebUI

> ***Note:*** *You must use the CLI to send probes manually on an HA link.*

CLI

```
exec nsrp probe ethernet8 00e02000080 count 5
```

Example: Sending Link Probes Automatically

In this example, the ethernet7 and ethernet8 interfaces on the NetScreen device are bound to the HA zone. You configure link probes to be automatically sent to both interfaces at three-second intervals. You also set the threshold value so that if there is no reply from the peer after sending four consecutive requests, the HA link is considered to be down.

WebUI

Network > NSRP > Link: Enter the following, and then click **Apply**:

Enable HA Link Probe: (select)

Interval: 3

Threshold: 5

CLI

```
set nsrp ha-link probe interval 3 threshold 4
```

SETUP PROCEDURE

To configure two NetScreen devices for high availability, you must cable them to the network and to each other, and then configure them for HA using NSRP.

Cabling for a Full-Mesh Configuration

The following diagrams illustrate the cabling of two NetScreen devices to each other and to redundant pairs of internal switches and external switches. The external switches are then cabled to a pair of redundant routers running VRRP, completing the full-mesh configuration. The first diagram shows two NetScreen devices with dedicated HA interfaces. The second diagram shows two NetScreen devices using network interfaces for HA traffic.

Note: *Depending on the topology in which you are deploying the NetScreen devices and the kinds of switches and routers you use, the cabling presented in the following diagram might differ from what your network requires.*

NetScreen Devices with Dedicated HA Interfaces

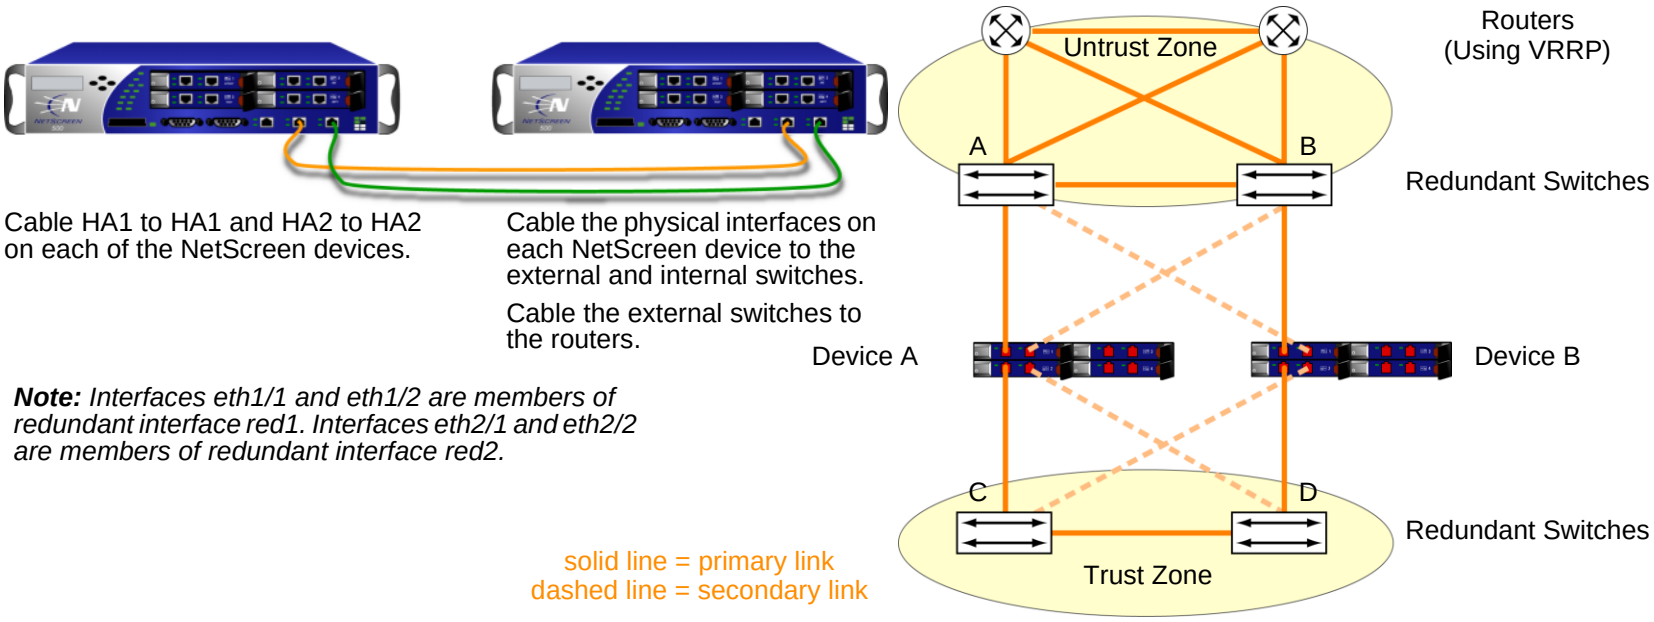

Cable two NetScreen devices (device A and device B) for NSRP in a full-mesh configuration as follows:

NetScreen A and NetScreen B: HA Links

1. Cable together the HA1 interfaces on each NetScreen device.
2. Cable together the HA2 interfaces on each NetScreen device.

NetScreen A: Redundant1 (eth1/1 and eth1/2), Untrust Zone

3. Cable ethernet1/1 to external switch A. (ethernet1/1 is one of the two physical interfaces bound to the redundant interface red1 in the Untrust zone.)
4. Cable ethernet1/2 to external switch B. (ethernet1/2 is the other physical interface bound to red1 in the Untrust zone.)

NetScreen A: Redundant2 (eth2/1 and eth2/2), Trust Zone

5. Cable ethernet2/1 to internal switch C. (ethernet2/1 is one of the two physical interfaces bound to the redundant interface red2 in the Trust zone.)
6. Cable ethernet2/2 to internal switch D. (ethernet2/2 is the other physical interface bound to red2 in the Trust zone.)

NetScreen B: Redundant1 (eth1/1 and eth1/2), Untrust Zone

7. Cable ethernet1/1 to external switch B. (ethernet1/1 is one of the two physical interfaces bound to the redundant interface red1 in the Untrust zone.)
8. Cable ethernet1/2 to external switch A. (ethernet1/2 is the other physical interface bound to red1 in the Untrust zone.)

NetScreen B: Redundant2 (eth2/1 and eth2/2), Trust Zone

9. Cable ethernet2/1 to internal switch D. (ethernet2/1 is one of the two physical interfaces bound to the redundant interface red2 in the Trust zone.)
10. Cable ethernet2/2 to internal switch C. (ethernet2/2 is the other physical interface bound to red2 in the Trust zone.)

Switches and Routers

11. Cable the redundant external switches together.
12. Cable the external switches to the redundant routers in the same configuration that you used to cable the NetScreen devices to the switches.
13. Cable the internal redundant switches together.

NetScreen Devices Using Network Interfaces for HA Links

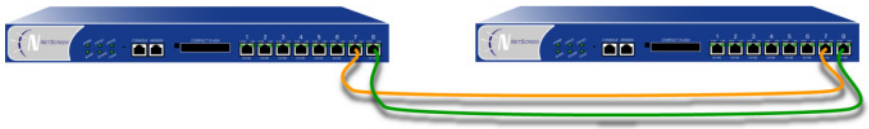

Bind ethernet7 and ethernet8 to the HA zone on each of the NetScreen devices.

Then cable together the interfaces bound to the HA zone:

- eth7 on device A to eth7 on device B
- eth8 on device A to eth8 on device B

Cable the physical interfaces on each NetScreen device to the external and internal switches.

Cable the external switches to the routers.

Note: *Interfaces eth1 and eth2 are members of redundant interface red1. Interfaces eth3 and eth4 are members of redundant interface red2.*

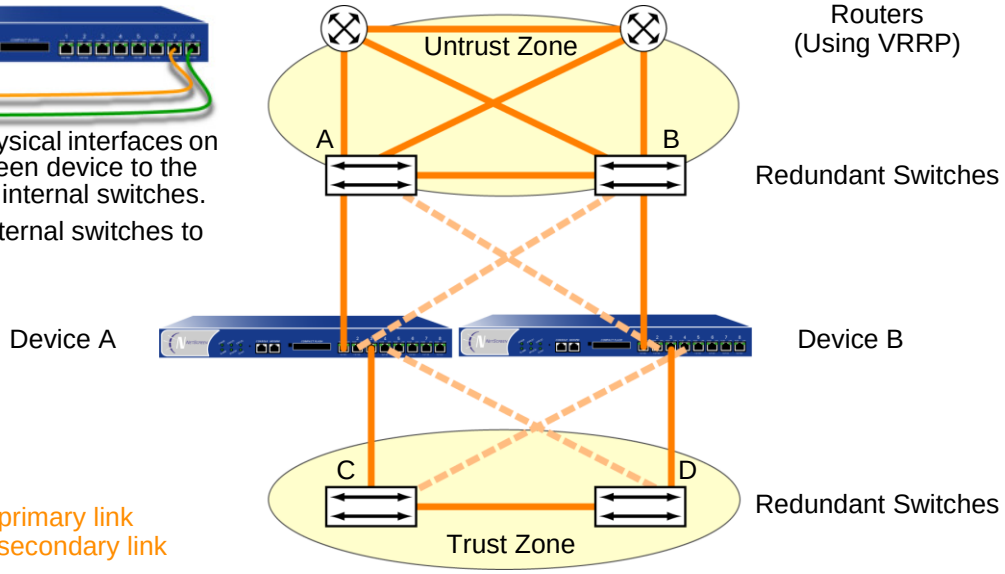

After binding ethernet7 and ethernet8 to the HA zone on both NetScreen devices (device A and device B), cable the NetScreen devices for NSRP in a full-mesh configuration as follows:

NetScreen A and NetScreen B: HA Links

1. Cable together the ethernet7 interfaces on each NetScreen device.
2. Cable together the ethernet8 interfaces on each NetScreen device.

NetScreen A: Redundant1 (ethernet1 and ethernet2), Untrust Zone

3. Cable ethernet1 to external switch A. (ethernet1 is one of the two physical interfaces bound to the redundant interface red1 in the Untrust zone.)
4. Cable ethernet2 to external switch B. (ethernet2 is the other physical interface bound to red1 in the Untrust zone.)

NetScreen A: Redundant2 (ethernet3 and ethernet4), Trust Zone

5. Cable ethernet3 to internal switch C. (ethernet3 is one of the two physical interfaces bound to the redundant interface red2 in the Trust zone.)
6. Cable ethernet4 to internal switch D. (ethernet4 is the other physical interface bound to red2 in the Trust zone.)

NetScreen B: Redundant1 (ethernet1 and ethernet2), Untrust Zone

7. Cable ethernet1 to external switch B. (ethernet1 is one of the two physical interfaces bound to the redundant interface red1 in the Untrust zone.)
8. Cable ethernet2 to external switch A. (ethernet2 is the other physical interface bound to red1 in the Untrust zone.)

NetScreen B: Redundant2 (ethernet3 and ethernet4), Trust Zone

9. Cable ethernet3 to internal switch D. (ethernet3 is one of the two physical interfaces bound to the redundant interface red2 in the Trust zone.)
10. Cable ethernet4 to internal switch C. (ethernet4 is the other physical interface bound to red2 in the Trust zone.)

Switches and Routers

11. Cable the redundant external switches together.
12. Cable the external switches to the redundant routers in the same configuration that you used to cable the NetScreen devices to the switches.
13. Cable the internal redundant switches together.

Active/Active NSRP Configuration

After cabling the NetScreen devices together and to the surrounding network devices, you can then configure them for HA. A complete active/active configuration involves the following steps:

1. Creating an NSRP cluster, which automatically includes the creation of VSD group 0
2. Creating a second VSD group within the cluster
3. Enabling device failure tracking methods—such as interface monitoring and path monitoring

Example: NSRP for an Active/Active Configuration

In this example, which builds upon the interfaces configured in “Example: Creating Redundant Interfaces for VSIs” on page 96, you create an NSRP cluster with ID 1 and name “cluster1” for two NetScreen devices—device A and device B—which do not have any other user-defined settings configured.

Note: *To enable command propagation, you must first define the cluster ID number on each device. The following settings are not propagated and must be configured on each device in the cluster: VSD group, VSD priority, authentication and encryption passwords, manage IP addresses, and IP tracking settings. All other commands are propagated among devices within the cluster.*

When you create the NSRP cluster, the NetScreen device automatically creates VSD group 0[19]. You define VSD group 1. You assign device A priority 1 in VSD group 0 and priority 100 (the default) in VSD group 1. You assign device B priority 1 in VSD group 1 and leave its priority at the default (100) in VSD group 0.

You set the interface monitoring option to monitor the two redundant interfaces—redundant1 and redundant2—for layer 2 network connectivity. If the primary physical interface for either of the monitored interfaces fails, the device immediately fails over to the secondary interface. If both physical interfaces comprising the members of a monitored redundant interface fail, the device fails over to the other device.

19. The VSD group ID “0” does not appear in the names of VSIs in VSD 0. Instead of *redundant1:0,* the VSI is identified simply as *redundant1.*

You define the ethernet2/1 interface as a secondary link for VSD heartbeat messages and the number of gratuitous ARPs after a device failover has occurred to 5. Because HA cables run directly between the two NetScreen devices, communication between members of the NSRP cluster is neither authenticated nor encrypted.

You also set a route to the default gateway (210.1.1.250) for each Untrust zone VSI, and a route to the internal network for each Trust zone VSI. All security zones are in the trust-vr routing domain.

Finally, after the configurations for both devices are synchronized, you enable RTO synchronization.

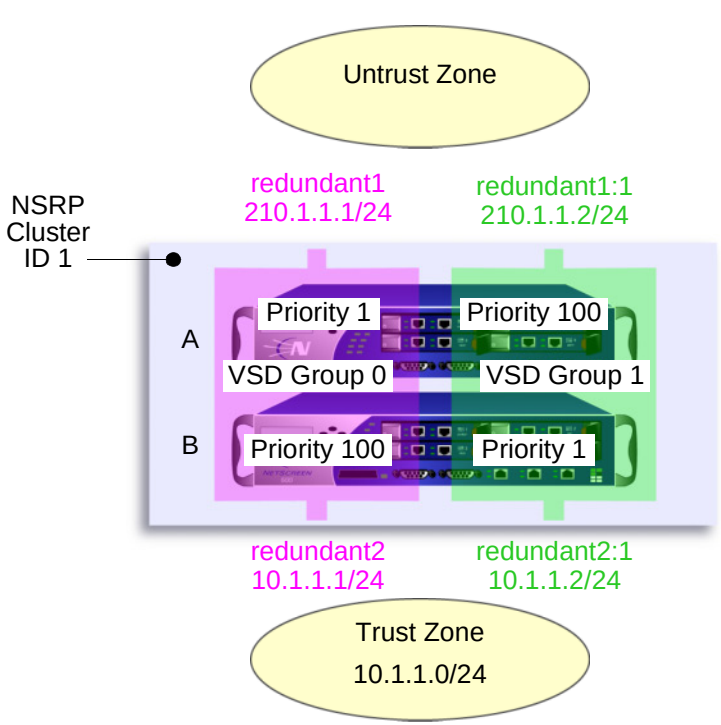

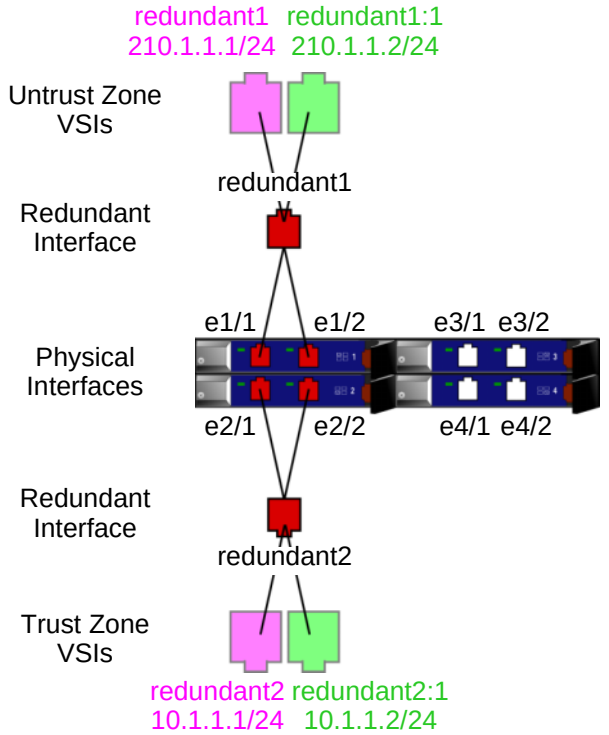

The IP address of the default gateway in the Untrust zone is 210.1.1.250.

The addresses and configuration shown here are identical on both NetScreen devices.

The only difference is the manage IP address.

On device A the manage IP is 10.1.1.21 and is on the redundant2 interface.

On device B the manage IP is 10.1.1.22, and is on the redundant2 interface.

WebUI (Device A)

1. Cluster and VSD Groups

Network > NSRP > Cluster: Type **1** in the Cluster ID field, and then click **Apply**.

Network > NSRP > VSD Group > Edit (for Group ID 0): Enter the following, and then click **OK**:

Priority: 1

Enable Preempt: (select)

Preempt Hold-Down Time (sec): 10[20]

Network > NSRP > VSD Group > New: Enter the following, and then click **OK**:

Group ID: 1

Priority: 100

Enable Preempt: (clear)

Preempt Hold-Down Time (s): 0

20. The hold-down time can be any length from 0 to 255 seconds, effectively delaying the failover to prevent a flurry of rapid failovers.

WebUI (Device B)

2. Cluster and VSD Groups

Network > NSRP > Cluster: Enter the following, and then click **Apply**[21]:

Cluster ID: 1

Number of Gratuitous ARPs to Resend: 5[22]

Network > NSRP > Link: Select **ethernet2/1** from the Secondary Link drop-down list, and then click **Apply**[23].

Network > NSRP > Synchronization: Select **NSRP RTO Synchronization**, and then click **Apply**.

Network > NSRP > VSD Group > New: Enter the following, and then click **OK**:

Group ID: 1

Priority: 1

Enable Preempt: (select)

Preempt Hold-Down Time (sec): 10

3. Redundant Interfaces and Manage IP

Network > Interfaces > New Redundant IF: Enter the following, and then click **OK**:

Interface Name: redundant1

Zone Name: Untrust

IP Address / Netmask: 210.1.1.1/24

Network > Interfaces > Edit (for ethernet1/1): Select **redundant1** in the "As member of" drop-down list, and then click **OK**.

---

21. You can only set the cluster name through the CLI.
22. This setting specifies that after a device failover, the new VSD group master sends 5 gratuitous ARP packets announcing the association of the VSI and virtual MAC address to the new master.
23. If both HA1 and HA2 links are lost, the VSD heartbeat messages pass via the ethernet2/1 in the Trust zone.

Network > Interfaces > Edit (for ethernet1/2): Select **redundant1** in the “As member of” drop-down list, and then click **OK**.

Network > Interfaces > New Redundant IF: Enter the following, and then click **Apply**:

Interface Name: redundant2

Zone Name: Trust

IP Address / Netmask: 10.1.1.1/24

> Enter **10.1.1.22** in the Manage IP field, and then click **OK**.

Network > Interfaces > Edit (for ethernet2/1): Select **redundant2** in the “As member of” drop-down list, and then click **OK**.

Network > Interfaces > Edit (for ethernet2/2): Select **redundant2** in the “As member of” drop-down list, and then click **OK**.

Network > NSRP > Monitor > Interface > VSD ID: Device Edit Interface: Select **redundant1** and **redundant2**, and then click **Apply**.

4. Virtual Security Interfaces

Network > Interfaces > New VSI IF: Enter the following, and then click **OK**:

Interface Name: VSI Base: redundant1

VSD Group: 1

IP Address / Netmask: 210.1.1.2/24

Network > Interfaces > New VSI IF: Enter the following, and then click **OK**:

Interface Name: VSI Base: redundant2

VSD Group: 1

IP Address / Netmask: 10.1.1.2/24

5. Routes

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: redundant1

Gateway IP Address: 210.1.1.250

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address: 0.0.0.0/0

Gateway: (select)

Interface: redundant1:1

Gateway IP Address: 210.1.1.250

WebUI (Device A)

6. Manage IP Address

Network > Interfaces > Edit (for redundant2): Enter **10.1.1.21** in the Manage IP field, and then click **OK**.

7. RTO Synchronization

Network > NSRP > Synchronization: Select **NSRP RTO Mirror Synchronization**, and then click **Apply**.

CLI (Device A)

1. Cluster and VSD Groups

```
set nsrp cluster id 1
set nsrp vsd-group id 0 preempt hold-down 10[24]
set nsrp vsd-group id 0 preempt
set nsrp vsd-group id 0 priority 1
set nsrp vsd-group id 1
set nsrp rto-mirror sync
save
```

CLI (Device B)

2. Cluster and VSD Groups

```
set nsrp cluster id 1[25]
set nsrp cluster name cluster1
set nsrp rto-mirror sync
set nsrp vsd-group id 1 priority 1[26]
set nsrp vsd-group id 1 preempt hold-down 10[27]
set nsrp vsd-group id 1 preempt
set nsrp secondary-path ethernet2/1[28]
set nsrp arp 5[29]
set arp always-on-dest[30]
```

24. The hold-down time can be any length from 0 to 255 seconds, effectively delaying the failover to prevent a flurry of rapid failovers.
25. Because devices A and B are both members of the same NSRP cluster, all subsequent commands (except those otherwise noted) that you enter on device B propagate to device A.
26. This command is not propagated.
27. This command is not propagated.
28. If both HA1 and HA2 links are lost, the VSD heartbeat messages pass via the ethernet2/1 in the Trust zone.
29. This setting specifies that after a device failover, the new VSD group master sends 5 gratuitous ARP packets announcing the association of the VSI and virtual MAC address to the new master.

3. Redundant Interfaces and Manage IP

```
set interface redundant1 zone untrust
set interface redundant1 ip 210.1.1.1/24
set interface ethernet1/1 group redundant1
set interface ethernet1/2 group redundant1
set interface redundant2 zone trust
set interface redundant2 ip 10.1.1.1/24
set interface redundant2 manage-ip 10.1.1.22
set interface ethernet2/1 group redundant2
set interface ethernet2/2 group redundant2
set nsrp monitor interface redundant1
set nsrp monitor interface redundant2
```

4. Virtual Security Interfaces

```
set interface redundant1:1 ip 210.1.1.2/24
set interface redundant2:1 ip 10.1.1.2/24
```

5. Routes

```
set vrouter trust-vr route 0.0.0.0/0 interface redundant1 gateway 210.1.1.250
set vrouter trust-vr route 0.0.0.0/0 interface redundant1:1 gateway 210.1.1.250
save
```

CLI (Device A)

6. Manage IP Address

```
set interface redundant2 manage-ip 10.1.1.21
```

7. RTO Synchronization

```
set nsrp rto-mirror sync
save
```

30. After you enter this command, the NetScreen device always does an ARP lookup to learn a destination MAC address instead of learning it from the source MAC in the originating ethernet frame. The external routers in this example are grouped as a virtual router running VRRP. Frames coming from this router use the virtual IP address as the source IP but the physical MAC address as the source MAC. If the router fails over and the NetScreen device has learned the MAC from the source MAC in the incoming frame, it would then direct return traffic to the wrong location. By doing an ARP lookup for the destination MAC, the NetScreen device can properly send traffic to the location of the new physical MAC address.

2

NSRP-Lite

NetScreen Redundancy Protocol (NSRP) is a proprietary protocol that is supported on select NetScreen devices to provide high availability (HA) services. NSRP-Lite is a lightweight version of standard NSRP and is supported only on some NetScreen devices running ScreenOS at Layer 3 in the OSI model (that is, the interfaces must be in either Route or NAT mode). NSRP-Lite only supports an Active/Passive configuration.

Note: *NSRP-Lite synchronizes configurations and files but not run-time objects (RTOs). If a failover occurs, all user sessions and VPN connections must be reestablished. Because of this, NetScreen recommends that you enable VPN monitoring with the rekey option on VPN tunnels so that they automatically reestablish themselves.*

This chapter explains the components of NSRP-Lite and describes how to configure a NetScreen device for HA using NSRP-Lite. The specific topics covered are as follows:

- "Introduction to NSRP-Lite" on page 59
  - "Clusters and VSD Groups" on page 60
  - "Default Settings" on page 61
- "Cluster" on page 62
  - "Cluster Name" on page 63
  - "Authentication and Encryption" on page 64
- "VSD Group" on page 65
  - "VSD Group Member States" on page 65
  - "Heartbeat Messages" on page 66
  - "Preempt Option" on page 67
- "Cabling and Configuring NSRP-Lite" on page 68

- "Configuration and File Synchronization" on page 76
  - "Synchronizing Configurations" on page 76
  - "Synchronizing Files" on page 77
  - "Disabling Configuration and File Synchronization" on page 78
- "Path Monitoring" on page 79
  - "Setting Thresholds" on page 80
  - "Weighting Tracked IP Addresses" on page 80
  - "IP Tracking for VPN Tunnel Failover" on page 81

INTRODUCTION TO NSRP-LITE

NSRP-Lite provides a simple high availability (HA) solution on some NetScreen devices. When you cable and configure two NetScreen devices for NSRP-Lite, one device acts as the master and actively processes network traffic. The other device acts as a backup, passively waiting to become master in the event that the current master cannot perform its functions. By connecting two NetScreen devices to your local network, configuring them for NSRP-Lite, and using a different Internet service provider (ISP) for each device, you can protect the local network against both device failure and ISP failure.

The master (Device A) actively processes traffic traversing the firewall between the local network and the Internet.

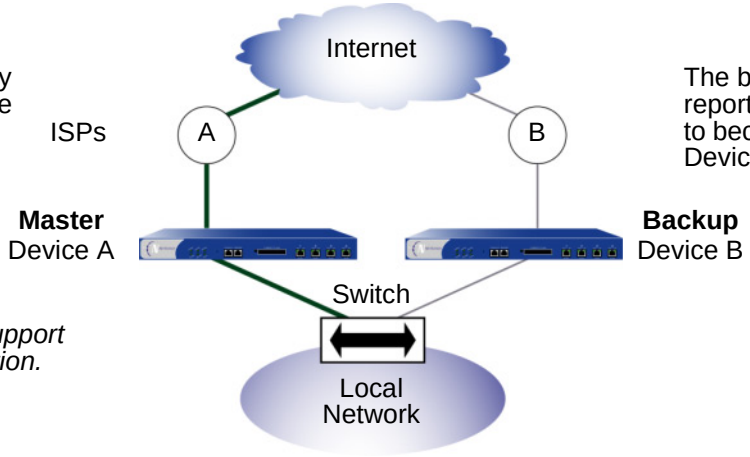

The backup (Device B) receives status reports from Device A and remains ready to become master should a failover occur. Device B does not process traffic.

Note: *NSRP-Lite does not support RTO or session synchronization.*

Device A fails

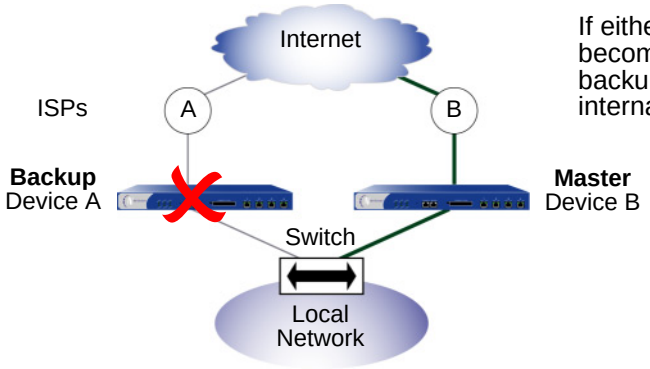

If either Device A or ISP A fails, Device B becomes master and Device A becomes backup (or it becomes inoperable if it has internal system problems).

ISP A fails

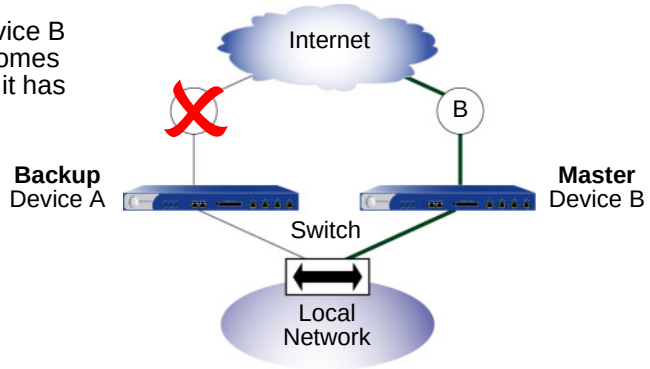

Clusters and VSD Groups

An NSRP-Lite cluster consists of a pair of NetScreen devices that comprise a single virtual security device (VSD) that provides redundant network connectivity. One physical device acts as the master of the VSD group and performs all of the network processes on traffic sent to the VSD. The other device acts as a backup to the master, and is ready to take over traffic processing should the current master fail or step down.

The two devices send VSD heartbeat messages to each other to provide status reports. If the backup receives a message that the master has experienced a network or system failure and has changed its status, the backup changes its status to master and begins actively processing traffic. This transition constitutes a failover.

Before two NetScreen devices can provide redundant services, you must group them in the same NSRP cluster by assigning a cluster ID between 1 and 7[1]. When a NetScreen device becomes a member of a cluster, it automatically becomes a member of VSD group 0, and all Trust zone interfaces become virtual security interfaces (VSIs) for VSD group 0.

A VSI binding can shift from the Trust zone interface of one physical device to that of another device, as the device acting as master of the VSD group shifts and claims the VSI for itself. The Untrust zone interface remains as a local interface, always dedicated to the physical NetScreen device that hosts it.

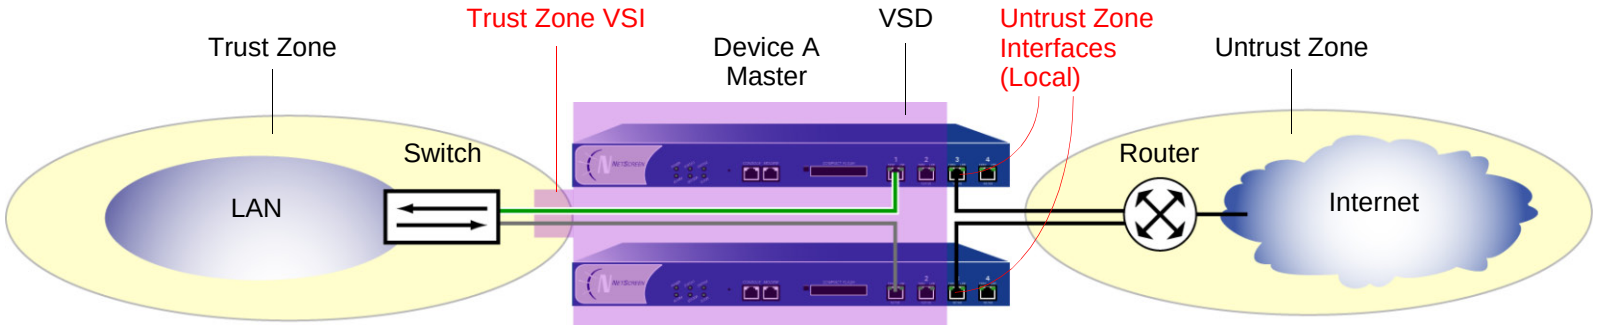

1. Assigning an ID of 0 removes a device from a cluster.

Default Settings

The basic NSRP configuration uses the following default settings:

- VSD Group Information
  - VSD group ID: 0
  - Device priority in the VSD group: 100
  - Preempt option: disabled
  - Preempt hold-down time: 0 seconds
  - Initial state hold-down time: 5 seconds
  - Heartbeat interval: 1000 milliseconds
  - Lost heartbeat threshold: 3
- NSRP Link Information
  - Number of gratuitous ARPs: 4
  - NSRP encryption: disabled
  - NSRP authentication: disabled
  - Interfaces monitored: none
  - Secondary path: none

When you set a NetScreen device in an NSRP cluster, the NetScreen device automatically creates VSD group 0 and transforms physical interfaces bound to the Trust zone into Virtual Security Interfaces (VSIs) for VSD group 0.

CLUSTER

An NSRP cluster consists of a group of NetScreen devices that enforce the same overall security policy and share the same configuration settings. When you assign a NetScreen device to an NSRP cluster, any changes made to the configuration on one member of the cluster propagate to the other[2]. Members of the same NSRP cluster maintain identical settings for the following:

- Policies and policy objects (such as addresses, services, VPNs, users, and schedules)
- System parameters (such as settings for authentication servers, DNS, SNMP, syslog, URL blocking, firewall detection options, and so forth)

Members of a cluster do not propagate the following configuration settings:

Non-Propagating Commands	
NSRP	• set/unset nsrp cluster id *number* • set/unset nsrp auth password *pswd_str* • set/unset nsrp encrypt password *pswd_str* • set/unset nsrp monitor interface *interface* • set/unset nsrp vsd-group id *id_num* { mode *string* | preempt | priority *number* } • set/unset nsrp rto-mirror …
Interface	• set/unset interface *interface* manage-ip *ip_addr* • set/unset interface *interface* phy … • set/unset interface *interface* bandwidth *number* • set/unset interface redundant *number* phy primary *interface* • All commands pertaining to local interfaces
Monitored Objects	• All IP tracking, zone monitoring, and interface monitoring commands

2. You can disable configuration and file synchronization. For information, see "Disabling Configuration and File Synchronization" on page 78.

Non-Propagating Commands	
Console Settings	• All console commands (set/unset console …)
Hostname	• set/unset hostname *name_str*
SNMP	• set/unset snmp name *name_str*
Virtual Router	• set/unset vrouter *name_str* router-id *ip_addr*
Clear*	• All clear commands (clear admin, clear dhcp, …)
Debug†	• All debug commands (debug alarm, debug arp, …)

* By default, NSRP cluster members do not propagate the **clear** commands. To propagate a clear command to all devices in an NSRP cluster, insert the keyword **cluster** into the command. For example, **clear cluster admin …**, **clear cluster dhcp …**

† By default, NSRP cluster members do not propagate the **debug** commands. To propagate a debug command to all devices in an NSRP cluster, insert the keyword **cluster** into the **debug** command. For example, **debug cluster alarm …**, **debug cluster arp …**

Cluster Name

Because NSRP cluster members can have different host names, a failover can disrupt SNMP communication and the validity of digital certificates because SNMP communication and certificates rely on the host name of a device to function properly.

To define a single name for all cluster members, type the following CLI command:

```
set nsrp cluster name name_str
```

Use the cluster name when configuring the SNMP host name for the NetScreen device (**set snmp name** *name_str*) and when defining the common name in a PKCS10 certificate request file.

The use of a single name for all cluster members allows SNMP communication and digital certificate use to continue without interruption after a device failover.

Authentication and Encryption

Because of the sensitive nature of NSRP communications, you can secure all NSRP traffic through encryption and authentication. For encryption and authentication, NSRP supports the DES and MD5 algorithms respectively.

Note: *When the devices are cabled directly to one another, there is no need to use authentication and encryption. However, if the devices are cabled through a switch to which other devices connect, you might consider implementing these additional security measures.*

To enable authentication or encryption, you must provide passwords on each device in the cluster.

WebUI

Network > NSRP > Cluster: Enter the following, and then click **Apply**:

NSRP Authentication Password: (select), *pswd_str*

NSRP Encryption Password: (select), *pswd_str*

CLI

```
set nsrp auth password pswd_str
set nsrp encrypt password pswd_str
```

VSD GROUP

A Virtual Security Device (VSD) group is a pair of physical NetScreen devices that collectively comprise a single VSD. One physical device acts as the master of the VSD group. The virtual security interface (VSI) of the VSD is bound to the Trust zone physical interface of the master. The other physical device acts as the backup[3]. If the master device fails, the VSD fails over to the backup and the VSI binding is transferred to the physical interface on the backup, which is instantly promoted to master[4].

VSD Group Member States

The members of a VSD group can be in one of six states:

- Master – The state of a VSD group member that processes traffic sent to the VSI.
- Primary Backup – The state of a VSD group member that becomes the master should the current master step down. The election process uses device priorities to determine which member to promote.
- Backup – The state of a VSD group member that monitors the status of the primary backup and elects one of the backup devices to primary backup if the current one steps down.
- Initial – The transient state of a VSD group member while it joins a VSD group, either when the device boots up or when it is added via the **set nsrp vsd-group id** *id_num* command.

  You can specify how long a VSD group member stays in the initial state with the **set nsrp vsd-group init-hold** *number* command. The default (and minimum) setting is 5. To determine the initial state hold-down time, multiply init-hold value by the VSD heartbeat-interval (init-hold x hb-interval = initial state hold-down time). For example, if the init-hold is 5 and the hb-interval is 1000 milliseconds, then the initial state hold-down time is 15,000 milliseconds, or 5 seconds (5 x 1000 = 5000).

  Note: *If you reduce the VSD heartbeat interval, you should increase the init-hold value. For information on configuring the heartbeat interval, see* *"Heartbeat Messages" on page 66**.*

3. In the current release, a VSD group can have two members. In later releases, there might be more than two members, in which case, one device acts as a master, another as a primary backup, and the remaining VSD group members as backups.
4. When using BGP and both the Trust and Untrust zones are in the same virtual routing domain, the NetScreen device advertises the subnet connected to the Trust zone VSIs of both the master (active) and backup (passive) VSD group members.

- Ineligible – The state that an administrator purposefully assigns to a VSD group member so that it cannot participate in the election process. To do this, use the **set nsrp vsd-group id** *id_num* **mode ineligible** command.
- Inoperable – The state of a VSD group member after a system check determines that the device has an internal problem (such as no processing boards) or a network connection problem (such as when an interface link fails).

Note: *When the device returns from either the ineligible state (when you use the* ***exec nsrp vsd-group id*** *id_num* ***mode { backup | init | master | pb }*** *command) or inoperable state (when the system or network problem has been corrected), it must first pass through the initial state.*

Heartbeat Messages

Heartbeat messages continually advertise the sender's member status, and the health of its system and network connectivity. Every VSD group member—even if it is in the initial, ineligible, or inoperable state—communicates with its group members by sending a heartbeat message every second. These messages allow every member to know the current state of every other member. The heartbeat message includes the following information:

- Unit ID of the device
- VSD group ID
- VSD group member status
- Device priority

The interval for sending VSD heartbeats is configurable (200, 600, 800, or 1000 milliseconds; 1000ms is the default). The CLI command—which applies globally to all VSD group members—is **set nsrp vsd-group hb-interval** *number*. You can also configure the lost heartbeat threshold that is used to determine when a VSD group member is considered as missing. The CLI command, which also applies globally to all VSD group members, is **set nsrp vsd hb-threshold** *number*. The minimum value for the lost heartbeat threshold is 3.

Preempt Option

You can determine whether a better priority number (closer to zero) can initiate a failover by setting the device that you want to be master in preempt mode. If you enable the preempt option on that device, it becomes the master of the VSD group if the current master has a lesser priority number (farther from zero). If you disable this option, a master with a lesser priority than a backup can keep its position (unless some other factor, such as an internal problem or faulty network connectivity, causes a failover).

To change the priority of a device (the default value is 100) and enable or disable the preempt option, use the following CLI commands:

set nsrp vsd-group id *number* **priority** *number*

unset nsrp vsd-group id *number* **priority**[5]

set/unset nsrp vsd-group id *number* **preempt**

Using the hold-down time to delay a failover can prevent a flurry of rapid failovers in the event of port-flickering on an adjacent switch and also ensure that surrounding network devices have sufficient time to negotiate new links before the new master becomes available. You can use the following CLI command to set the hold-down time—used for delaying the preempted failover—to any length from 0 to 600 seconds:

set nsrp vsd-group id *number* **preempt hold-down** *number*

5. This command returns the priority to its default value of 100.

CABLING AND CONFIGURING NSRP-LITE

To set up two NetScreen devices for high availability, you must cable them to the network and configure them for NSRP-Lite.

Use an RJ-45 ethernet cable to connect the Untrust zone interface on both NetScreen devices to the external router. Use another RJ-45 ethernet cable to connect one of the Trust zone ports to the internal switch on the local area network (LAN). Because the heartbeat messages used for NSRP communications is a proprietary protocol, these messages cannot be routed at Layer 3 in the OSI model. Therefore, use only a Layer 2 switch or hub to connect the devices in the Trust zone.

1. Cable the Untrust zone port on each NetScreen device to the external router.

2. Cable the Trust zone port(s) on each NetScreen device to the internal switch.

3. Connect other hosts on the internal LAN to the switch.

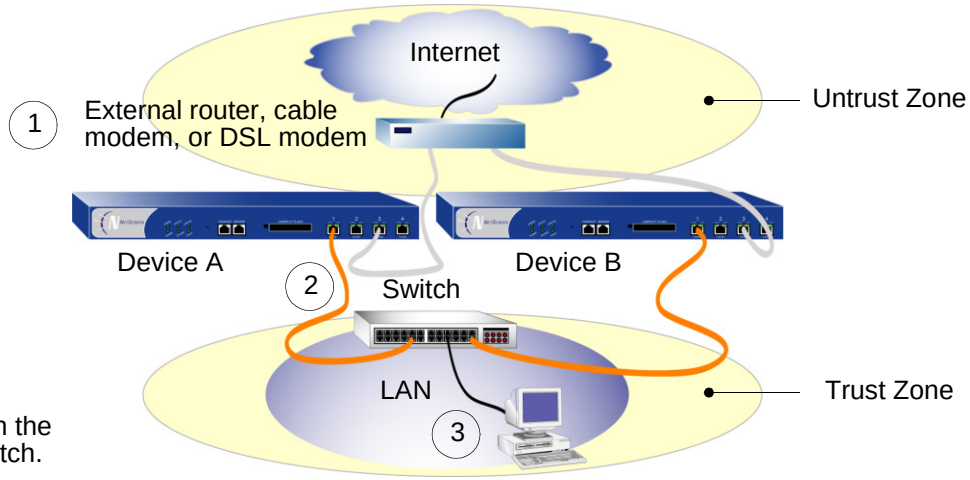

Example: Configuring NSRP-Lite

In this example, you configure two NetScreen devices for high availability using NSRP-Lite. The IP addresses for the interfaces are as follows:

- NetScreen-A
  - ethernet3 – Untrust zone interface, 1.1.1.1/24, manage IP: 1.1.1.2

    This is a local interface, not a VSI.
  - ethernet1:1 – Trust zone interface, 10.1.1.1/24, NAT mode

    This is a VSI for VSD group 1.
  - ethernet4 – HA zone interface

    This is for the control link for HA communications between the two devices. You also set ethernet1 as the interface for secondary path control link communications in the event that ethernet4 fails.
- NetScreen-B
  - ethernet3 – Untrust zone interface, 1.1.2.1/24, manage IP: 1.1.2.2

    This is a local interface, not a VSI.
  - ethernet1 – Trust zone interface, 10.1.1.1/24, NAT mode

    This is a VSI for VSD group 1.
  - ethernet4 – HA zone interface

    This is for the control link for HA communications between the two devices. You also set ethernet1 as the interface for secondary path control link communications in the event that ethernet4 fails.

You want NetScreen-A to be the master of VSD group 1, so you give it a priority of 1 and leave the priority for NetScreen-B at the default value of 100. You set the preempt hold-down time for NetScreen-A to become master after 10 seconds.

You set both devices to monitor interfaces ethernet1 and ethernet3, and assign each a weight of 255 (the default value). If either interface fails, a device-level failover occurs.

You define two default routes for each Untrust zone interface. For ethernet3 on NetScreen-A, the default route points to an external router with IP address 1.1.1.250. For ethernet3 on NetScreen-B, the default route points to an external router with IP address 1.1.2.250. All security zones are in the trust-vr routing domain.

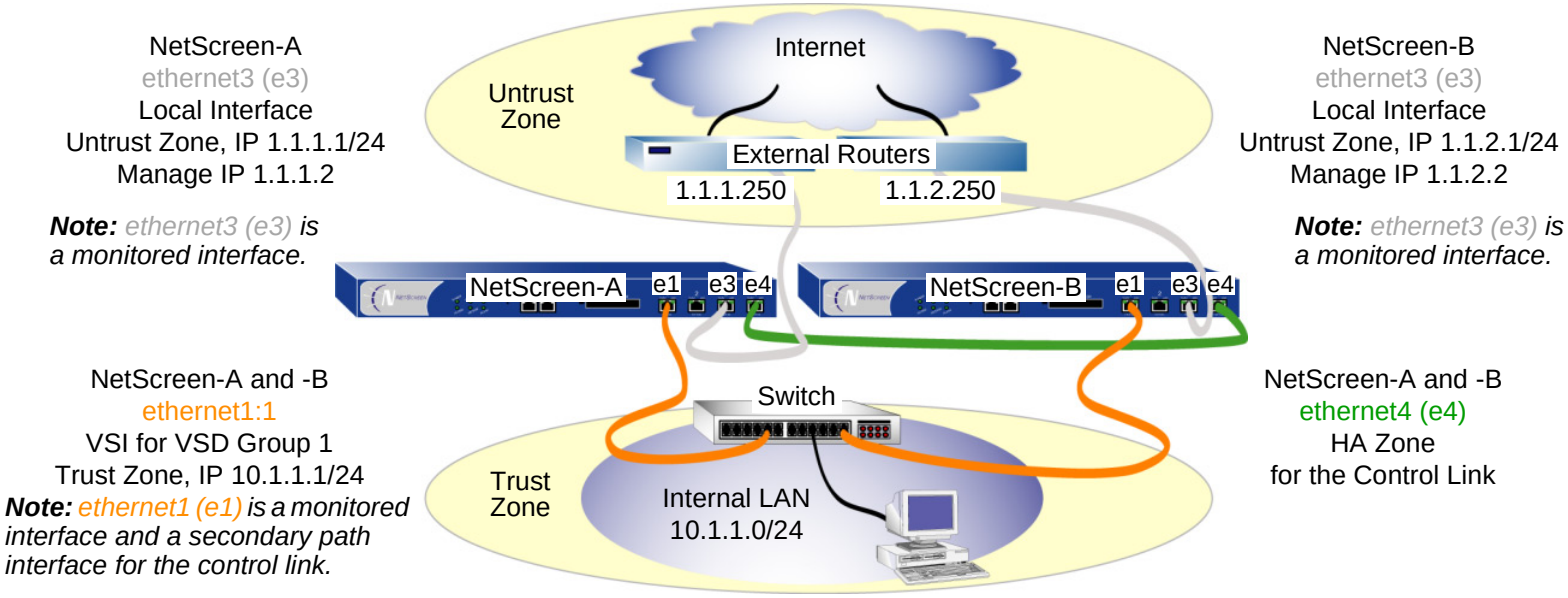

WebUI (NetScreen-A)

1. Interfaces (NetScreen-A)

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **OK**:

Zone Name: Trust

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select)

IP Address/Netmask: 1.1.1.1/24

Manage IP: 1.1.1.2

Network > Interfaces > Edit (for ethernet4): Enter the following, and then click **OK**:

Zone Name: HA

2. NSRP (NetScreen-A)

Network > NSRP > Cluster: Enter the following, and then click **Apply**:

Cluster ID: (select), 1

Network > NSRP > VSD Group: Click **Remove** for VSD group 0. When prompted to confirm the removal, click **OK**.

Note: *At the time of this release, you must use the following CLI command to create a VSD group with an ID other than 0:* ***set nsrp vsd-group id 1****. After you create a VSD group with ID 1, you can use the WebUI to modify its priority, preempt option, and hold-down time interval.*

Network > NSRP > VSD Group > Edit (for VSD Group 1): Enter the following, and then click **OK**:

Priority: 1

Enable Preempt: (select)

Preempt Hold-Down Time (sec): 10

Network > Interfaces > New VSI IF: Enter the following, and then click **OK**:

Interface Name:

VSI Base: ethernet1

VSD Group: 1

IP Address / Netmask: 10.1.1.1/24

Network > NSRP > Link: Select **ethernet1** from the Secondary Link drop-down list, and then click **Apply**.

Network > NSRP > Monitor > Interface > VSD ID: Device Edit Interface: Enter the following, and then click **Apply**:

ethernet1: (select), Weight: 255

ethernet3: (select), Weight: 255

3. Route (NetScreen-A)

Network > Routing > Routing Entries > (trust-vr) New: Enter the following, and then click **OK**:

Network Address / Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 1.1.1.250

WebUI (NetScreen-B)

4. Interfaces (NetScreen-B)

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select)

IP Address/Netmask: 1.1.2.1/24

Manage IP: 1.1.2.2

Network > Interfaces > Edit (for ethernet4): Enter the following, and then click **OK**:

Zone Name: HA

5. NSRP (NetScreen-B)

Network > NSRP > Cluster: Enter the following, and then click **Apply**:

Cluster ID: (select), 1

Network > NSRP > VSD Group: Click **Remove** for VSD group 0. When prompted to confirm the removal, click **OK**.

Note: *At the time of this release, you must use the following CLI command to create a VSD group with an ID other than 0:* ***set nsrp vsd-group id 1****.*

You must use the CLI to synchronize the configuration from NetScreen-A to NetScreen-B: ***exec nsrp sync global-config save****. Then reset the device with the CLI command* ***reset****.*

You must also use the CLI to disable configuration synchronization: ***unset nsrp config sync****.*

Network > NSRP > Link: Select **ethernet1** from the Secondary Link drop-down list, and then click **Apply**.

Network > NSRP > Monitor > Interface > VSD ID: Device Edit Interface: Enter the following, and then click **Apply**:

ethernet1: (select), Weight: 255

ethernet3: (select), Weight: 255

6. Route (NetScreen-B)

Network > Routing > Routing Entries > (trust-vr) New: Enter the following, and then click **OK**:

Network Address / Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 1.1.2.250

CLI (NetScreen-A)

1. Interfaces (NetScreen-A)

```
set interface ethernet1 zone trust
set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.1.1/24
set interface ethernet3 manage-ip 1.1.1.2
set interface ethernet4 zone ha
```

2. NSRP (NetScreen-A)

```
set nsrp cluster id 1
unset nsrp vsd-group id 0
set nsrp vsd-group id 1
set nsrp vsd-group id 1 priority 1
set nsrp vsd-group id 1 preempt hold-down 10
set nsrp vsd-group id 1 preempt
set interface ethernet1:1 ip 10.1.1.1/24
set nsrp secondary-path ethernet1
set nsrp monitor interface ethernet1
set nsrp monitor interface ethernet3
```

3. Route (NetScreen-A)

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
save
```

CLI (NetScreen-B)

4. Interfaces (NetScreen-B)

```
set interface ethernet3 zone untrust
set interface ethernet3 ip 1.1.2.1/24
set interface ethernet3 manage-ip 1.1.2.2
set interface ethernet4 zone ha
```

5. NSRP (NetScreen-B)

```
set nsrp cluster id 1
unset nsrp vsd-group id 0
set nsrp vsd-group id 1
save
exec nsrp sync global-config save
reset
```

The following prompt appears: “Configuration modified, save? [y] / n”

Press the **N** key.

The following prompt appears: “System reset, are you sure? y / [n] n”

Press the **Y** key.

The system reboots.

```
unset nsrp config sync
set nsrp secondary-path ethernet1
set nsrp monitor interface ethernet1
set nsrp monitor interface ethernet3
```

6. Route (NetScreen-B)

```
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.2.250
save
```

CONFIGURATION AND FILE SYNCHRONIZATION

When you add a new device to an active NSRP cluster, you can synchronize the configuration and files (such as PKI public/private key files) from the master of the VSD group or groups to the new device. By default, two devices synchronize configurations and files when they first enter the cluster, begin NSRP communications, and establish the roles of master and backup in VSD group 0. You can also synchronize configurations and files after a member of a cluster becomes unsynchronized for any reason.

Synchronizing Configurations

If you make any configuration changes on one device while another in the cluster reboots (or if either of the Trust zone interfaces—through which all NSRP communications pass—fails), it is possible that the configuration settings can become unsynchronized. To discover if the configuration of one device is out of sync with that of another, use the **exec nsrp sync global-config check-sum** command. The output states whether the configurations of the two devices are in or out of sync and provides the checksum of the local and remote devices.

If the configurations are out of sync, use the following command to resynchronize them: **exec nsrp sync global-config save**. If you do not use the **unset all** command on the local device before synchronizing the configurations, the local device appends the settings from the remote device to its existing settings. However, after synchronizing the configurations, every duplicate setting produces an error message. To avoid generating error messages while synchronizing configurations, you can do the following:

1. Download both the local and remote configurations to a workstation.
2. Use an application such as WinDiff to discern the differences between the files.
3. Manually enter the settings on the local device that had been added, modified, or deleted on the remote device.

Note: *Because NetScreen devices use the NetScreen Reliable Transport Protocol (NRTP), which is very similar to TCP—only more lightweight—configurations on active devices in a cluster rarely become unsynchronized.*

Synchronizing Files

If you need to synchronize a specific file, enter the following command on the device to which you are synchronizing the file: **exec nsrp sync file name** *name_str* **from peer**. If you need to synchronize all files, enter **exec nsrp sync file from peer**.

Example: Adding a Device to an Active NSRP Cluster

In this example, you add device A, which had previously been functioning as a single security appliance, to VSD groups 0 in NSRP cluster with cluster ID 1 and name “cluster1”. You unset the previous configurations on device A, reboot it, and then synchronize the configuration and files from the master VSD group 0. You then assign device A as the master of VSD group 0.

WebUI

> ***Note:*** *The configuration synchronization feature is only available through the CLI.*

CLI

Device A

`unset all`[6]

The following prompt appears: “Erase all system config, are you sure y / [n]?”

Press the **Y** key.

The system configuration is returned to the factory default settings.

`reset`

The system reboots.

6. If you do not first use the **unset all** command, the **exec nsrp sync global-config** command appends new configuration settings to existing settings. (Note: The NetScreen device generates an error message for each duplicate setting that is synchronized.)

```
set nsrp cluster id 1
set nsrp cluster name cluster1
exec nsrp sync file
exec nsrp sync global-config
exec nsrp vsd-group id 0 mode master
save
```

Disabling Configuration and File Synchronization

By default, devices placed in an NSRP cluster synchronize configurations and files. You can disable the automatic synchronization of configurations and files; for example, if you want all configuration changes to originate from NetScreen-Security Manager. To accomplish this, do either of the following:

WebUI

Network > NSRP > Synchronization: Clear the **NSRP Configuration Synchronization** check box, and then click **Apply**.

CLI

```
unset nsrp config sync
save
```

PATH MONITORING

Path monitoring checks the layer 2 and layer 3 network connections between a NetScreen interface and the interface of another device. Path monitoring is a useful tool for devices within a redundant group to determine whether the network connectivity of the device is acceptable. If the connectivity is unacceptable and passes a defined threshold, a failover occurs—either at the VSD group level or at the device level. For information about the distinction between these two levels of failover, see "Failover" on page 129.

Layer 2 path monitoring functions by checking that the physical ports are active and connected to other network devices. You can interfaces on a per-interface basis or on a per-zone basis. Per-interface:

- WebUI: Click **Network > NSRP > Monitor > Interface > VSD ID:** { **Device** | *number* } **Edit Interface**, and then select the interfaces you want the NetScreen device to monitor.
- CLI: **set nsrp** [**vsd-group id** *number*] **monitor interface** *interface*

Per zone (that is, the NetScreen device monitors all the interface bound to the selected zone):

- WebUI: Click **Network > NSRP > Monitor > Zone > VSD ID:** { **Device** | *number* } **Edit Zone**, and then select the zones you want the NetScreen device to monitor.
- CLI: **set nsrp** [**vsd-group id** *number*] **monitor zone** *zone*

Layer 3 path monitoring, or IP tracking, functions by sending ping or ARP requests to up to 16 specified IP addresses at user-determined intervals and then monitoring if the targets respond. If the total of tracked IP failures for a device acting as master (but not for the device acting as its backup) exceeds the device failover threshold, then the backup is automatically promoted to master, and the deposed master enters the inoperable state. (The inoperable VSD group member continues its IP path tracking activity. When the results no longer exceed the failover threshold, it transitions from the inoperable to initial state, and then to the backup state[7].)

Note: *When routers are grouped in a redundant cluster using the Virtual Router Redundancy Protocol (VRRP), the router functioning as the master cannot respond to ping requests to the virtual IP address if it is not the IP address owner (which might be the case after a failover). However, the master virtual router must respond to ARP requests with the virtual MAC address regardless of IP address ownership. (See RFC 2338 for details.) To use ARP when IP tracking, the polled device must be on the same physical subnet as the NetScreen manage IP address.*

7. If the VSD group is in preempt mode and the device has a better priority than the current master, it transitions from the inoperable to initial to master state.

When tracking IP addresses, you can send ping or ARP requests from a manage IP address on an interface. For VSIs, the manage IP address must be different from the interface IP address and must be unique per device. For local interfaces, the manage IP address can be the same as or different from the interface IP address.

Setting Thresholds

IP path tracking involves two kinds of thresholds: a tracked IP failure threshold and a device failover threshold.

Tracked IP Failure Threshold – The number of consecutive failures to elicit a ping or ARP response from a specific IP address required to be considered a failed attempt. Not exceeding the threshold indicates an acceptable level of connectivity with that address; exceeding it indicates an unacceptable level. You can set this threshold per IP address at any value from 1 to 200. The default value is 3.

Device Failover Threshold – The total weight of the cumulative failed attempts required to cause a VSD group master to step down. (For information on how to assign a weight to a tracked IP address, see the next section, "Weighting Tracked IP Addresses" on page 80.) You can set the device failover threshold at any value between 1 and 255. The default is 255.

Weighting Tracked IP Addresses

By applying a weight, or a value, to a tracked IP address, you can adjust the importance of connectivity to that address in relation to reaching other tracked addresses. You can assign relatively greater weights to relatively more important addresses, and lesser weights to relatively less important addresses. The assigned weights come into play when a tracked IP failure threshold is reached. For example, exceeding the tracked IP failure threshold for an address weighted 10 brings the master closer to device failover than would a tracked IP failure for an address weighted 1. You can assign weights from 1 to 255. The default is 255.

IP Tracking for VPN Tunnel Failover

By configuring a VPN tunnel on each device to reach the same remote network through two different remote VPN gateways and then tracking IP addresses at the remote site through the tunnels, you can protect VPN traffic from local and remote gateway device failure, tunnel failure, and ISP failure.

The master (Device A) actively processes VPN traffic between the local and remote networks through VPN tunnel A. Device A monitors the health of its system, its network connectivity, and VPN tunnel A.

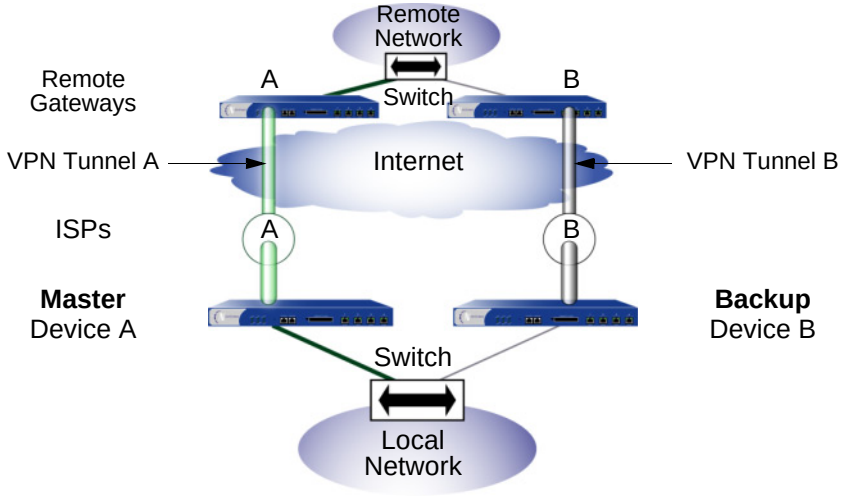

The backup (Device B) receives status reports from Device A and remains ready to become master should a failover occur. VPN tunnel B is up but inactive.

If any of the following events occur, Device B becomes master and Device A becomes backup (or Device A becomes inoperable if it has internal system problems).

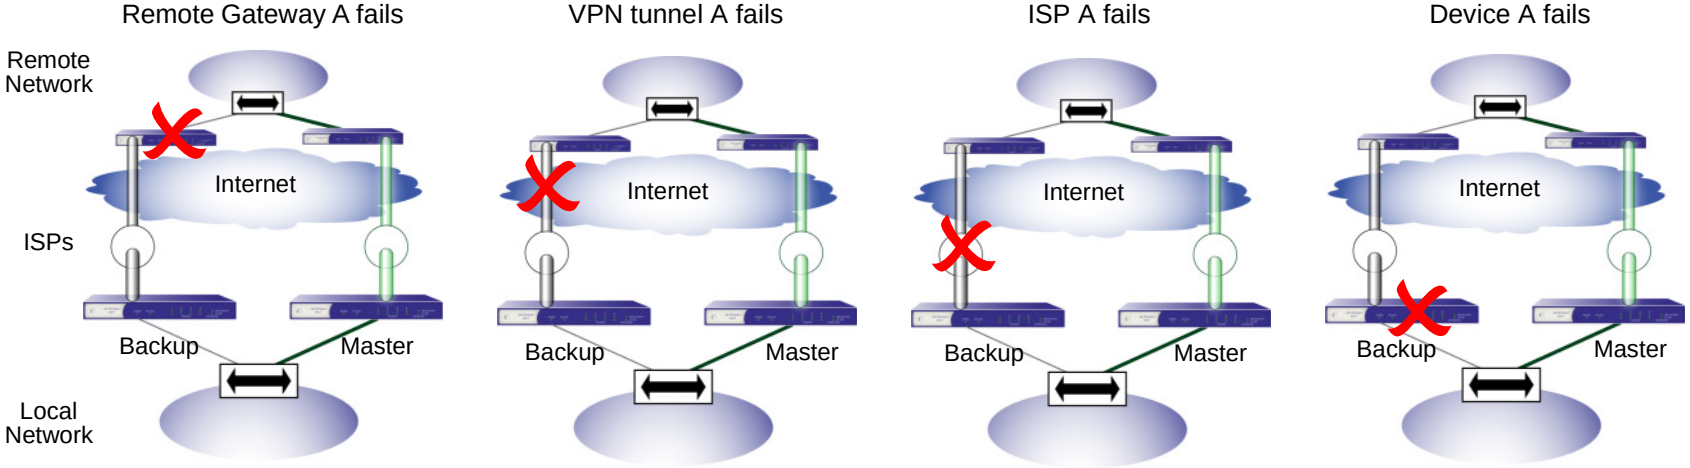

Example: IP Tracking through a VPN Tunnel

In this example, you configure two VPN tunnels[8], one for each of two NetScreen devices in an NSRP cluster. Then you configure both devices to track the IP address of two servers at the remote end of the tunnel: 10.2.2.50 and 10.2.2.60.

Note: *This example builds upon the configuration in* *"Example: Configuring NSRP-Lite" on page 69**.*

You configure each VPN tunnel as routing-based and bind it to an unnumbered tunnel interface named *tunnel.1*. Both tunnels use a preshared key (the keys are different for the two tunnels in this example, but they can also be identical). Phase 1 negotiations are in Main mode, and you enable replay protection for Phase 2 negotiations. You use the security level predefined as "Compatible" for both Phase 1 and Phase 2 proposals[9]. You also enable VPN monitoring with the rekey option. (For more information about routing-based VPN tunnels, see Volume 5, "VPNs")

The settings you define for each tracked address are as follows:

- Server at 10.2.2.50
  - Interval: 10
  - Threshold: 5
  - Weight: 16
- Server at 10.2.2.60
  - Interval: 10
  - Threshold: 5
  - Weight: 16

Not receiving a ping response after 5 consecutive attempts to one of the servers is considered a failed attempt and contributes a weighted value of 16 toward the total failover threshold.

Because the device failover threshold is 31, both tracked IP addresses must fail before a device failover occurs. If you are not willing to tolerate that amount of failure, you can lower the threshold to a more acceptable level.

8. The configuration of the two tunnels on the devices at the remote site is not included in this example.
9. The four Phase 1 compatible security level proposals are pre-g2-3des-sha, pre-g2-3des-md5, pre-g2-des-sha, and pre-g2-des-md5. The four Phase 2 compatible security level proposals are nopfs-esp-3des-sha, nopfs-esp-3des-md5, nopfs-esp-des-sha, and nopfs-esp-des-md5

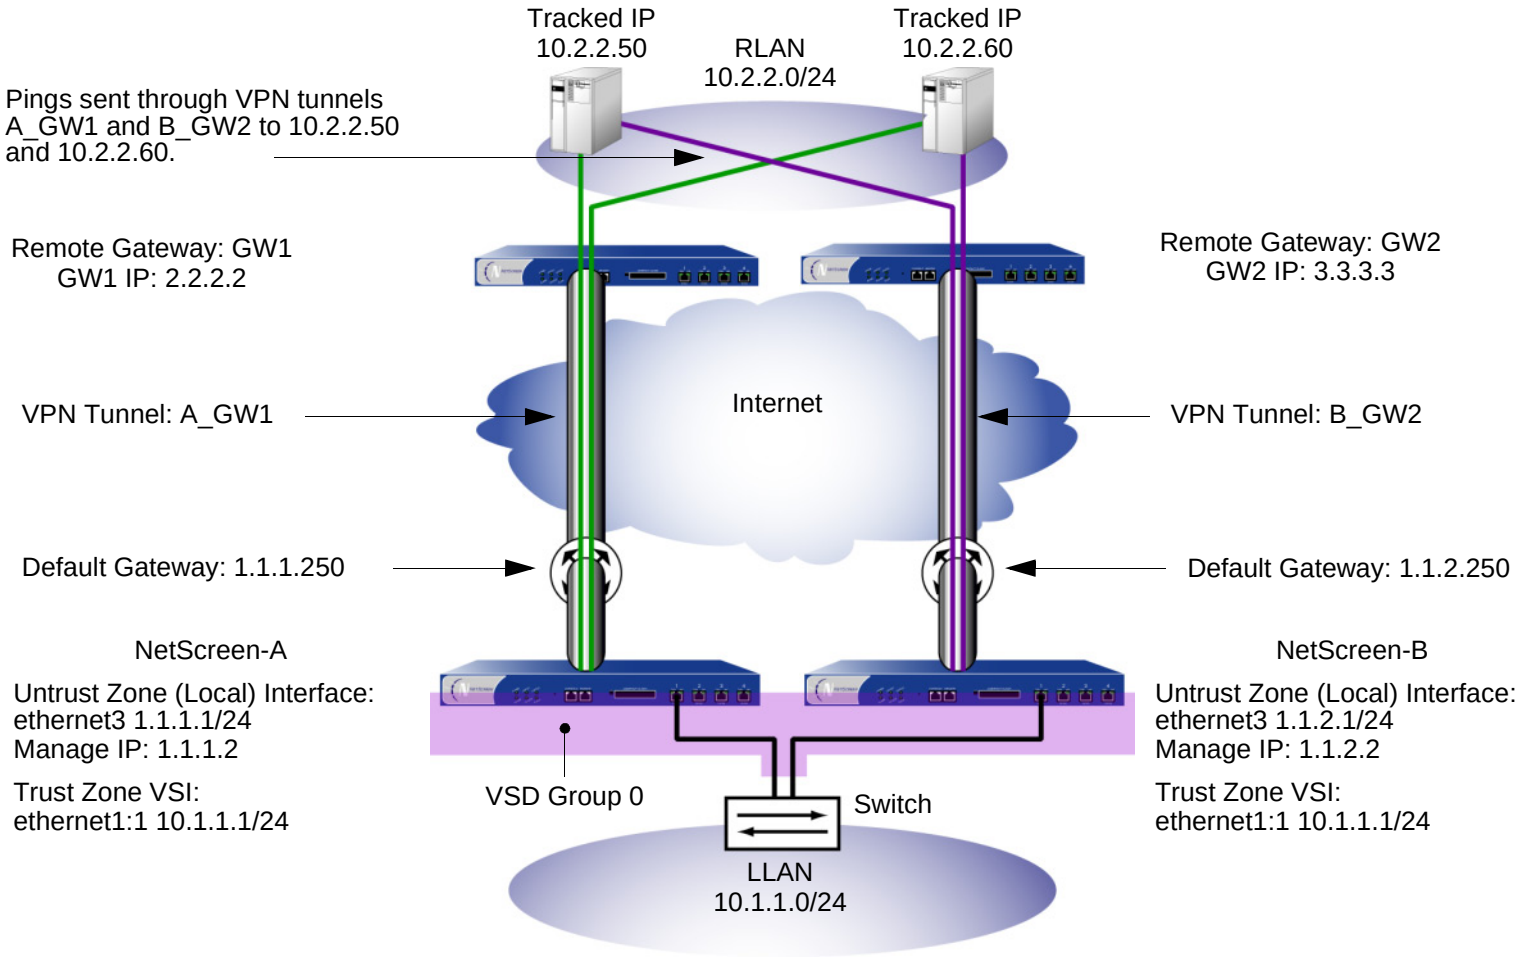

WebUI (NetScreen-A)

1. VPN Tunnel (NetScreen-A)

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: LLAN

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.0/24

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: RLAN

IP Address/Domain Name:

IP/Netmask: (select), 10.2.2.0/24

Zone: Untrust

Network > Interfaces > Tunnel IF New: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.1

Zone (VR): Untrust (trust-vr)

Unnumbered: (select)

Interface: ethernet3 (trust-vr)[10]

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: A_gw1

Security Level: Compatible

Remote Gateway: Create a Simple Gateway: (select)

Gateway Name: gw1

Type: Static IP (select), Address/Hostname: 2.2.2.2

Preshared Key: h1p8A24nG5

Security Level: Compatible

Outgoing Interface: ethernet3

10. The source interface must be in the same virtual routing domain to which the tunnel interface is bound; in this case, the trust-vr. The unnumbered tunnel interface borrows the IP address of the specified security zone interface.

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Security Level: Compatible

Replay Protection: (select)

Bind to: Tunnel Interface: tunnel.1

Proxy-ID: (select)

Local IP / Netmask: 10.1.1.0/24

Remote IP / Netmask: 10.2.2.0/24

Service: ANY

VPN Monitor: (select)

Source Interface: Default

Destination IP: 2.2.2.2

Optimization: (clear)

Rekey: (select)

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 10.2.2.0/24

Gateway: (select)

Interface: tunnel.1

Gateway IP Address: 0.0.0.0

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 1.1.1.250

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:
Address Book Entry: (select), LLAN
Destination Address:
Address Book Entry: (select), RLAN
Service: ANY
Action: Permit
Position at Top: (select)

2. IP Tracking (NetScreen-A)

Network > NSRP > Monitor > Track IP > New: Enter the following, and then click **OK**:

Track IP: 10.2.2.50
Method: Ping
Weight: 16
Interval (sec): 10
Threshold: 5
Interface: tunnel.1
VSD Group ID: Device

Network > NSRP > Monitor > Track IP > New: Enter the following, and then click **OK**:

Track IP: 10.2.2.60
Method: Ping
Weight: 16
Interval (sec): 10
Threshold: 5
Interface: tunnel.1
VSD Group ID: Device

Network > NSRP > Monitor > Track IP > Edit (for VSD: Device): Select **Enable Track IP**, and enter **31** in the Failover Threshold field.

3. VPN Tunnel (NetScreen-B)

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: LLAN

IP Address/Domain Name:

IP/Netmask: (select), 10.1.1.0/24

Zone: Trust

Objects > Addresses > List > New: Enter the following, and then click **OK**:

Address Name: RLAN

IP Address/Domain Name:

IP/Netmask: (select), 10.2.2.0/24

Zone: Untrust

Network > Interfaces > Tunnel IF New: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.1

Zone (VR): Untrust (trust-vr)

Unnumbered: (select)

Interface: ethernet3 (trust-vr)[11]

11. The source interface must be in the same virtual routing domain to which the tunnel interface is bound; in this case, the trust-vr. The unnumbered tunnel interface borrows the IP address of the specified security zone interface.

VPNs > AutoKey IKE > New: Enter the following, and then click **OK**:

VPN Name: B_gw2

Security Level: Compatible

Remote Gateway: Create a Simple Gateway: (select)

Gateway Name: gw2

Type: Static IP (select), Address/Hostname: 3.3.3.3

Preshared Key: ih38CvE3g9

Security Level: Compatible

Outgoing Interface: ethernet3

> Advanced: Enter the following advanced settings, and then click **Return** to return to the basic AutoKey IKE configuration page:

Security Level: Compatible

Replay Protection: (select)

Bind to: Tunnel Interface: tunnel.1

Proxy-ID: (select)

Local IP / Netmask: 10.1.1.0/24

Remote IP / Netmask: 10.2.2.0/24

Service: ANY

VPN Monitor: (select)

Source Interface: Default

Destination IP: 3.3.3.3

Optimization: (clear)

Rekey: (select)

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 10.2.2.0/24

Gateway: (select)

Interface: tunnel.1

Gateway IP Address: 0.0.0.0

Network > Routing > Routing Entries > trust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: ethernet3

Gateway IP Address: 1.1.2.250

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), LLAN

Destination Address:

Address Book Entry: (select), RLAN

Service: ANY

Action: Permit

Position at Top: (select)

4. IP Tracking (NetScreen-B)

Network > NSRP > Monitor > Track IP > New: Enter the following, and then click **OK**:

Track IP: 10.2.2.50
Method: Ping
Weight: 16
Interval (sec): 10
Threshold: 5
Interface: tunnel.1
VSD Group ID: Device

Network > NSRP > Monitor > Track IP > New: Enter the following, and then click **OK**:

Track IP: 10.2.2.60
Method: Ping
Weight: 16
Interval (sec): 10
Threshold: 5
Interface: tunnel.1
VSD Group ID: Device

Network > NSRP > Monitor > Track IP > Edit (for VSD: Device): Select **Enable Track IP**, and enter **31** in the Failover Threshold field.

CLI

1. VPN Tunnel (NetScreen-A)

```
set address trust LLAN 10.1.1.0/24
set address untrust RLAN 10.2.2.0/24
set interface tunnel.1 zone untrust
set interface tunnel.1 ip unnumbered interface ethernet3
set ike gateway gw1 ip 2.2.2.2 main outgoing-interface ethernet3 preshare
    h1p8A24nG5 sec-level compatible
set vpn A_gw1 gateway gw1 replay sec-level compatible
set vpn A_gw1 bind interface tunnel.1
set vpn A_gw1 proxy-id local-ip 10.1.1.0/24 remote-ip 10.2.2.0/24 any
set vpn A_gw1 monitor source-interface ethernet3 destination-ip 2.2.2.2 rekey
set vrouter trust-vr route 10.2.2.0/24 interface tunnel.1
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.1.250
set policy top from trust to untrust LLAN RLAN any permit
```

2. IP Tracking (NetScreen-A)

```
set interface tunnel.1 track-ip ip 10.2.2.50
set interface tunnel.1 track-ip ip 10.2.2.50 interval 10
set interface tunnel.1 track-ip ip 10.2.2.50 threshold 5
set interface tunnel.1 track-ip ip 10.2.2.50 weight 16
set interface tunnel.1 track-ip ip 10.2.2.60
set interface tunnel.1 track-ip ip 10.2.2.60 interval 10
set interface tunnel.1 track-ip ip 10.2.2.60 threshold 5
set interface tunnel.1 track-ip ip 10.2.2.60 weight 16
set nsrp track-ip threshold 31
set nsrp track-ip
save
```

3. VPN Tunnel (NetScreen-B)

```
unset nsrp config sync
set address trust LLAN 10.1.1.0/24
set address untrust RLAN 10.2.2.0/24
set interface tunnel.1 zone untrust
set interface tunnel.1 ip unnumbered interface ethernet3
set ike gateway gw2 ip 3.3.3.3 main outgoing-interface ethernet3 preshare
    ih38CvE3g9 sec-level compatible
set vpn B_gw2 gateway gw2 replay sec-level compatible
set vpn B_gw2 bind interface tunnel.1
set vpn B_gw2 proxy-id local-ip 10.1.1.0/24 remote-ip 10.2.2.0/24 any
set vpn B_gw2 monitor source-interface ethernet3 destination-ip 3.3.3.3 rekey
set vrouter trust-vr route 10.2.2.0/24 interface tunnel.1
set vrouter trust-vr route 0.0.0.0/0 interface ethernet3 gateway 1.1.2.250
set policy top from trust to untrust LLAN RLAN any permit
```

4. IP Tracking (NetScreen-B)

```
set interface tunnel.1 track-ip ip 10.2.2.50
set interface tunnel.1 track-ip ip 10.2.2.50 interval 10
set interface tunnel.1 track-ip ip 10.2.2.50 threshold 5
set interface tunnel.1 track-ip ip 10.2.2.50 weight 16
set interface tunnel.1 track-ip ip 10.2.2.60
set interface tunnel.1 track-ip ip 10.2.2.60 interval 10
set interface tunnel.1 track-ip ip 10.2.2.60 threshold 5
set interface tunnel.1 track-ip ip 10.2.2.60 weight 16
set nsrp track-ip threshold 31
set nsrp track-ip
save
```

3

Interface Redundancy

This chapter describes the various ways in which NetScreen devices provide interface redundancy. The chapter is divided into the following sections:

- "Redundant Interfaces" on page 94
- "Aggregate Interfaces" on page 101
- "Dual Untrust Interfaces" on page 103
  - "Interface Failover" on page 104
  - "Determining Interface Failover" on page 105
- "Serial Interface" on page 118
  - "Modem Settings" on page 119
  - "ISP Configuration" on page 121
  - "Serial Interface Failover" on page 123

REDUNDANT INTERFACES

For HA interface redundancy, NetScreen devices either provide dedicated physical redundant HA interfaces or you can bind two generic interfaces to the HA zone. See "Dual HA Interfaces" on page 38 for more information. You can also create redundant security zone interfaces, as described in this section.

Applying a similar kind of virtualization that allows a VSI to shift its binding from the physical interface on one device to the physical interface on another device, the redundant interface can shift its binding from one physical interface to another physical interface on the same device. For example, if the link from the primary interface to the switch becomes disconnected, the link fails over to the secondary interface, thereby preventing a device failover from the VSD master to backup.

You can also set a holddown time for a physical interface to wait before becoming the primary interface after an interface failover occurs. To set a holddown time for a member of a redundant interface, use the following command, in which the interface name is that of a physical interface: **set interface** *interface* **phy holddown** *number*. Note that you must enter this command before making the interface a member of a redundant group.

You can bind a VSI to any of the following interface types:

- A subinterface
- A physical interface
- A redundant interface, which in turn is bound to two physical interfaces[1]

> ***Note:*** *You cannot group subinterfaces to a redundant interface. However, you can define a VLAN on a redundant interface in the same way that you can define a VLAN on a subinterface. For information on subinterfaces and VLANs, see "Defining Subinterfaces and VLAN Tags" on page* ***7****-23.*

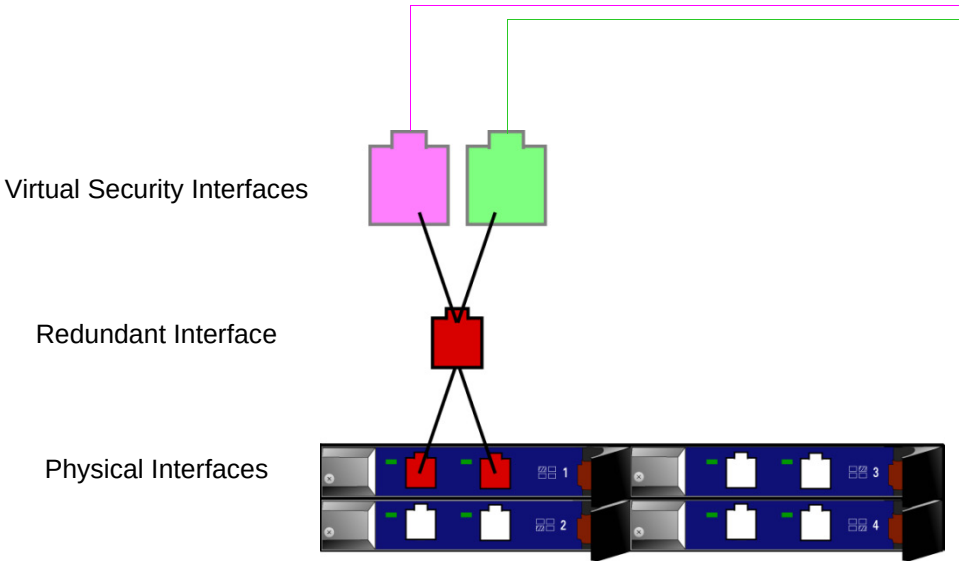

The Untrust zone VSI for VSD group 1 and the Untrust zone VSI for VSD group 2 are each bound to a redundant interface.

The redundant interface is bound to a pair of physical interfaces.

One physical interface acts as the primary interface and actively handles traffic. The other physical interface acts as the secondary interface standing by as a backup for the primary.

1. You can configure VSI on a loopback interface, but you cannot bind two loopback interfaces to a redundant interface.

Example: Creating Redundant Interfaces for VSIs

In this example, devices A and B are members of two VSD groups—VSD group 0 and VSD group 1—in an active/active configuration. Device A is the master of VSD group 0 and the backup in VSD group 1. Device B is the master of VSD group 1 and the backup in VSD group 0. The NetScreen devices are linked to two pairs of redundant switches—switches A and B in the Untrust zone, and switches C and D in the Trust zone.

Note: *This example only presents the creation of redundant interfaces on device A. Because devices A and B are members of the same NSRP cluster, device A propagates all interface configurations to device B except the manage IP address, which you enter on the redundant2 interface on both devices: device A 10.1.1.21, device B 10.1.1.22.*

You put ethernet1/1 and ethernet1/2 in redundant1, and ethernet2/1 and ethernet2/2 in redundant2. On the redundant2 interface, you define a manage IP of 10.1.1.21 for device A and a manage IP of 10.1.1.22 for device B on this interface.

The physical interfaces that are bound to the same redundant interface connect to different switches:

- Physical interfaces bound to a redundant interface in the Untrust zone: ethernet1/1 to switch A, ethernet1/2 to switch B
- Physical interfaces bound to a redundant interface in the Trust zone: ethernet2/1 to switch C, ethernet2/2 to switch D

Note: *The physical interfaces do not have to be in the same security zone as the redundant interface to which you bind them.*

By putting ethernet1/1 and ethernet2/1 in their respective redundant interfaces first, you designate them as primary interfaces. (You can change the primary status assignments via the CLI command **set interface** *redundant1* **primary** *interface1/1*.) If the link to a primary interface becomes disconnected, the NetScreen device reroutes traffic through the secondary interface to the other switch without requiring the VSD master device to fail over.

In this example, the cable from ethernet1/1 becomes disconnected, causing a port failover to ethernet1/2. Consequently, all the traffic to and from devices A and B passes through switch B. Reconnecting the cable from ethernet1/1 on device A to switch A automatically causes that interface to regain its former priority.

The IP addresses for the VSIs:

VSIs for VSD Group 0		VSIs for VSD Group 1	
redundant1	210.1.1.1/24	redundant1:1	210.1.1.2/24
redundant2	10.1.1.1/24	redundant2:1	10.1.1.2/24

Note: *IP addresses for multiple VSIs can be in the same subnet or in different subnets if the VSIs are on the same redundant interface, physical interface, or subinterface. If the VSIs are on different interfaces, they must be in different subnets.*

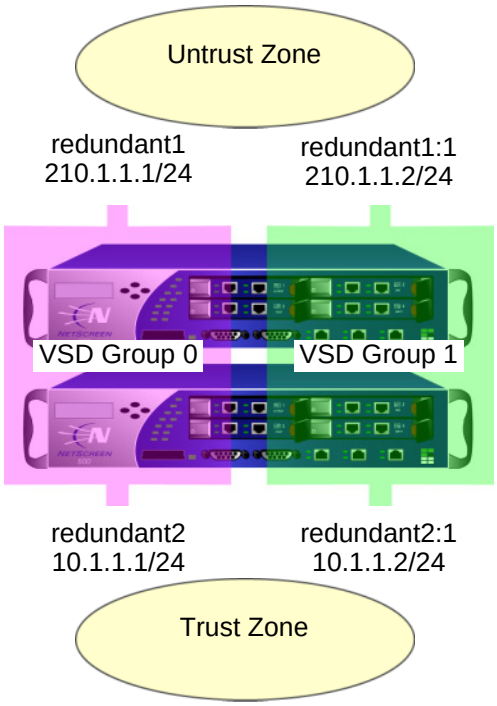

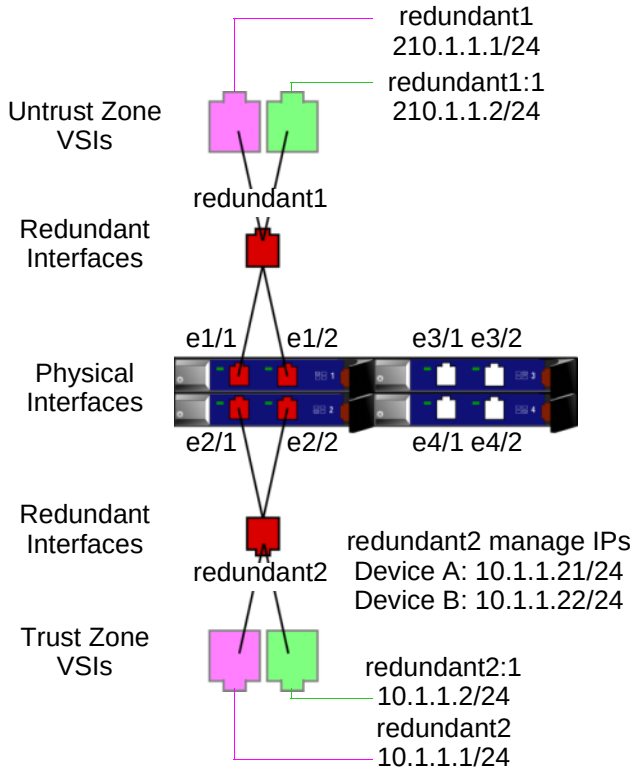

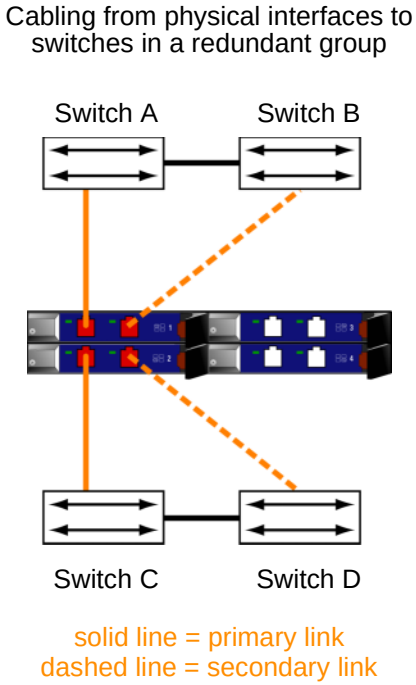

WebUI (Device A)

Redundant Interfaces

Network > Interfaces > New Redundant IF: Enter the following, and then click **OK**:

Interface Name: redundant1

Zone Name: Untrust

IP Address / Netmask: 210.1.1.1/24

Network > Interfaces > Edit (for ethernet1/1): Select **redundant1** in the “As member of” drop-down list, and then click **OK**.

Network > Interfaces > Edit (for ethernet1/2): Select **redundant1** in the “As member of” drop-down list, and then click **OK**.

Network > Interfaces > New Redundant IF: Enter the following, and then click **Apply**:

Interface Name: redundant2

Zone Name: Trust

IP Address / Netmask: 10.1.1.1/24

> Enter **10.1.1.21** in the Manage IP field, and then click **OK**.

Network > Interfaces > Edit (for ethernet2/1): Select **redundant2** in the “As member of” drop-down list, and then click **OK**.

Network > Interfaces > Edit (for ethernet2/2): Select **redundant2** in the “As member of” drop-down list, and then click **OK**.

Virtual Security Interfaces

Network > Interfaces > New VSI IF: Enter the following, and then click **OK**:

Interface Name: VSI Base: redundant1

VSD Group: 1

IP Address / Netmask: 210.1.1.2/24

Network > Interfaces > New VSI IF: Enter the following, and then click **OK**:

Interface Name: VSI Base: redundant2

VSD Group: 1

IP Address / Netmask: 10.1.1.2/24

WebUI (Device B)

Network > Interfaces > Edit (for redundant2): Type **10.1.1.22** in the Manage IP field, and then click **OK**.

Note: *You must enter static routes to addresses beyond the immediate subnet of a VSI for each VSI in each VSD. For an example showing the addition of a default route for two Untrust zone VSIs, see* *"Example: NSRP for an Active/Active Configuration" on page 49**.*

CLI (Device A)

Redundant Interfaces

```
set interface redundant1 zone untrust
set interface redundant1 ip 210.1.1.1/24

set interface ethernet1/1 group redundant1
set interface ethernet1/2 group redundant1

set interface redundant2 zone trust
set interface redundant2 ip 10.1.1.1/24
set interface redundant2 manage-ip 10.1.1.21
set interface redundant2 nat

set interface ethernet2/1 group redundant2
set interface ethernet2/2 group redundant2

set interface redundant1 primary ethernet1/1

set interface redundant2 primary ethernet2/1
```

Virtual Security Interfaces

```
set interface redundant1:1 ip 210.1.1.2/24
set interface redundant2:1 ip 10.1.1.2/24
save
```

CLI (Device B)

```
set interface redundant2 manage-ip 10.1.1.22
save
```

Note: *You must enter static routes to addresses beyond the immediate subnet of a VSI for each VSI in each VSD. For an example showing the addition of a default route for two Untrust zone VSIs, see* *"Example: NSRP for an Active/Active Configuration" on page 49*.

AGGREGATE INTERFACES

NetScreen-5000 systems allow you to combine two or more physical ports into a single virtual port. This virtual port is known as an *aggregate interface*. Only Secure Port Modules (SPMs) support this feature.

- On a 5000-8G SPM, you can create up to four aggregate interfaces.
- On a 5000-24FE SPM, you can create up to five aggregate interfaces.

The 5000-8G SPM supports only certain combinations of ports for aggregate interfaces. For example, a 5000-8G SPM residing in Slot 2 only supports the following port combinations:

- ethernet2/1 and ethernet2/2
- ethernet2/3 and ethernet2/4
- ethernet2/5 and ethernet2/6
- ethernet2/7 and ethernet2/8

You must assign one of the following names to the aggregate interface: **aggregate1**, **aggregate2**, **aggregate3**, or **aggregate4**.

Note: *As with most other ports and interfaces, you must assign the aggregate interface an IP address so that other hosts on the network can reach it.*

Example: Configuring an Aggregate Interface

In the following example, you combine two Gigabit Ethernet mini-GBIC ports, each running at 1 Gbps, into an aggregate interface aggregate1 running at 2-Gbps. The aggregate interface consists of Ethernet ports 1 and 2 on a 5000-8G SPM (residing in Slot 2) and is bound to the Trust zone.

Note: *To see the physical ports that are available on the system, go to the Network > Interfaces screen in the WebUI or enter the CLI command* ***get interface****.*

WebUI

Network > Interfaces > Aggregate IF > New: Enter the following, and then click **Apply**:

Interface Name: aggregate1

Zone Name: Trust (select)

IP Address / Netmask: 10.1.1.0/24

Enter the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet2/1): Enter the following, and then click **OK**:

As member of: aggregate1 (select)

Network > Interfaces > Edit (for ethernet2/2): Enter the following, and then click **OK**:

As member of: aggregate1 (select)

CLI

```
set interface aggregate1 zone trust
set interface aggregate1 ip 10.1.1.0/24
set interface aggregate1 nat

set interface ethernet2/1 aggregate aggregate1
set interface ethernet2/2 aggregate aggregate1
save
```

DUAL UNTRUST INTERFACES

You can select a *port mode* for some NetScreen appliances. The port mode automatically sets different port, interface, and zone bindings for the device. Certain port modes bind a second, backup interface to the Untrust zone (see “Port Modes” on page **2**-55). For these port modes, the backup interface is used only when there is a failure on the connection through the primary interface or when you manually force traffic from the primary interface to the backup. For example, on the NetScreen-5XT, the Dual Untrust and Combined port modes provide a backup interface to the Untrust zone.

Dual Untrust

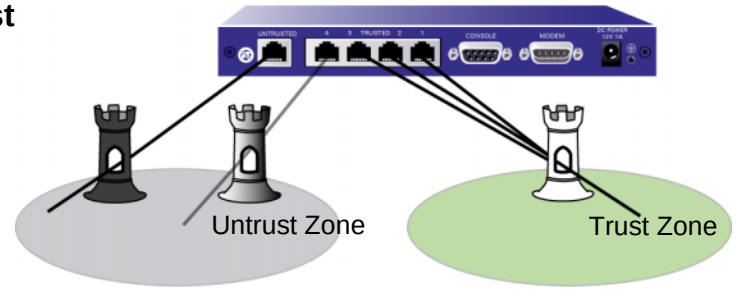

Combined

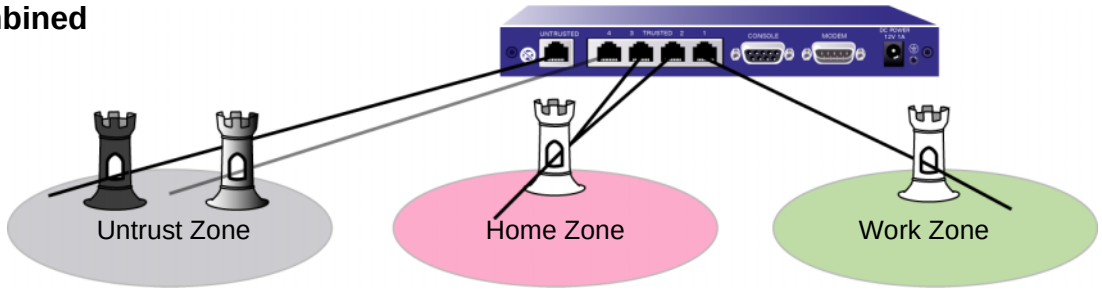

Interface Failover

When there are both primary and backup interfaces bound to the Untrust zone (see “Setting the Port Mode on NetScreen Appliances” on page **2**-59), you can manually force traffic from the primary interface to the backup interface through the WebUI or the CLI. You can also configure the NetScreen device to automatically forward traffic to the backup interface if ScreenOS detects a failure on the primary interface connection.

Example: Manually Forcing Traffic from the Primary to the Backup Interface

To force traffic from the primary interface to the backup interface:

WebUI

Network > Untrust Failover: Click **Force to Failover**.

CLI

```
exec failover force
```

When the primary interface is again available, you need to use the WebUI or the CLI to switch traffic from the backup to the primary interface.

Example: Manually Forcing Traffic from the Backup to the Primary Interface

To force traffic from the backup interface to the primary interface:

WebUI

Network > Untrust Failover: Click **Force to Revert**.

CLI

```
exec failover revert
```

Example: Automatically Switching Traffic between the Primary and Backup Interface

You can configure the NetScreen device to automatically switch traffic to the backup interface if ScreenOS detects a failure on the primary interface connection. By default, there is a 30-second interval before the switchover occurs. In automatic interface failover mode, when the connection through the primary interface is restored, ScreenOS automatically switches traffic from the backup interface to the primary.

To configure ScreenOS for automatic interface failover:

WebUI

Network > Untrust Failover: Select **Automatic Failover**, and then click **Apply**.

CLI

```
set failover auto
save
```

Determining Interface Failover

An interface failover can occur when ScreenOS detects a physical link problem on the primary interface connection, such as an unplugged cable. You can also define the following types of interface failover:

- When certain IP addresses become unreachable through a given interface using IP tracking
- When certain VPN tunnels on the primary untrust interface become unreachable using VPN tunnel monitoring

Interface Failover with IP Tracking

You can specify an interface failover when certain IP addresses become unreachable through a given interface, even if the physical link is still active. ScreenOS uses layer 3 path monitoring, or *IP tracking*, similar to that used for NSRP, to monitor the connection through the primary interface. For example, if an interface connects directly to a router, you can track the next-hop address on the interface to determine if the router is still reachable. Note that you can configure IP tracking without configuring automatic interface failover.

You can configure up to four IP addresses for ScreenOS to track. For each IP address to be tracked, you specify the following:

- Interval, in seconds, at which the pings are sent to the specified IP address.
- Number of consecutive unsuccessful ping attempts before the connection to the IP address is considered failed.
- Weight of the failed IP connection (once the sum of the weights of all failed IP connections crosses a specified threshold, ScreenOS initiates a switchover to the backup link).
- IP address of the next-hop (gateway) to be used to reach the tracked IP address. You must specify the gateway if the tracked IP address is on a different subnet. If you do not specify a gateway address, ScreenOS uses the default route for the interface to reach the tracked IP address.

Instead of tracking specific IP addresses, you can configure ScreenOS to track the interface's default gateway. The dynamic next-hop option is useful with L2TP or DHCP.

Note that when you configure an IP address for ScreenOS to track, a host route for that IP address is not added to the routing table.

There are two types of configurable thresholds in tracking IP addresses:

- Track IP address failure threshold — The number of consecutive failures to elicit a ping response from a specific IP address that constitutes a failure. Not exceeding the threshold indicates an acceptable level of connectivity with the address; exceeding the threshold indicates an unacceptable level. You set this threshold for each IP address at any value between 1 and 200. The default value is 3.
- Interface failover threshold — The total weight of the cumulative failed attempts to reach IP addresses on the interface that constitutes an interface failure. You can set this threshold at any value between 1 and 255. The default value is 1, which means any failure to reach an IP address can cause an interface failover.

By applying a *weight*, or a value, to a tracked IP address, you can adjust the importance of connectivity to that address in relation to reaching other tracked addresses. You can assign comparatively greater weights to relatively more important addresses, and less weights to relatively less important addresses. Note that the assigned weights only come into play when a tracked IP address failure threshold is reached. For example, failure of a tracked IP address with a weight of 10 brings the interface closer to a failover than would the failure of a tracked IP address with a weight of 1. You can assign weights from 1 to 255. The default weight is 1.

Example: Configuring Automatic Failover with IP Tracking

In this example, you first configure the NetScreen-5XT for Dual Untrust mode. You then configure the device for automatic failover. If automatic failover from the primary interface to the backup interface occurs, the backup interface carries all traffic to and from the Untrust zone until the primary interface is restored. For the primary interface, ScreenOS monitors three IP addresses to determine when failover occurs; each tracked IP address has the following weight:

- 100.100.100.100 6
- 200.200.200.200 4
- 210.210.210.210 3

For each of the above tracked IP addresses, the failure threshold is the default value 3. This means that if ScreenOS is unable to obtain a ping response from 100.100.100.100 three or more consecutive times, it considers the IP address unreachable through the primary interface. For the primary interface, failover occurs when the interface failover threshold reaches 10. This means that if both IP addresses 100.100.100.100 and 200.200.200.200 become

unreachable through the primary interface, the cumulative weights of the failures equal 10, which causes an automatic failover to the backup interface. Note that if IP addresses 200.200.200.200 and 210.210.210.210 both become unreachable through the primary interface, the cumulative weights of the failures equal 7, and no failover occurs.

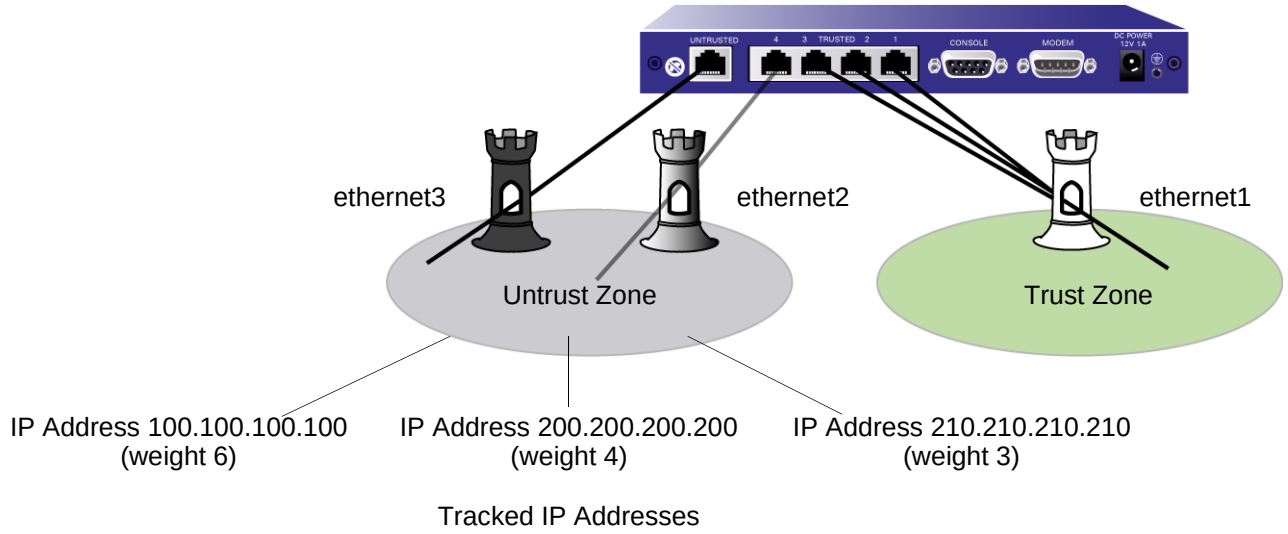

WebUI

1. Port Mode

Configuration > Port Mode: Select **Dual-Untrust** from the drop-down list, and then click **Apply**.

The following prompt appears:

Operational mode change will erase current configuration and reset the device, continue?

Click **OK**, which causes the NetScreen device to reboot.

2. Log In and Interfaces

Log in again, and set the interface IP addresses. Then continue with the following configuration:

3. Automatic Failover and IP Tracking

Network > Untrust Failover: Select **Automatic Failover**, and then click **Apply**.

Network > Interfaces > Edit (for ethernet3) > Track IP: Enter the following, and then click **Apply**:

Track IP: 100.100.100.100

Weight: 6

Enter the following, and then click **Apply**:

Track IP: 200.200.200.200

Weight: 4

Enter the following, and then click **Apply**:

Track IP: 210.210.210.210

Weight: 3

Network > Interface > Edit (for ethernet3) > Track IP Options: Enter the following, and then click **OK**:

Enable Track IP: (select)

Failover Threshold: 10

CLI

1. Port Mode

```
exec port-mode dual-untrust
```

The following prompt appears:

```
Change port mode from <trust-untrust> to <dual-untrust> will erase system
    configuration and reboot box
Are you sure y/[n] ?
```

Press the **Y** key, which causes the NetScreen device to reboot.

2. Log In and Interfaces

Log in again, and set the interface IP addresses. Then continue with the following configuration:

3. Automatic Failover and IP Tracking

```
set interface failover auto
set interface ethernet3 track-ip
set interface ethernet3 track-ip threshold 10
set interface ethernet3 track-ip ip 100.100.100.100 weight 6
set interface ethernet3 track-ip ip 200.200.200.200 weight 4
set interface ethernet3 track-ip ip 210.210.210.210 weight 3
save
```

Interface Failover with VPN Tunnel Monitoring

You can specify an interface failover when certain VPN tunnels on the primary interface are determined to be "down." For each VPN tunnel, you specify a failover weight, in percent. The assigned weights only come into play when the status of one or more monitored tunnels is "down". If the cumulative weight of the down VPN tunnels reaches or exceeds 100%, ScreenOS fails over to the backup interface.

By applying a *weight*, or a value, to a VPN tunnel, you can adjust the importance of the tunnel status in relation to other tunnels. You can assign comparatively greater weight to relatively more important tunnels, and less weight to relatively less important tunnels. Note that the accumulated weights of *all* monitored VPN tunnels determine when interface failover occurs. For example, failure of a VPN tunnel with a weight of 50 brings the primary interface closer to a failover than would the failure of a VPN tunnel with a weight of 10. Also note that tunnels that are in "inactive," "ready," or undetermined state are counted as 50% of the assigned weight. That is, if you assign a weight of 50 to a tunnel that is in inactive state, the tunnel's weight that is counted toward interface failover is 25.

If failover to the backup interface occurs, ScreenOS can still try to establish new VPN tunnel(s) on the primary interface if the VPN monitor rekey feature is enabled. If one or more VPN tunnels on the primary interface returns to "up" status so that the accumulated failover weight is less than 100%, ScreenOS can revert traffic back to the primary interface. Enable the VPN monitor rekey feature to allow ScreenOS to switch traffic from the backup interface to the primary.

Example: Configuring Automatic Failover with VPN Tunnel Monitoring

In this example, you first configure the NetScreen-5XT for Dual Untrust mode. You then configure three VPN tunnels with the primary Untrust zone interface (ethernet3) as the outgoing interface. For the primary interface, the NetScreen device monitors three VPN tunnels to determine when a failover occurs. Each VPN tunnel has the following failover weight:

- to_remote1 60
- to_remote2 40
- to_remote3 40

You also configure the device for automatic failover. If automatic failover from the primary interface to the backup interface occurs, the backup interface carries all traffic to and from the Untrust zone until the primary interface is restored. Primary interface failover occurs when the cumulative failover weight reaches or exceeds 100%. This means that if both to_remote1 and to_remote2 are down, the cumulative weight of the failures would be 100%, which would cause an automatic failover to the backup interface. Note that if only to_remote2 and to_remote3 are down, the cumulative weight of the failures would be 80%, and no failover would occur.

In this example, you also enable the VPN monitor rekey feature. In the event of a failover, this feature allows the NetScreen device to revert traffic from the backup interface to the primary if the accumulated weight of the VPN tunnels on the primary interface becomes less than 100%.

WebUI

1. Port Mode

Configuration > Port Mode: Select **Dual-Untrust** from the drop-down list, and then click **Apply**.

The following prompt appears:

> Operational mode change will erase current configuration and reset the device, continue?

Click **OK**, which causes the NetScreen device to reboot.

2. Log In and Interfaces

Log back in to the NetScreen device. Then continue with the following configuration:

Network > Interfaces > Edit (for ethernet1): Enter the following, and then click **Apply**:

Static IP: (select)

IP Address/Netmask: 10.1.1.1/24

Enter the following, and then click **OK**:

Interface Mode: NAT

Network > Interfaces > Edit (for ethernet2): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select)

IP Address/Netmask: 1.2.2.1/24

Network > Interfaces > Edit (for ethernet3): Enter the following, and then click **OK**:

Zone Name: Untrust

Static IP: (select)

IP Address/Netmask: 1.1.1.1/24

Network > Interfaces > New Tunnel IF: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.1

Zone (VR): Untrust (trust-vr)

Fixed IP: (select)

IP Address / Netmask: 1.1.1.1/24

Network > Interfaces > New Tunnel IF: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.2

Zone (VR): Untrust (trust-vr)

Fixed IP: (select)

IP Address / Netmask: 2.2.2.2/24

Network > Interfaces > New Tunnel IF: Enter the following, and then click **OK**:

Tunnel Interface Name: tunnel.3

Zone (VR): Untrust (trust-vr)

Fixed IP: (select)

IP Address / Netmask: 3.3.3.3/24

3. VPN Tunnels

VPNs > AutoKey Advanced > Gateway > New: Enter the following, and then click **OK**:

Gateway Name: remote_a

Security Level: Basic

Remote Gateway Type:

Static IP Address: (select), Address/Hostname: 4.4.4.4

Preshared Key: netscreen1

Outgoing Interface: Untrust

VPNs > Autokey IKE > New: Enter the following, and then click **OK**:

VPN Name: to_remote1

Security Level: Basic

Remote Gateway:

Predefined: (select), remote_a

> Advanced: Enter the following advanced settings, and then click Return to return to the basic AutoKey IKE configuration page:

Bind to: Tunnel Interface: (select), tunnel.1

VPN Monitor: (select)

Rekey: (select)

VPNs > Autokey IKE > New: Enter the following, and then click **OK**:

VPN Name: to_remote2

Security Level: Basic

Remote Gateway:

Predefined: (select), remote_a

> Advanced: Enter the following advanced settings, and then click Return to return to the basic AutoKey IKE configuration page:

Bind to: Tunnel Interface: (select), tunnel.2

VPN Monitor: (select)

Rekey: (select)

VPNs > Autokey IKE > New: Enter the following, and then click **OK**:

VPN Name: to_remote3

Security Level: Basic

Remote Gateway:

Predefined: (select), remote_a

> Advanced: Enter the following advanced settings, and then click Return to return to the basic AutoKey IKE configuration page:

Bind to: Tunnel Interface: (select), tunnel.3

VPN Monitor: (select)

Rekey: (select)

4. Tunnel Failover

Network > Untrust Failover > Select the following, and then click **Apply**:

Failover Type: Tunnel Interface (select)

Automatic Failover (select)

Network > Untrust Failover > Edit Weight: Enter the following, and then click **Apply**:

VPN to_remote1 (bound to tunnel interface tunnel.1) weight: 60

VPN to_remote2 (bound to tunnel interface tunnel.2) weight: 40

VPN to_remote3 (bound to tunnel interface tunnel.3) weight: 40

CLI

1. Port Mode

```
exec port-mode dual-untrust
```

The following prompt appears:

```
Change port mode from <trust-untrust> to <dual-untrust> will erase system
    configuration and reboot box
Are you sure y/[n] ?
```

Press the **Y** key, which causes the NetScreen device to reboot.

2. Log In and Interfaces

Log back in to the NetScreen device. Then continue with the following configuration:

```
set interface ethernet1 ip 10.1.1.1/24
set interface ethernet1 nat

set interface ethernet2 ip 1.2.2.1/24

set interface ethernet3 ip 1.1.1.1/24

set interface tunnel.1 zone untrust
```

```
set interface tunnel.1 ip 1.1.1.1/24

set interface tunnel.2 zone untrust
set interface tunnel.2 ip 2.2.2.2/24

set interface tunnel.3 zone untrust
set interface tunnel.3 ip 3.3.3.3/24
```

3. VPN Tunnels

```
set ike gateway remote_a ip 4.4.4.4 outgoing-interface ethernet3 preshare
    netscreen1 sec-level basic

set vpn to_remote1 gateway remote_a sec-level basic
set vpn to_remote1 bind interface tunnel.1
set vpn to_remote1 monitor rekey

set vpn to_remote2 gateway remote_a sec-level basic
set vpn to_remote2 bind interface tunnel.2
set vpn to_remote2 monitor rekey

set vpn to_remote3 gateway remote_a sec-level basic
set vpn to_remote3 bind interface tunnel.3
set vpn to_remote3 monitor rekey
```

4. Tunnel Failover

```
set failover type tunnel-if
set failover auto
set vpn to_remote1 failover-weight 60
set vpn to_remote2 failover-weight 40
set vpn to_remote3 failover-weight 40
save
```

SERIAL INTERFACE

You can connect an external modem to the RS-232 serial port on certain NetScreen devices to allow the device to establish a PPP connection to an ISP. This provides a dial-up backup interface for traffic to the Untrust zone if there is a failure on the connection through the primary interface. The dial backup feature is enabled by default for the Trust-Untrust and Home-Work port modes (see “Port Modes” on page **2** -55).

The dial backup feature allows two interfaces to the Untrust zone:

- The primary physical interface is the Untrusted Ethernet port. In ScreenOS, the primary logical interface is the Untrust interface in Trust-Untrust port mode and the ethernet3 interface in the Home-Work port mode.
- The backup physical interface is the modem port. In ScreenOS, the backup interface is the serial interface in either Trust-Untrust or Home-Work port modes. By default, the serial interface is bound to the Null zone and you need to bind it to the Untrust zone to use it as the backup interface.

You configure ScreenOS to dial through the modem to an existing ISP account when traffic is switched to the serial interface. When a switch to the serial interface occurs, the modem does not dial unless there is traffic[2] to be sent or the modem idle timeout is set to 0. ScreenOS can queue up to 16 packets while the dial-up link is brought up, so there is minimal data loss when traffic is switched to the serial interface.

By default, interface failover on the NetScreen device is manual. With manual failover, you need to force ScreenOS to switch traffic from one interface to the other using the CLI or WebUI. When the primary interface is again available, you need to use the CLI or WebUI to direct ScreenOS to switch traffic from the backup to the primary interface.

The NetScreen device can automatically fail over to the serial interface, including dialing and authenticating to a pre-existing ISP account. When the connection through the primary interface is restored, ScreenOS can automatically switch traffic from the serial interface back to the primary interface.

2. Only policy-enabled through (user-generated) traffic causes the modem to dial. Management or routing protocol related messages such as OSPF hellos do not cause modem dialup.

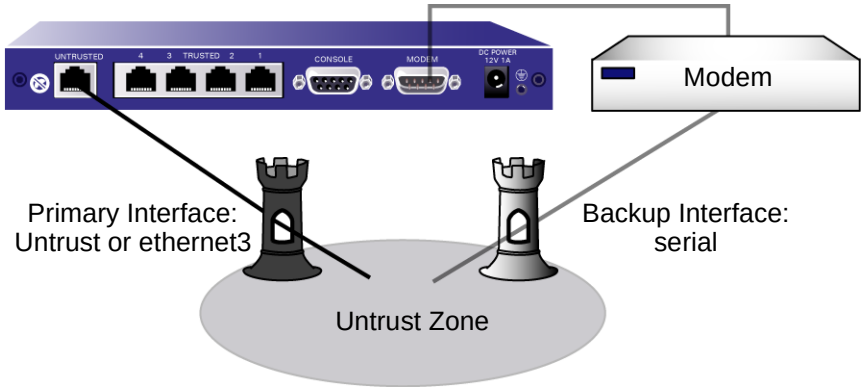

Modem Settings

The modem you use for the dial-up connection must support the following features:

- Hardware flow control
- Provide clear to send (CTS) signals
- Able to respond to request to send (RTS) signals
- Asynchronous only
- Support AT command set

You can configure the following serial link parameters in ScreenOS:

- The maximum amount of time that the serial link can be idle before ScreenOS automatically disconnects the modem (the default is 10 minutes)
- The number of times ScreenOS retries the dial-up connection if the line is busy or there is no response (the default is 3 times)
- The interval, in seconds, between dial-up retries (the default is 10 seconds)
- The maximum baud rate for the serial link (the default rate is 115200 bps)

ScreenOS uses a default modem initialization string. You can configure up to four modem initialization strings, but you can activate only one of the configured initialization strings at a time. The modem initialization string must meet the following requirements:

- Hardware flow control is recommended, but not required (you can specify no flow control)
- Software flow control is not used
- Result code must be displayed in verbal mode

Example: Configuring Modem Settings

In this example, you configure the modem idle time to be 20 minutes. You also define a modem initialization string for a new modem setting, *mod1*, and activate it.

WebUI

Network > Interfaces > Edit (for serial) > Modem: Enter the following, and then click **OK**:

Modem Name: mod1

Init String: AT&FS7=255S32=6

Status: Enable (select)

Inactivity Timeout: 20

CLI

```
set modem idle-time 20
set modem settings mod1 init-strings AT&FS7=255S32=6
set modem settings mod1 active
save
```

ISP Configuration

You configure the NetScreen device to dial to an ISP account if a failover to the serial interface occurs and there is traffic to be sent. You can configure up to four ISP connections, assigning each a different priority number (1 is the highest priority). The priority number determines the order that ScreenOS uses in attempting the dial-up connection; ScreenOS dials up the ISP with the highest priority first. If ScreenOS is unable to log in to the ISP account with the highest priority, it will dial the ISP with the next highest priority number, and so on, until there are no more ISP configurations.

Note: *By default, ScreenOS attempts to dial to a configured ISP account up to three times (see* *"Modem Settings" on page 119* *for information on modem parameters). If ScreenOS is not able to connect to any configured ISP account, it sends a connect fail message and waits until the primary interface is available again.*

For each ISP configuration, you specify the following:

- Account login and password.[3]
- Primary phone number and, optionally, an alternate phone number. If the modem uses pulse dial by default but you want to use tone dial, precede the phone number with a **T**. If the modem uses tone dial by default but you want to use pulse dial, precede the phone number with a **P**.
- Priority for this connection, relative to other configured ISP connections.

3. The ISP account must be a standard Point-to-Point Protocol (PPP) account that only requires a username and password for login.

Example: Configuring ISP Information

In this example, you configure information for two different ISP accounts: the *isp1* account has a priority value of 1, while the *isp2* account has a priority value of 2. This means that ScreenOS will always dial up the isp1 account first if failover to the serial interface occurs.

WebUI

Network > Interfaces > Edit (for serial) > ISP: Enter the following, and then click **OK**:

ISP Name: isp1

Primary Number: 4085551111

Alternative Number: 4085552222

Login Name: kgreen

Login Password: 98765432

Priority: 1

Network > Interfaces > Edit (for serial) > ISP: Enter the following, and then click **OK**:

ISP Name: isp2

Primary Number: 4085551212

Alternative Number:

Login Name: kgreen

Login Password: 12345678

Priority: 2

CLI

```
set modem isp isp1 account login kgreen password 98765432
set modem isp isp1 primary-number 4085551111 alternative-number 4085552222
set modem isp isp1 priority 1
set modem isp isp2 account login kgreen password 12345678
set modem isp isp2 primary-number 4085551212
set modem isp isp2 priority 2
save
```

Serial Interface Failover

By default, you must use the WebUI or CLI to force ScreenOS to switch over to the serial interface when the primary interface (Untrust or ethernet3 interface) connection fails and to switch back to the primary interface when the primary is again available. You can configure the interface failover to be automatic. You can also configure IP tracking to monitor failure on the Untrust or ethernet3 interfaces. See “Interface Failover” on page 104 for more information.

By default, policies that are enabled for traffic from the Trust zone to the Untrust zone or from the Untrust zone to the Trust zone are still active after a failover to the serial interface. But traffic through the primary interface could be so heavy that it cannot be handled by the dialup link. When you define a policy, you can specify whether or not the policy should be active if ScreenOS switches to the serial interface. See “Example: Specifying a Policy as Inactive for Serial Interface Failover” on page 128 for information on how to configure this in the WebUI and the CLI.

The serial interface is bound by default to the Null zone and you need to explicitly bind it to the Untrust zone to use it as a backup interface. If you bind the serial interface to the Untrust zone using the WebUI, ScreenOS automatically adds a default route for the serial interface. If you bind the serial interface to the Untrust zone using the CLI, ScreenOS does *not* add a default route to the serial interface and you must explicitly add a default route for the serial interface if traffic is to be routed through the serial interface. See “Example: Deleting a Default Route for the Serial Interface” on page 127 for information on how to configure this in the WebUI and the CLI.

Example: Configuring Dial Backup in the Trust-Untrust Mode

In this example, you first bind the serial interface to the Untrust zone. The serial interface becomes the backup interface to the primary (the Untrust interface). You then configure ScreenOS to automatically fail over to the serial interface when the primary interface connection fails.

You configure IP tracking to determine failure of the primary interface—if IP addresses 100.100.100.100 and 200.200.200.200 become unreachable through the primary interface, ScreenOS automatically switches over to the backup interface.

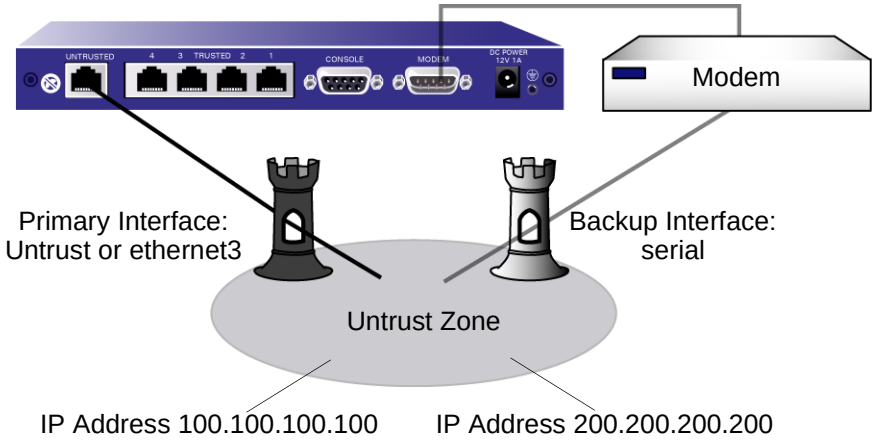

WebUI

Network > Interfaces > Edit (for serial): Enter the following, and then click **OK**:

Zone Name: (select) Untrust

Network > Interfaces > Edit (for serial) > Modem: Enter the following, and then click **OK**:

Modem Name: mod1

Init String: AT&FS7=255S32=6

Inactivity Timeout: 20

Network > Interfaces > Edit (for serial) > ISP: Enter the following, and then click **OK**:

ISP Name: isp1

Primary Number: 4085551111

Alternative Number: 4085552222

Login Name: kgreen

Login Password: 98765432

Priority: 1

Network > Interfaces > Edit (for serial) > ISP: Enter the following, and then click **OK**:

ISP Name: isp2

Primary Number: 4085551212

Alternative Number:

Login Name: kgreen

Login Password: 12345678

Priority: 2

Network > Untrust Failover > Automatic Failover: (select), and then click **Apply**.

Network > Interface > Edit (for ethernet3) > Track IP: Enter the following, and then click **Apply**:

Track IP: 100.100.100.100

Weight: 6

Enter the following, and then click **Apply**:

Track IP: 200.200.200.200

Weight: 4

Enter the following, and then click **Apply**:

Track IP: 210.210.210.210

Weight: 3

Network > Interface (ethernet3) > Edit > Track IP Options: Enter the following, and then click **OK**:

Enable Track IP: (select)

Failover Threshold: 10

CLI

```
set interface serial zone untrust
set failover auto

set modem idle-time 20
set modem settings mod1 init-strings AT&FS7=255S32=6
set modem settings mod1 active
set modem isp isp1 account login kgreen password 98765432
set modem isp isp1 primary-number 4085551111 alternative-number 4085552222
set modem isp isp1 priority 1
set modem isp isp2 account login kgreen password 12345678
set modem isp isp2 primary-number 4085551212
set modem isp isp2 priority 2

set interface ethernet3 track-ip
set interface ethernet3 track-ip threshold 10
set interface ethernet3 track-ip ip 100.100.100.100 weight 6
set interface ethernet3 track-ip ip 200.200.200.200 weight 4
set interface ethernet3 track-ip ip 210.210.210.210 weight 3
save
```

Example: Deleting a Default Route for the Serial Interface

If you bind the serial interface to the Untrust zone using the WebUI, ScreenOS automatically adds a default route for the serial interface. In this example, you use the WebUI to bind the serial interface to the Untrust zone. You then delete the default route that has been automatically created for the serial interface.

WebUI

Network > Interfaces > Edit (for serial): Enter the following, and then click **OK**:

Zone Name: (select) Untrust

Network > Routing > Routing Entries: In the Configure column, click **Remove** for the default route to 0.0.0.0/0 through the serial interface.

Example: Adding a Default Route for the Serial Interface

If you bind the serial interface to the Untrust zone using the CLI, ScreenOS does *not* add a default route to the serial interface and you must explicitly add a default route for the serial interface if you want the NetScreen device to route traffic through the serial interface. In this example, you use the CLI to bind the serial interface to the Untrust zone. You then add a default route for the serial interface, which is bound to the Untrust zone.

CLI

```
set interface serial zone untrust
set route 0.0.0.0/0 interface serial
save
```

Example: Specifying a Policy as Inactive for Serial Interface Failover

In this example, normal traffic through the primary interface (ethernet3) to the Untrust zone includes large files transferred via FTP from host22 in the Trust zone to ftp_srv in the Untrust zone. If a failover to the serial interface occurs, the dialup link might drop such large FTP traffic. Whenever there is a failover to the serial interface, any policy that is configured to be inactive for the serial interface becomes invalid and the policy lookup procedure continues to the next policy.

WebUI

Policies > (From: Trust, To: Untrust) New: Enter the following, and then click **OK**:

Source Address:

Address Book Entry: (select), host22

Destination Address:

Address Book Entry: (select), ftp_srv

Service: FTP

Action: Permit

> Advanced: Clear **Valid for Serial**, and then click **Return** to set the advanced options and return to the basic configuration page.

CLI

```
set policy from trust to untrust host22 ftp_srv ftp permit no-session-backup
save
```

4

Failover

Redundancy ensures that the functions of a certain component can still be performed even if the primary component becomes unavailable. NetScreen features, such as NSRP, provide for redundancy in devices, VSD groups, VPNs, and interfaces. Where there are redundant components, failover is the operational mode where the functions of a primary component are automatically assumed by the backup when the primary becomes unavailable.

The specific topics covered are as follows:

- "Device Failover (NSRP)" on page 130
- "VSD Group Failover (NSRP)" on page 131
- "Configuring Object Monitoring for Device or VSD Group Failover" on page 132
  - "Configuring Monitored Objects" on page 134
- "Virtual System Failover" on page 144

DEVICE FAILOVER (NSRP)

When you configure two NetScreen devices in an NSRP cluster, the master device synchronizes all configuration and state information with the backup device so that the backup can assume the master role when necessary. For example, if the master device in a cluster fails, the backup is promoted to master and takes over traffic processing. If the original master is restored to its pre-failure status, it can once again take over traffic processing.

There can be many different conditions that can cause a master device in an NSRP cluster to fail over to the backup. These conditions can include physical problems with the master device itself, such as a system crash, loss of power, down link, or removal of CPU or memory boards in the device. In addition, there are administrator-defined conditions that can cause a master device to fail over to the backup. For example, you can specify that a loss of connection to certain gateways or servers causes a master device to fail over to the backup.

You can configure NSRP to monitor different objects so that the failure of one or more of the monitored objects causes a failover of the master device. See “Configuring Object Monitoring for Device or VSD Group Failover” for more information about these objects and how to configure them.

In the event that there a multiple failovers in a cluster, there should always be at least one device that remains as the master. If a device is the last device in a cluster that has not failed or become ineligible to be master, that device continues to act as master. Under certain conditions, the failure of monitored objects can cause both devices in a cluster to become ineligible, which results in a traffic “black hole”. To ensure that one device is still elected as master and can forward traffic, issue the CLI command **set nsrp vsd-group master-always-exist**. This allows a device in the NSRP cluster to continue to forward traffic even if all units in the cluster are deemed to have failed due to NSRP object monitoring. If all devices in a cluster are simultaneously deemed to be in a failed state, a new master is elected based on the preempt and priority values configured for the devices.

VSD GROUP FAILOVER (NSRP)

In addition to device failover, you can configure NSRP for VSD group failover. Like device failover, failure of one or more monitored objects can cause the master device in a VSD group to fail over to the backup device for the group. See "Configuring Object Monitoring for Device or VSD Group Failover" for information about the objects and how to configure them. For VSD failover, you can configure the same objects to be monitored as device failover.

In the following example, if a port on a master device in a VSD group fails, the entire device does not necessarily fail over to the backup device. In the following configuration, if the interface ethernet 2/1 fails, VSD 1 fails over from the primary VSD group on device A to the backup VSD group on device B. VSD 2 remains active on device A.

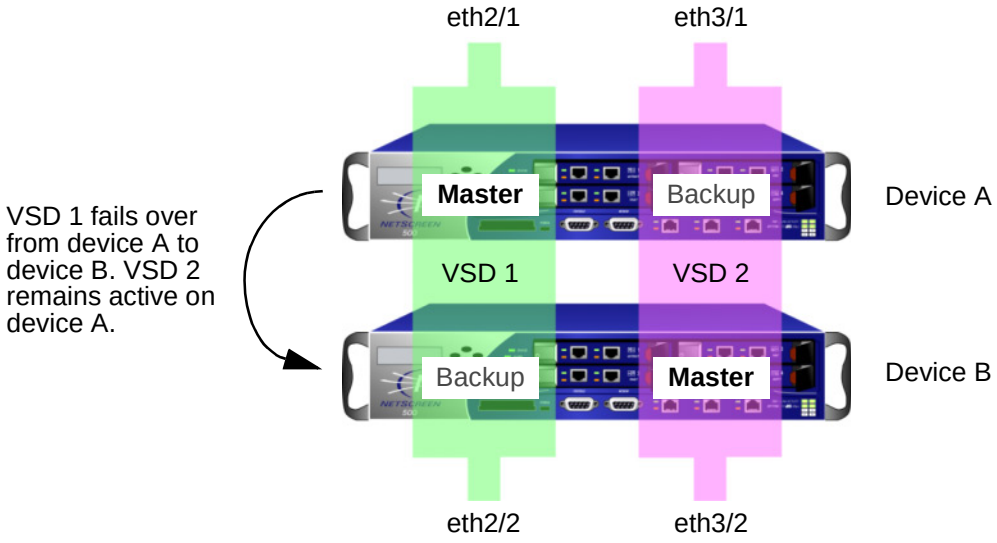

CONFIGURING OBJECT MONITORING FOR DEVICE OR VSD GROUP FAILOVER

With NSRP, you can monitor certain objects to determine failover of the NetScreen device or of a VSD group. NSRP monitored objects can include:

- **Physical interfaces** – The NetScreen device uses NSRP to check that the physical ports are active and connected to other devices.
- **Zones** – The NetScreen device uses NSRP to check that all physical ports in a zone are active.
- **Specific target IP addresses** – The NetScreen device sends ping or ARP requests to up to 16 specified IP addresses at specified intervals and then monitors responses from the targets. All the IP addresses configured on the device or for a specified VSD group constitute a single monitored object.

Configuring device or VSD group failover with monitored objects involves setting the following:

- **Device or VSD failover threshold** – The device or VSD group failover threshold is the total weight of failed monitored objects that is required to cause either VSD group on a device or a device in an NSRP cluster to step down as master. If the cumulative weight of the failures of all monitored objects exceeds the threshold, then the VSD groups or the device fails over to the backup VSD groups or device. You can set the device or VSD failover threshold at any value between 1 and 255. The default threshold is 255.
- **Failure weight of each object being monitored** – Each monitored object has a configurable failure weight, which is the weight that the failure of the monitored object contributes toward the device or VSD failover threshold. You can set the object failure weight at any value between 1 and 255. The default failure weight for monitored objects is 255. If you want to monitor an object but do not want the failure of the object to affect failover of the device or VSD, set the failure weight of the object to 0. ScreenOS logs the failure of any monitored object, even if the failure weight of the object is 0. The following sections describe how to set failure weights for monitored objects.

For tracked IP addresses, you need to specify individual IP addresses and how they are to be monitored. You also need to define what constitutes the failure of each tracked IP address (the threshold) and the weight that the failed IP address carries. For the tracked IP object, you also specify a failure threshold. This threshold is the sum of the weights of all failed tracked IP addresses required for the tracked IP object to be considered failed.

Note that objects that are monitored for a VSD group are independent from the objects monitored for the device. That is, you can configure a specific set of objects, weights, and thresholds for a VSD group and a different set for a device. You can also configure independent sets of monitored objects for different VSD groups. For example, you can configure the same monitored objects for two VSD groups with different weights and thresholds specified for each VSD group for the object.

The following diagram shows the relationship of various monitored objects to the device or VSD group failover. The weights of all failed monitored objects contribute toward the device or VSD failover threshold. If you do not change the default weight of a monitored object or the device or VSD failover threshold, failure of any monitored object can cause the device or VSD to fail over. For tracked IP addresses, the weights of all failed tracked IP addresses contribute toward the tracked IP object failure threshold. If the tracked IP object failure threshold is reached, the tracked IP object failure weight is compared to the device or VSD failover threshold.

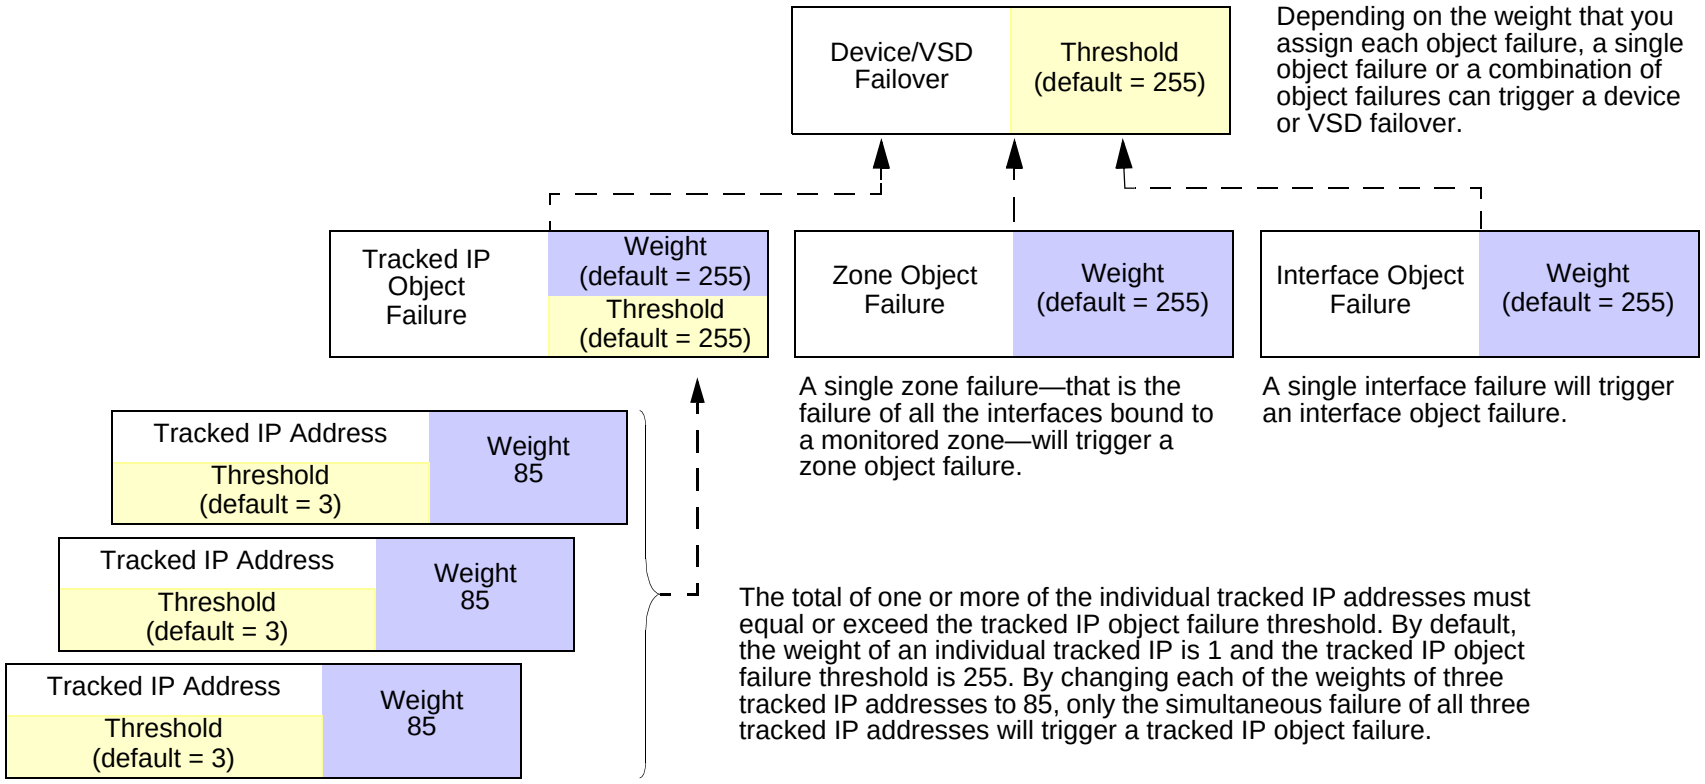

Configuring Monitored Objects

This section describes how to configure monitored objects, including setting failure weights.

Physical Interface Objects

Layer 2 path monitoring functions by checking that the physical ports are active and connected to other network devices. Failure of a physical interface object occurs when the port is no longer active.

Example: Monitoring an Interface

In this example, you enable the monitoring of ethernet2/1 for a possible device failover. You set a failure weight of 100 for the interface.

WebUI

Network > NSRP > Monitor > Interface > VSD ID: Device Edit Interface: Enter the following, and then click **Apply**:

Interface Name: ethernet2/1 (select)

Weight: 100

CLI

```
set nsrp monitor interface ethernet2/1 weight 100
save
```

Zone Objects

Failure of a zone object occurs only when *all* physical interfaces in a monitored zone are down. There is no zone failure as long as there is still an active port in the zone. If a monitored zone has no interfaces bound to it, the zone object does not fail. If a down interface is the only interface bound to a monitored zone, the zone object is failed; if you unbind the interface from the zone, the zone object is no longer failed. If you unbind an active interface from a monitored zone where the remaining interfaces are down, the zone will be considered failed.

Example: Monitoring an Interface

In this example, you enable the monitoring of the Trust zone for a possible device failover. You set a failure weight of 100 for the zone.

WebUI

Network > NSRP > Monitor > Zone > VSD ID: Device Edit Zone: Enter the following, and then click **Apply**:

Zone Name: Trust (select)

Weight: 100

CLI

```
set nsrp monitor zone trust weight 100
save
```

Tracked IP Objects

IP tracking functions by sending ping or ARP requests to up to 16 specified IP addresses at user-determined intervals and then monitoring if the targets respond. When you configure IP tracking, the device sends either ping or ARP requests from a manage IP address that is bound to a physical interface, redundant interface, or subinterface. (The manage IP address must be a different IP address from the IP address of the interface.) Note that you cannot use a VSI for IP tracking because that address can shift its bindings among multiple devices.

Note: *When routers are grouped in a redundant cluster using the Virtual Router Redundancy Protocol (VRRP), the router functioning as the master cannot respond to ping requests to the virtual IP address if it is not the IP address owner (which might be the case after a failover). However, the master virtual router must respond to ARP requests with the virtual MAC address regardless of IP address ownership. (See RFC 2338 for details.) To use ARP when IP tracking, the polled device must be on the same physical subnet as the NetScreen manage IP address.*

For each tracked IP address, you specify the following:

- **Tracked IP Failure Threshold** – This is the number of consecutive failures to elicit a ping or ARP response from a specific IP address that constitutes a failed attempt. Not exceeding the threshold indicates an acceptable level of connectivity with the address; exceeding it indicates an unacceptable level. You can set the threshold to any value between 1-200. The default value is 3.
- **Tracked IP Failure Weight** – This is the weight that failure to elicit a response from the tracked IP address contributes to the tracked IP object failure weight. By applying a weight to a tracked IP address, you can adjust the importance of connectivity to that address in relation to reaching other tracked IP addresses. You can assign greater weights to relatively more important addresses, and lesser weights to relatively less important addresses. The assigned weights come into play when a tracked IP failure threshold is reached. For example, exceeding the tracked IP failure threshold for an address weighted 10 adds more to the tracked IP object failure weight than would a tracked IP failure for an address weighted 1. You can assign weights from 1 to 255. The default is 1.

You also configure a failure threshold for the tracked IP object which contributes to the device or VSD failover threshold. If one or more tracked IP addresses exceed their failure thresholds, then the weights for the individual failed addresses are totaled. If the sum reaches or exceeds the failure threshold for the tracked IP object, then the tracked IP object failure weight is applied to the device or VSD failover threshold. Note that only the failure weight of the tracked IP object is applied to the device or VSD failover threshold; failure weights of individual tracked IP addresses are never applied to the device or VSD failover threshold. Consider the following example:

Tracked IP Addresses	Failure Weights			
10.10.10.250	100	Tracked IP Object Failure Threshold	Tracked IP Object Failure Weight	Device Failover Threshold
1.1.1.30	75	125	255	255
2.2.2.40	75			

If the tracked IP address 10.10.10.250 fails, then the tracked IP failure weight (100) is compared to the tracked IP object failure threshold (125). Since the tracked IP failure weight is less than the tracked IP object failure threshold, the tracked IP object is not considered failed. If both tracked IP addresses 1.1.1.30 and 2.2.2.40 fail, then the combined failure weight (150) is compared to the tracked IP object failure threshold (125). Since the combined failure weight exceeds the tracked IP object failure weight, the tracked IP object is considered failed. The tracked IP object failure weight (255) is compared to the device failover threshold (255). Since the tracked IP object failure weight equals the device failover threshold, the device fails over.

To set a failure weight of 100 for the tracked IP address 10.10.10.250, enter the following:

WebUI

Network > NSRP > Track IP > New: Enter the following, and then click **OK**:

Track IP: 10.10.10.250

Weight: 100

CLI

```
set nsrp track-ip ip 10.10.10.250 weight 100
save
```

To set a failure threshold of 125 for the tracked IP object for a possible device failover, enter the following:

WebUI

Network > NSRP > Monitor > Track IP > VSD ID: Device Edit: Enter the following, and then click **Apply**:

Enable Track IP: (select)

Failover Threshold: 125

CLI

```
set nsrp monitor track-ip threshold 125
save
```

Example: Track IP for Device Failover

Two NetScreen devices are in an active/active configuration. Every 10 seconds, both devices send ARP requests to the physical IP addresses[1] of two external routers running VRRP in a redundant cluster in the Untrust zone and ping requests to two Web servers in the Trust zone. The tracked IP object failure threshold is 51. The tracked IP object weight and the device failover threshold are the default values (255). The weights and failure thresholds of the tracked IP addresses are as follows:

- Redundant routers in the Untrust zone
  - 210.1.1.250 – Weight: 16, threshold 5
  - 210.1.1.251 – Weight: 16, threshold 5
- Web servers in the Trust zone
  - 10.1.1.30 – Weight 10, threshold 3
  - 10.1.1.40 – Weight 10, threshold 3

Not receiving an ARP response after 5 consecutive attempts to one of the routers is considered a failed attempt and contributes a weighted value of 16 toward the total failover threshold. Not receiving a ping response after 3 consecutive attempts to one of the Web servers is considered a failed attempt and contributes a weighted value of 10 toward the total failover threshold.

Because the device failover threshold is 51, all four tracked IP addresses must fail before a device failover occurs. If you are not willing to tolerate that amount of failure, you can lower the threshold to a more acceptable level.

In this example, device A has a 100% success rate, while device B has failed to receive three consecutive responses from 10.1.1.40, contributing a value of 10 toward the total failover threshold of 51.

Note: *All NSRP monitoring settings apply to the local unit only. The IP tracking settings do not propagate to other devices in a VSD group. You must enter the same settings on all devices in the group if necessary.*

The Untrust zone interface is ethernet1 and the Trust zone interface is ethernet2 on both devices. The ethernet1 manage IP address is 210.1.1.1 on device A, and 210.1.1.2 on device B. The ethernet2 manage IP address is 10.1.1.1 on device A, and 10.1.1.2 on device B. All the security zones are in the trust-vr routing domain.

1. The physical IP addresses are the addresses dedicated to the physical routers that comprise the VRRP cluster.

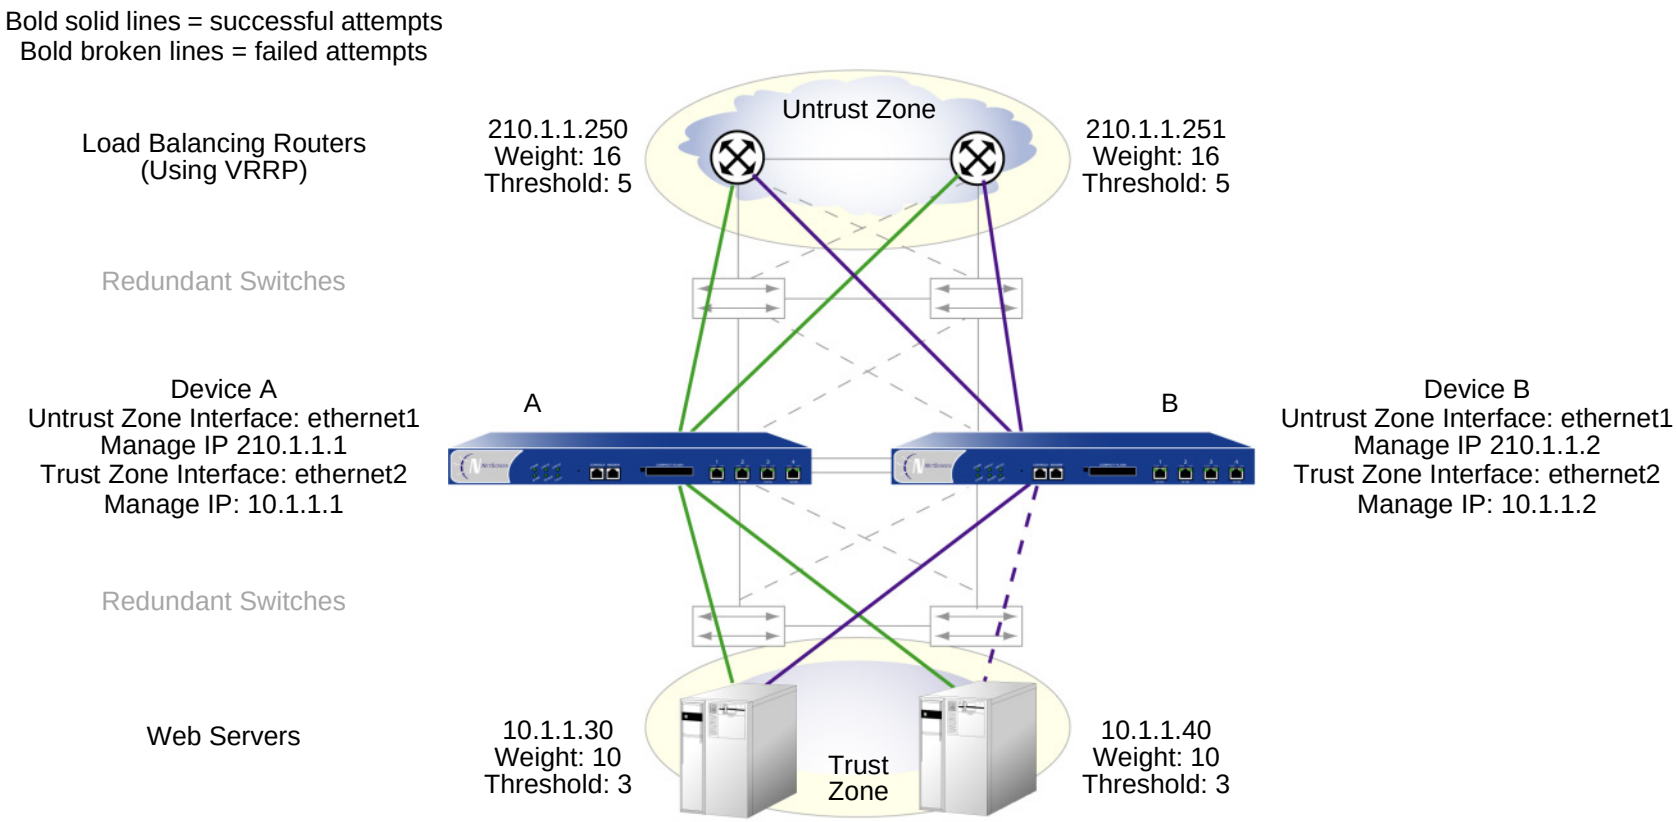

WebUI

1. Track IP Addresses

Network > NSRP > Monitor > Track IP > New: Enter the following, and then click **OK**:

Track IP: 210.1.1.250

Method: ARP

Weight: 16

Interval (sec): 10

Threshold: 5

Interface: ethernet1

VSD Group ID: Device

Network > NSRP > Monitor > Track IP > New: Enter the following, and then click **OK**:

Track IP: 210.1.1.251

Method: ARP

Weight: 16

Interval (sec): 10

Threshold: 5

Interface: ethernet1

VSD Group ID: Device

Network > NSRP > Monitor > Track IP > New: Enter the following, and then click **OK**:

Track IP: 10.1.1.30

Method: Ping

Weight: 10

Interval (sec): 10

Threshold: 3

Interface: ethernet2

VSD Group ID: Device

Network > NSRP > Monitor > Track IP > New: Enter the following, and then click **OK**:

Track IP: 10.1.1.40

Method: Ping

Weight: 10

Interval (sec): 10

Threshold: 3

Interface: ethernet2

VSD Group ID: Device

2. Track IP Object Failure Threshold

Network > NSRP > Monitor > Track IP > Edit (for VSD: Device): Enter the following, and click **Apply**:

Enable Track IP: (select)

Failover Threshold: 51

CLI

1. Track IP Addresses

```
set nsrp track-ip ip 210.1.1.250 interface ethernet1
set nsrp track-ip ip 210.1.1.250 interval 10
set nsrp track-ip ip 210.1.1.250 method arp
set nsrp track-ip ip 210.1.1.250 threshold 5
set nsrp track-ip ip 210.1.1.250 weight 16
set nsrp track-ip ip 210.1.1.251 interface ethernet1
set nsrp track-ip ip 210.1.1.251 interval 10
set nsrp track-ip ip 210.1.1.251 method arp
set nsrp track-ip ip 210.1.1.251 threshold 5
set nsrp track-ip ip 210.1.1.251 weight 16
set nsrp track-ip ip 10.1.1.30 interface ethernet2
set nsrp track-ip ip 10.1.1.30 interval 10
```

```
set nsrp track-ip ip 10.1.1.30 method ping[2]
set nsrp track-ip ip 10.1.1.30 threshold 3
set nsrp track-ip ip 10.1.1.30 weight 10
set nsrp track-ip ip 10.1.1.40 interface ethernet2
set nsrp track-ip ip 10.1.1.40 interval 10
set nsrp track-ip ip 10.1.1.40 method ping
set nsrp track-ip ip 10.1.1.40 threshold 3
set nsrp track-ip ip 10.1.1.40 weight 10
set nsrp track-ip
```

2. Track IP Object Failure Threshold

```
set nsrp track-ip threshold 51
save
```

2. By default, pinging is the method for IP tracking and a tracked IP failure threshold value is 3; therefore, you do not need to specify them. The commands **set nsrp track-ip ip 10.1.1.30** and **set nsrp track-ip ip 10.1.1.40** are sufficient.

VIRTUAL SYSTEM FAILOVER

For a virtual system to fail over, it must be in a VSD group. For a VSD group to support virtual systems, you must create VSIs for each virtual system. A virtual system has its own Trust zone VSI, and it can have its own Untrust zone VSI. A virtual system can also share the Untrust zone VSI with the root level. When virtual systems have their own Untrust zone VSIs, they must be in different subnets from each other and from the Untrust zone VSI at the root level. All Trust zone virtual system VSIs must also be in different subnets from one another.

Example: VSIs for Inter-Virtual System Load Sharing

Two NetScreen devices (device A and device B) are in an active/active full-mesh configuration. You have already configured the root system of device A as the master of VSD 0 and that of device B as the master of VSD group 1. The Trust and Untrust zone VSIs for VSDs 0 and 1 in the root system are as follows:

VSIs for VSD Group 0		VSIs for VSD Group 1	
redundant1	210.1.1.1/24	redundant1:1	210.1.1.2/24
redundant2	10.1.1.1/24	redundant2:1	10.1.1.2/24

(For the complete configuration of the root system VSD groups, see "Example: NSRP for an Active/Active Configuration" on page 49.)

In this example, you configure two virtual systems (vsys1 and vsys2) for NSRP. To provide load sharing[3] for incoming traffic to the virtual systems, VSD membership is apportioned as follows:

- Vsys1 is a member of VSD group 0.
- Vsys2 is a member of VSD group 1.

The NetScreen devices share the incoming traffic load by distributing the VSD groupings of the virtual systems. Because of the initial design of configuring vsys1 on device A and vsys2 on device B, incoming traffic to these virtual systems is directed to the device that contains it.

3. Note that in this example the load is not evenly distributed; that is, it is not load balanced. The two NetScreen devices share the load, with devices A and B receiving incoming traffic in dynamically shifting proportions (60/40%, 70/30%, and so on).

The root system is in VSD groups 0 and 1, and is active in both NetScreen devices.

Vsys1 is in VSD group 0, and is active only in device A.

Vsys2 is in VSD group 1, and is active only in device B.

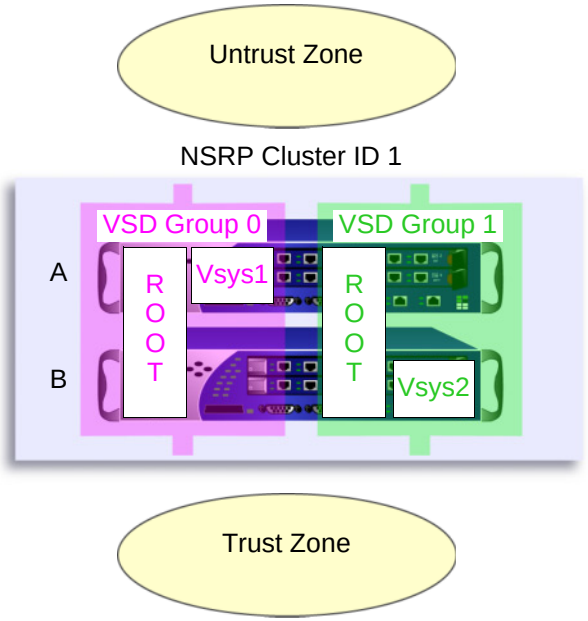

The default gateway for outbound traffic is different for the root system and each virtual system:

- Root: 210.1.1.250
- Vsys1: 210.11.1.250
- Vsys2: 210.12.1.250

Because this example builds on “Example: NSRP for an Active/Active Configuration” on page 49, in which you set up VSD groups 0 and 1 and set the devices in NSRP cluster ID 1, NSRP is already enabled. Therefore, the settings you configure on device A automatically propagate to device B.

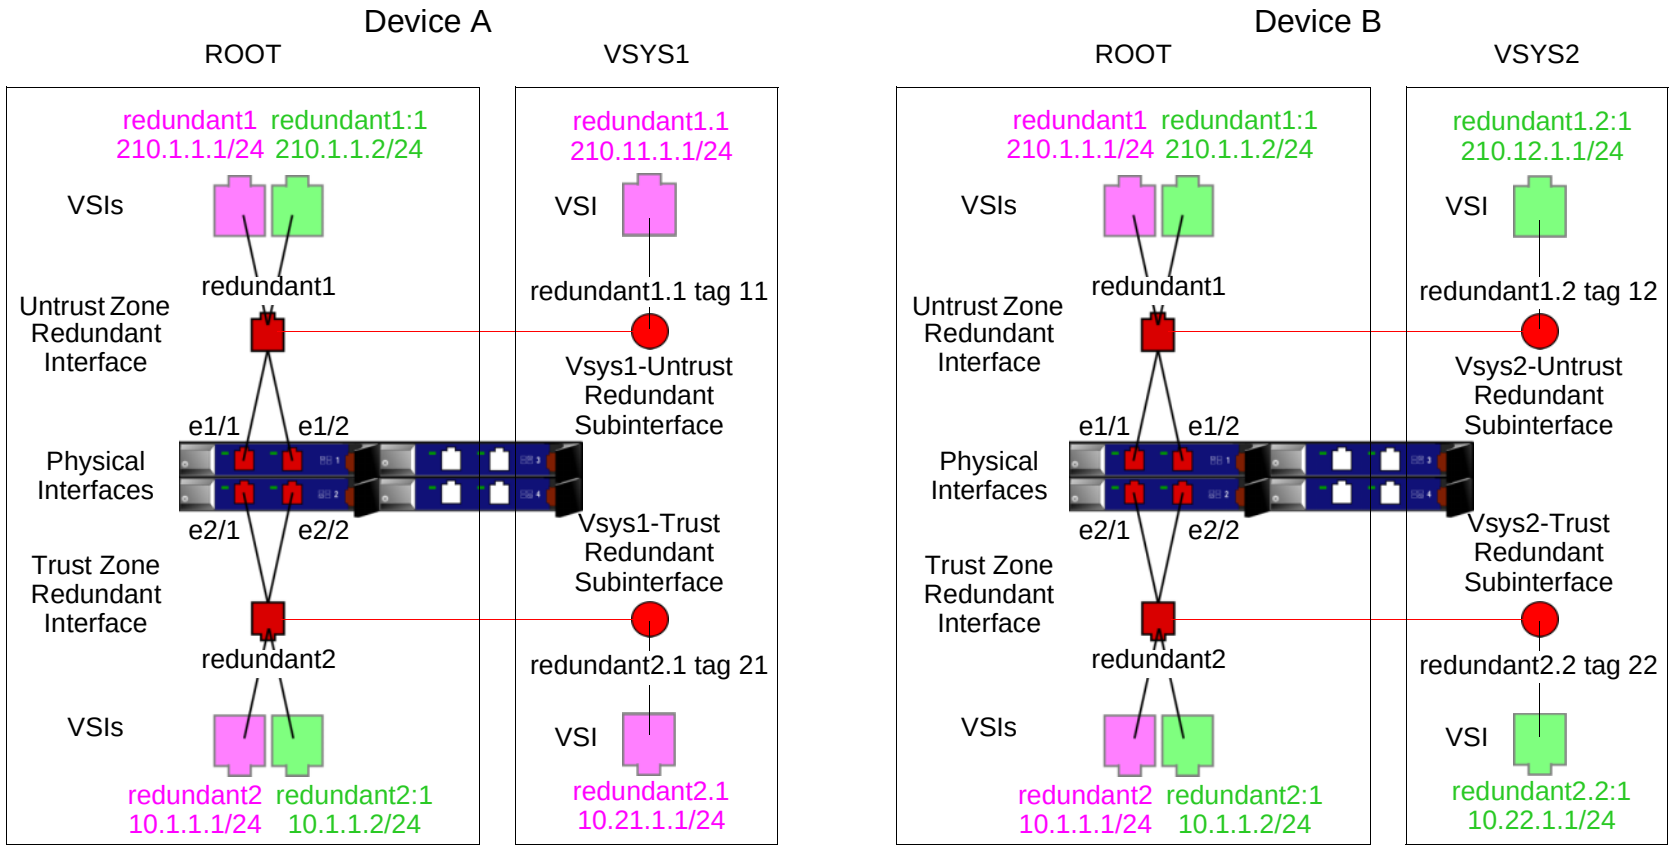

VSD Group 0 = Magenta (Note: The VSIs for VSD 0 do not display their VSD ID number.)
VSD Group 1 = Green (Note: The VSIs for VSD 1 indicate their VSD ID by colon+1.)

WebUI

1. Device A: Root

Note: *The NSRP configuration for the root system is identical to that in* "Example: NSRP for an Active/Active Configuration" on page 49*.*

2. Device A: Vsys1

Vsys > New: Enter the following, and then click **OK**:

VSYS Name: vsys1[4]

Vsys > Enter (vsys1) > Network > Interface > New Sub-IF: Enter the following, and then click **OK**:

Interface Name: Redundant1.1

Zone Name: Untrust

VLAN Tag: 11

Network > Interfaces > New VSI IF: Enter the following, and then click **OK**:

VSI Base: Redundant1.1

VSD Group: 0

IP Address / Netmask: 210.11.1.1/24

Network > Interfaces > New Sub-IF: Enter the following, and then click **OK**:

Interface Name: Redundant2.1

Zone Name: Trust-vsys-vsys1

VLAN Tag: 21

Network > Interfaces > New VSI IF: Enter the following, and then click **OK**:

VSD Group ID: 0

IP Address / Netmask: 10.21.1.1/24

Interface Mode: Route[5]

4. If you do not define a vsys admin, the NetScreen device automatically creates one by appending “vsys_” to the vsys name. In this example, the vsys admin for vsys1 is vsys_vsys1.

5. Virtual systems can be in either Route or NAT mode, independent of the mode you set at the root level.

Network > Routing > Routing Entries > untrust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: Redundant1

Gateway IP Address: 210.11.1.250

Click **Exit Vsys** to return to the root level.

3. Device A: Vsys2

Vsys > New: Enter the following, and then click **OK**:

VSYS Name: vsys2

Vsys > Enter (vsys2) > Network > Interface > New Sub-IF: Enter the following, and then click **OK**:

Interface Name: Redundant1.2

Zone Name: Untrust

VLAN Tag: 12

Network > Interfaces > New VSI IF: Enter the following, and then click **OK**:

VSI Base: Redundant1.2

VSD Group: 1

IP Address / Netmask: 210.12.1.1

Network > Interfaces > New Sub-IF: Enter the following, and then click **OK**:

Interface Name: Redundant2.2

Zone Name: Trust-vsys-vsys2

VLAN Tag: 22

Network > Interfaces > New VSI IF: Enter the following, and then click **OK**:

VSD Group ID: 1

IP Address / Netmask: 10.22.1.1/24

Interface Mode: Route

Network > Routing > Routing Entries > untrust-vr New: Enter the following, and then click **OK**:

Network Address/Netmask: 0.0.0.0/0

Gateway: (select)

Interface: Redundant1

Gateway IP Address: 210.12.1.250

Click **Exit Vsys** to return to the root level.

4. Device B

Note: *Because device A propagates the other configuration settings to device B, you do not need to enter them again in device B.*

CLI

1. Device A: Root

Note: *The NSRP configuration for the root system is identical to that in* “Example: NSRP for an Active/Active Configuration” on page 49*.*

2. Device A: VSYS 1

```
set vsys vsys1
ns(vsys1)-> set interface redundant1.1 tag 11 zone untrust
ns(vsys1)-> set interface redundant1.1 ip 210.11.1.1/24
ns(vsys1)-> set interface redundant2.1 tag 21 zone trust-vsys1
ns(vsys1)-> set interface redundant2.1 ip 10.21.1.1/24
```

```
ns(vsys1)-> set interface redundant2.1 route[6]
ns(vsys1)-> set vrouter untrust-vr route 0.0.0.0/0 interface redundant1 gateway
    210.11.1.250
ns(vsys1)-> save
ns(vsys1)-> exit
```

3. Device A: VSYS 2

```
set vsys vsys2
ns(vsys2)-> set interface redundant1.2 tag 12 zone untrust
ns(vsys2)-> set interface redundant1.2:1 ip 210.12.1.1/24
ns(vsys2)-> set interface redundant2.2 tag 22 zone trust-vsys2
ns(vsys2)-> set interface redundant2.2:1 ip 10.22.1.1/24
ns(vsys2)-> set interface redundant2.2:1 route
ns(vsys2)-> set vrouter untrust-vr route 0.0.0.0/0 interface redundant1 gateway
    210.12.1.250
ns(vsys2)-> save
ns(vsys2)-> exit
```

4. Device B

Note: *Because device A propagates the other configuration settings to device B, you do not need to enter them again in device B.*

6. Virtual systems can be in either Route or NAT mode, independent of the mode you set at the root level.

Index

A

aggregate interfaces 65
ARP 56, 100
  broadcasts 18
  path monitoring 137
authentication
  NSRP 7, 18
  NSRP-Lite 122

C

character types, ScreenOS supported x
CLI
  conventions vi
  set arp always-on-dest 56
cluster name, NSRP 17, 121
clusters 16–20, 49, 118–121
control messages 38
  HA messages 40
  HA physical link heartbeats 39
  RTO heartbeats 40
  VSD heartbeats 40
conventions
  CLI vi
  illustration ix
  names x
  WebUI vii

D

data messages 40
device failover 94
dual Untrust interfaces 67

E

encryption
  NSRP 7, 18
  NSRP-Lite 122

F

failover
  device 94
  dual Untrust interfaces 68, 69
  object monitor 96
  serial interface 87
  virtual system 108
  VSD group 95
full-mesh configuration 108

H

HA
  active/active failover 6
  active/passive failover 4
  aggregate interfaces 65
  cabling 45–48
  cabling for dedicated HA interfaces 45
  cabling network interfaces as HA links 47
  control link 38
  data link 41
  dual Untrust interfaces 67
  HA LED 25
  IP tracking 100, 137
  link probes 42
  messages 40
  path monitoring 137
  redundant interfaces 58
  secondary path 25
  serial interface 82
High Availability
  See HA

I

illustration
  conventions ix
interfaces
  aggregate 65
  dual Untrust 67
  HA, dual 38–41
  monitoring 18
  redundant 58
  serial 82
  Virtual HA 47
  VSIs 28
IP tracking 100, 137
  device failover threshold 138
  ping and ARP 100, 137
  tracked IP failure threshold 97, 138
  tunnel failover 139
  weights 138
ISP configuration for serial interface 85

L

LED indicators, HA 25
load sharing 108

M

manage IP
  VSD group 0 8
modem configuration for serial interface 83

N

names
  conventions x
NetScreen Redundancy Protocol
  See NSRP
NetScreen Reliable Transport Protocol
  See NRTP
NRTP 33, 134
NSRP
  ARP 56
  ARP broadcasts 18
  backup 4
  cabling 45–48
  clear cluster command 16, 121
  cluster name 17, 121
  clusters 16–20, 49
  config sync 33

control link 38
control messages 38, 39
data link 41
data messages 40
debug cluster command 16, 121
default settings 9, 119
files, sync 34
full-mesh configuration 45, 108
HA cabling, dedicated interfaces 45
HA cabling, network interfaces 47
HA interfaces 39
HA LED 25
HA ports, redundant interfaces 58
HA session backup 21
hold-down time 51, 55
interface monitoring 18
load sharing 108
manage IP 100, 138
master 4
NAT and Route modes 8
NTP synchronization 37
overview 3
packet forwarding and dynamic routing 41
port failover 58
port monitoring 98
preempt mode 23
priority numbers 23
redundant ports 38
RTO states 22
RTOs 21–22, 49
RTOs, sync 34
secondary path 18, 25
secure communications 7, 18
synchronization, PKI 34
Transparent mode 8
virtual systems 108–114
VSD groups 5, 23–27, 49, 137
VSIs 5
VSIs, static routes 28, 63, 64
NSRP-Lite 115–136
  cabling 126
  clusters 118–121
  config synchronization 134
  disabling synchronization 136
  file synchronization 135
  port monitoring 137
  preempt mode 125
  secure communications 122
  VSD groups 123–125
NTP
  NSRP synchronization 37

O

object monitoring 96

P

path monitoring 137
  tunnel failover 139
ports
  monitoring 98, 137
  port failover 58
  primary trusted and untrusted 58
  redundant 38
  secondary trusted and untrusted 58
preempt mode 23, 125
protocols
  NRTP 33, 134
  NSRP 1, 115
  VRRP 100, 137

R

RTOs 21–22
  operational states 22
  RTO peer 24
run-time objects
  See RTOs

S

secondary path 18, 25
serial interface 82
  failover 87
  ISP configuration 85
  modem configuration 83
synchronization
  configuration 33
  files 34
  PKI objects 34
  RTOs 34

V

Virtual HA interface 47
virtual security device groups
  See VSD groups
virtual security interface
  See VSI
virtual system
  failover 108
  load sharing 108
  NSRP 108
VRRP 100, 137
VSD groups 5, 23–27, 123–125
  failover 95
  heartbeats 18, 25, 124
  hold-down time 51, 55
  member states 24, 123–124, 137
  priority numbers 23
VSIs 5, 23, 123
  multiple VSIs per VSD group 108
  static routes 28

W

WebUI
  conventions vii